Diss. ETH No. 23907

Low-power Wireless Systems Coexistence

A thesis submitted to attain the degree of

DOCTOR OF SCIENCES of ETH ZURICH

(Dr. sc. ETH Zurich)

presented by

ANWAR HITHNAWI
M.Sc. in Software Systems Engineering, RWTH Aachen

accepted on the recommendation of
Prof. Dr. Friedemann Mattern, examiner
Dr. Simon Duquennoy, co-examiner
Prof. Dr. James Gross, co-examiner
Prof. Dr. Lothar Thiele, co-examiner

2016

Abstract

The convergence of networked embedded devices, wearables, and sensing technologies has expedited the emergence of an array of innovative services and applications that are radically changing the way we perceive and interact with the physical world. Wireless communication is the natural substrate connectivity means for a wide variety of these applications. For these applications to perform correctly, they require the underlying wireless communication to be reliable and energy-efficient. Meeting these requirements is, however, challenging. Particularly, as we witness an unprecedented demand for wireless access, more wireless technologies and devices need to share the scarcely available radio spectrum. This is especially a growing problem for devices operating in the unlicensed spectrum, where the density and heterogeneity of radios operating in this spectrum are surging. Consequently, interference between the heterogeneous radio systems is growing in unpredictable ways. The emerging spectrum crunch necessitates the design and development of innovative wireless systems that enhance spectrum utilization and are apprehensive of the uncoordinated wireless coexistence problem.

In this dissertation, we take an alternative approach to deal with Cross-Technology Interference (CTI). Instead of avoiding interference, we adopt an interdisciplinary approach combining a cross-layer design and machine learning techniques to build cognitive low-power wireless systems that can cope with Cross-Technology Interference. We begin this dissertation by acquiring a good understanding of how various interfering wireless signals interact, and we harness this understanding in our designs. We then introduce a family of algorithms and system architectures that improve the robustness of low-power wireless networks operating in interference-rich environments. The introduced systems embody a cross-layer design and a cognitive engine that radios can exploit to intelligently share the spectrum and implement CTI-aware mitigation schemes. In particular, we present three novel systems contributing to low-power wireless systems coexistence:

i) Technology-Independent Interference Mitigation (TIIM): Interfering radio technologies differ widely in the way they affect wireless links. Cross-Technology Interference has a complex impact on wireless links, which needs to be taken into account when treating interference.

To address this challenge, we present **TIIM**, a system that identifies, quantifies, and reacts to CTI in real-time. In the design of **TIIM**, we follow an unconventional approach, where we employ lightweight machine learning techniques to assist wireless nodes in recovering from interference. Within **TIIM**, we develop a lightweight classifier which is trained to select a coexistence solution that works most effectively for the current channel fingerprint.

ii) CrossZig: Current wireless designs still largely impose layer isolation. Thereby, conventional approaches to tackle wireless performance have focused on separately optimizing different layers of the networking stack. This rigid design fails to harness the rich ambient information embedded in the physical signals. Hence, reliability solutions targeting layers in isolation are typically suboptimal. In recent years, cross-layer optimizations were profoundly advocated in the wireless community. In this work, we pursue this research direction. We show how physical layer information and primitives can be coupled with the link layer to enhance low-power wireless systems coexistence and performance under interference. Notably, we present **CrossZig**, a cross-layer wireless design, that enables low-power wireless networks to exploit fine-grained physical layer information to make informed decisions that can help them recover from varying sources of interference. **CrossZig** utilizes physical layer information to detect the presence of external interference in corrupted packets and to apply an adaptive packet recovery which incorporates a novel cross-layer based packet merging scheme and an adaptive channel coding.

iii) Controlled Interference Generator (CIG): Wireless research testbed infrastructures often lack proper tools for enabling repeatable replay of realistic radio interference commonly found in real-world deployments. Hence, benchmarking wireless coexistence solutions is often cumbersome, time-consuming, and even infeasible in remote testbeds. To facilitate Cross-Technology Interference and wireless coexistence experimentations, we develop **CIG**, a framework that extends wireless testbed infrastructures with the capability of reproducing heterogeneous external interference. In the design of **CIG**, we consider a unified approach that incorporates a careful selection of interferer technologies (implemented in software), to expose networks to realistic interference patterns.

The systems presented in this dissertation demonstrate that incorporating cognitive and cross-layer wireless designs is adequate to mitigate the problem of uncoordinated wireless coexistence.

Zusammenfassung

Die Konvergenz vernetzter eingebetteter Geräte, Wearables sowie Sensortechnologien hat eine Reihe innovativer Dienste und Anwendungen ermöglicht, welche die Art und Weise, wie wir mit der physischen Welt interagieren und diese wahrnehmen, radikal verändert haben. Für die Mehrzahl dieser Anwendungen stellt dabei die Funktechnologie das natürliche Kommunikationsmedium dar. Damit derartige Anwendungen einwandfrei funktionieren, muss die zugrunderliegende drahtlose Kommunikationstechnologie daher zuverlässig und energieeffizient sein. Dies stellt eine Herausforderung dar, da der Bedarf an drahtloser Kommunikationsmöglichkeit ständig steigt, neue Funktechnologien etabliert werden und sich immer mehr Geräte das ohnehin spärlich verfügbare Frequenzband teilen müssen. Dies stellt insbesondere ein zunehmendes Problem für solche Systeme dar, die in den lizenzfreien Frequenzbändern operieren, wo die Dichte und Heterogenität der Geräte schnell anwächst. Als Folge davon steigt das Interferenzpotential zwischen heterogenen drahtlosen Systemen stark an. Die zunehmende Ressourcenknappheit erfordert daher innovative Konzepte, welche eine verbesserte Frequenzbandnutzung ermöglichen und sich dediziert des Problems der unkoordinierten drahtlosen Koexistenz annehmen.

In diesem Sinne verfolgen wir in der vorliegenden Dissertation einen neuartigen Ansatz zur Behandlung der sogenannten technologieübergreifenden Interferenz (Cross-Technology Interference, CTI). Anstatt auf Interferenzvermeidung zu setzen, kombinieren wir ein schichtübergreifendes Konzept mit Ansätzen des maschinellen Lernens, um energieeffiziente kognitive Funksysteme realisieren zu können, die mit CTI gut zurechtkommen. Zu Beginn der Dissertation wird unser erzieltes Verständnis dazu, wie interferenzverursachende Signale wirken, ausführlich dargelegt. Die gewonnenen Erkenntnisse machen wir uns sodann in unseren Systementwürfen zunutze. Wir stellen dazu eine Familie von Algorithmen und Systemarchitekturen vor, welche die Robustheit energieeffizienter Funkkommunikation in interferenzreichen Umgebungen steigert. Die vorgestellten Systeme realisieren ein schichtübergreifendes Konzept und verkörpern eine „cognitive engine“, die von funkbasierten Geräten auf intelligente Weise genutzt werden kann, um Frequenzbänder untereinander zu teilen und Massnahmen gegen CTI zu treffen. Insbesondere stellen wir drei neue Konzepte vor, welche zur

Koexistenz von energieeffizienten drahtlosen Systemen beitragen:

i) Technologieunabhängige Interferenz-Gegenmassnahmen (Technology Independent Interference Mitigation, TIIM): Interferenzverursachende Drahtlostechnologien unterscheiden sich deutlich in der Art und Weise, wie sie funkbasierte Kommunikationsverbindungen beeinflussen. Technologieübergreifende Interferenz wirkt in komplexer Weise auf drahtlose Verbindungen ein, was bei Behandlung von Interferenz berücksichtigt werden muss. Um dieser Herausforderung zu begegnen, entwickelten wir **TIIM**, ein System, welches CTI in Echtzeit identifiziert, quantifiziert und darauf reagiert. Beim Entwurf von **TIIM** verfolgen wir einen unkonventionellen Ansatz, indem wir leichtgewichtige Mechanismen des maschinellen Lernens einsetzen, um die jeweiligen Knoten beim Beherrschen von Interferenz zu unterstützen. Hierfür entwickelten wir einen leichtgewichtigen Klassifikator, welcher daraufhin trainiert ist, Kanalzustände zu erkennen, bei denen eine bestimmte Koexistenzmassnahme am effektivsten wirkt.

ii) CrossZig: Existierende Konzepte der Drahtloskommunikation unterliegen weitgehend dem Prinzip der Schichtenisolation. Daher haben herkömmliche Ansätze zur Effizienzsteigerung den Fokus auf die getrennte Optimierung der jeweiligen Schicht des Netzwerkstacks gelegt. Diese starre Vorgehensweise kann jedoch nicht das reichhaltige Kontextwissen, welches in den Funksignalen enthalten ist, nutzen. Daher sind die nach diesem Entwurfsmodell entwickelten Lösungen zur Steigerung der Zuverlässigkeit in der Regel suboptimal. In den vergangenen Jahren sind jedoch schichtübergreifende Optimierungen bei der Drahtloskommunikation stark propagiert worden. In der vorliegenden Dissertation greifen wir diesen Ansatz auf. Wir zeigen, wie die Informationen und Grundelemente aus der Bitübertragungsschicht, wenn sie mit der Sicherungsschicht geeignet verbunden werden, zu einer verbesserten Koexistenz und höheren Leistung energieeffizienter drahtloser Systeme bei Vorliegen von Interferenz führen können. Dazu diskutieren wir **CrossZig**, ein schichtübergreifendes Entwurfskonzept, das es energieeffizienten drahtlosen Netzen ermöglicht, feingranulare Informationen aus der Bitübertragungsschicht zu nutzen, um fundierte Entscheidungen zur Beherrschung von Interferenz unterschiedlicher Provenienz zu treffen. **CrossZig** nutzt bei fehlerbehafteten Paketen Information aus der Bitübertragungsschicht, um das Vorhandensein externer Interferenz zu erkennen und um einen adaptiven Paketwiederherstellungsmechanismus anzuwenden, welcher einen neuartigen schichtübergreifenden Paketverschmelzungsansatz und eine adaptive Kanalkodierung umfasst.

iii) Steuerbarer Interferenzgenerator (Controlled Interference Generator, CIG): Bei den Testinfrastrukturen im Rahmen der Forschung an drahtlosen Systemen fehlen meistens geeignete Werkzeuge, die Interferenzen so realistisch reproduzieren können, wie sie in der realen Welt anzutreffen sind. Daher ist das Benchmarking von Lösungen zur Koexistenz drahtloser Systeme oft mühsam, zeitaufwendig und auch nicht in abgesetzten Testumgebungen durchführbar. Um Experimente zur Koexistenz bei CTI zu ermöglichen, entwickelten wir daher das **CIG**-Framework, welches Testumgebungen für drahtlose Systeme um die Möglichkeit der Reproduktion heterogener externer Interferenz erweitert. Bei der Ausgestaltung von **CIG** wurde ein einheitlicher Ansatz verfolgt, der eine sorgsame Auswahl (softwaremässig implementierter) interferenzverursachender Technologien umfasst, damit drahtlose Netze realistischen Interferenzen ausgesetzt werden können.

Die in dieser Dissertation vorgestellten Systeme zeigen, dass die Kombination von kognitiven und schichtübergreifenden Konzepten geeignet ist, das Problem der unkoordinierten drahtlosen Koexistenz wesentlich zu entschärfen.

Contents

List of Figures

List of Tables

1

Introduction

The field of wireless communication has been established over more than a hundred years ago, dating back to around 1897 with Marconi's successful demonstration of the first commercial wireless telegraphy system based on radio waves. The topic has been extensively studied since then. In the last decade, however, radio-based wireless communication has experienced an unprecedented surge in research activities and phenomenal growth. This is attributed to a confluence of several factors. *(i)* The substantial increase in demand for ubiquitous wireless connectivity, driven by affordable wireless-enabled mobile computers, and wireless handheld devices. *(ii)* The pronounced progress in chip design and software radios which has enabled efficient implementation of complex signal processing algorithms and realization of theoretical work that was not feasible before. *(iii)* The concomitant maturation of mobile ecosystems such as Windows Mobile, Apple's iOS and Google's Android systems leading to increased interest in high throughput wireless. *(iv)* Finally, the increasingly growing interest in the Internet of Things has boosted the development of ultra-low-power wireless technologies. However, there is a fundamental factor that is throttling this growth; wireless transmissions are inherently broadcast by nature and the radio spectrum is fundamentally a shared and scarce resource. If the wireless growth continues at this pace, wireless demand soon will exceed the capacity of allocated radio spectrum. This can lead to serious consequences on the dependently and performance of wireless networks.

In this dissertation, we primarily focus on the low-power wireless systems coexistence in the unlicensed spectrum. Low-power wireless

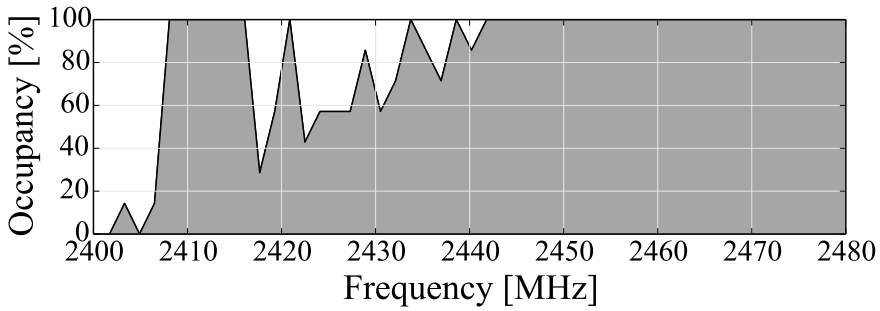

Figure 1.1: Averaged channel occupancy in the 2.4 GHz band over one week (26.-31. August 2013). Data from Microsoft Spectrum Observatory in an enterprise building in Brussels, Belgium.

communication has been a key enabling technology for a class of innovative applications in the last decade. It is the primary choice of connectivity for embedded low-power devices. These devices are increasingly integrated into objects and environments surrounding us, paving the way for the Internet of Thing's vision of digitizing the physical world. These devices are utilized in a range of performance-sensitive applications, such as health systems, general monitoring and tracking, home automation, etc. While it is not expected that exchanged data will be at high volume in these applications, a myriad of devices will be connected at the same time which will create a set of new challenges. Low-power wireless technologies (e.g., Bluetooth Low Energy, IEEE 802.15.4, and backscatter communication) employed by these applications are expected to endure interference from other radio technologies. The Cross-Technology Interference (CTI) problem is exacerbated for these low-power networks, where energy and complexity constraints prohibit the use of sophisticated interference suppression and cancellation techniques that are finding their ways into unconstrained wireless systems. To date, much of the devised radio frequency interference solutions focused on resolving interference between devices of the same technology. As a consequence, current wireless systems are short of mechanisms to identify and adapt to dynamic sources of external interference. Utilizing non-overlapping segments of the spectrum has been the natural solution to avoid/tackle interference between different technologies. However, as the density of radio devices continue to increase, this solution will no longer suffice (see Figure 1.1).

This dissertation presents novel system designs and mechanisms that enable low-power wireless systems to gracefully coexist in the crowded spectrum and mitigate the effects of CTI. The introduced systems embody a cross-layer design and a cognitive engine that radios can harness to intelligently share the spectrum and implement CTI mitigation schemes.

1.1 Challenges in Low-Power Wireless Coexistence

How wireless devices can navigate their way through the dense and diverse spectral environments and deal with interference is fundamentally central to the design of wireless systems. The coexistence of low-power wireless systems in the crowded radio spectrum is particularly challenging. At the core of typical low-power wireless systems are embedded devices with limited computation and communication capabilities. The constrained resources available on these systems poses several challenges to the design of dependable wireless protocols and coexistence mechanisms. Furthermore, low-power wireless systems are expected to compete for the shared spectrum with a wide range of devices that are typically less constrained and often adopt greedy practices in sharing the spectrum. In the following, we further highlight these challenges.

1.1.1 Resource-Constrained Devices

Typical platforms for low-power wireless systems feature a low-power Microcontroller (MCU), limited memory, tight energy budget, and a short-range, low-rate wireless radio transceiver. To have a better understanding of the extent of available resources in such systems, we discuss two example of typical low-power platforms: (i) a TelosB device [119] (also known as Tmote Sky) is a common platform in wireless sensor network deployments. This platform features a 16-bit MSP430 microcontroller operating at frequencies of up to 8 MHz and is equipped with a TI CC2420 wireless radio transceiver compliant with the IEEE 802.15.4 standard, featuring 10 kB of RAM and 48 kB of ROM. (ii) OpenMotes are based on the TI CC2538 microcontroller [113] (similar platforms are used in a range of IoT applications, such as health-tracking wristbands). They feature a 32-bit ARM Cortex-M3 SoC at 32 MHz. They are equipped with IEEE 802.15.4 standard compliant radio transceivers, 32 kB of RAM, and 512 kB of ROM. The resources available in these systems are several orders of magnitude lower than what is the norm on modern mobile and computing devices. Such scarcity of resources limits the ability of wireless protocols to perform computationally intensive and memory hungry operations on these devices (i.e., buffering and sophisticated signal processing). Hence, the constrained nature of these devices poses serious challenges to the design of wireless communication protocols, ultimately hampering their ability to withstand CTI. Moreover, energy efficiency is a crucial factor for these systems; such systems are anticipated to operate continuously and

unattended for long periods that can range from a few weeks to several years [114, 28, 16] with a limited energy budget. This constraint largely influences the design of wireless protocols, where radios should operate in a low-power sleep mode for most of the time and transmit at a low power to maximize their lifespan.

Researchers and industry players are pushing for microscopic low-power platform designs [86, 94] (i.e., cubic-millimeter) so sensing and computation capabilities can be fabricated and concealed in all objects (i.e., smart objects). This will fuel an array of innovative services and applications that will change the way we perceive and interact with the physical world. The constant push towards miniaturization of technology implies that these devices are likely to remain extremely resource-constrained to confine with the space requirements.

1.1.2 Spectrum Crunch

Radio spectrum is a globally finite resource, that needs to be effectively allocated and shared. Despite the notable recent advances in radio spectrum efficiency, it is anticipated that demand for wireless services is likely to outstrip radio spectrum capacity in the near future [47, 11]. This high demand trend is present for both the unlicensed and licensed bands but is more pressing for the unlicensed bands. Failing to address this spectrum crunch will lead to serious consequences on wireless service quality. Meeting this high demand necessitates new technologies that intensify frequency use and for the government to reallocate and open new spectrum bands.

In 1985 the Federal Communications Commission (FCC) devised to open up the 900, 2400, and 5800 MHz bands for the unlicensed use in data communications industry. Since then these bands have been home for a wide range of standardized and propriety wireless technologies and devices. The IEEE 802.11 (also referred to as WiFi) is a prime example of a high throughput pervasive technology that exists in these bands. These bands are typically much less regulated and wireless networks coexist without any form of coordination. Hence, interference is inevitable in these bands. The coexistence of heterogeneous co-located wireless systems is a technically challenging issue both for medium access control (MAC) and physical (PHY) layer designs. The inherent problem in radios today is that they do not know much about their neighboring networks (unaware of the technology type, proximity, and spectrum usage patterns of co-located networks) to act upon optimally. This is mainly due to the absence of communication means between these diverse technologies (i.e., speak different PHY protocols). Hence, coexistence between radios of different types is anarchic.

Low-power wireless systems have to compete for their share in these anarchic settings. Some of the communication properties and aspects that are adopted by radio technologies populating the unlicensed bands make it particularly challenging for low-power wireless technologies to coexist in the shared spectrum. For instance, (i) Wide-band: many devices use wide-band channels that affect a large segment of the available spectrum. Moreover, wireless systems are in general increasingly moving to wider frequency bands to cope with the high throughput demands. For example, to cope with the immense demand for high throughput over WiFi, the recent amendments of 802.11 allow the configuration of 40 MHz-wide channels in the 2.4 GHz band. Also, recent efforts in the 802.11 community are advocating to discard the notion of channelization to allow nodes to access a wider spectrum in order to improve load demand [79, 124]. Another case of wide-band occupancy are microwave ovens; they typically affect a large segment of the available spectrum in the 2.4 GHz band. (ii) High-power: due to the inherent application requirements, devices operating in the unlicensed bands transmit at different power levels. Low-power radios typically transmit at less than 1 mW for energy efficiency requirements, others, such as analog phones, can transmit at the maximum allowed power (i.e., 1000 mW). This severe power asymmetry poses significant coexistence problems, where high-power interferers can completely starve low-power technologies. That is because a typical high-power interferer might fail to detect the transmission of a nearby low-power transmitter, thus can interfere with the low-power node's transmission and monopolize the shared channel. Although Electromagnetic Compatibility (EMC) regulators lightly regulate this aspect by setting an upper limit of 30 dBm for transmit-power in the unlicensed bands, energy leaks from microwave ovens can reach up to 60 dBm. This is significantly higher than the typical output power of low-power wireless systems which is 0 dBm. (iii) Mobile-phone carriers presence in the unlicensed bands: Wireless carriers are developing systems (i.e., LTE-Unlicensed (LTE-U) or Licensed Assisted Access (LAA) [1, 3, 2]) that allow them to offload large segment of mobile data traffic into the unlicensed spectrum. This would further exacerbate congestion in the unlicensed bands.

1.2 State of the Art

In the following, we cover prominent related work on interference mitigation, cognitive communication, and wireless experimentation. The work in this dissertation contributes to and builds on ideas from these

research areas. Additionally, each chapter exhibits a dedicated related work section to cover more specific related work.

1.2.1 Interference Management and Mitigation

Wireless interference is (and has long been) an important topic in the wireless communication research. Wireless is inherently a broadcast medium and the unlicensed radio spectrum is fundamentally a scarce resource that an increasing number of devices have to share, in an uncoordinated manner. Consequently, exacerbating the interference problem. Recent years have seen significant and fundamental contributions to the state-of-the-art interference management. A large body of work on radio interference mitigation is available in the literature and there is a various set of strategies that can be employed to tackle this challenge [49]. Many of these solutions have been successfully used in the past or are being used in contemporary wireless systems. In the following, we highlight prominent research directions to address wireless interference.

Interference Avoidance. The most widely adopted approach to deal with interference is to avoid it. Devices employ mechanisms that can facilitate transmitting signals in segmented non-overlapping time slots, spaces, or frequency bands to avoid interference. These mechanisms can be part of the physical layer (PHY) or the media access control (MAC) layer. Frequency-based isolation is the most common isolation approach employed in wireless systems. This approach embodies mechanisms such as employing spectrum sensing to identify interference-free channels [126, 162, 31] and adaptive frequency fragmentation techniques [29, 127, 163, 79, 124, 106]. While these mechanisms focus on the careful tuning of signals to realize frequency isolation, others exploit frequency diversity to increase resilience against interference [107, 9, 40, 105, 83, 60] where the radio signal is spread over multiple channels. Other efforts focused on avoiding interference in time by exploiting channel temporal diversity. Interferers generally exhibit some regularity that can be learned and harnessed by nodes to schedule their transmissions in idle periods [80, 21, 161, 30]. Finally, researchers harness opportunities arising from space diversity, by exploiting directionality offered by antenna beam steering to avoid interfering with a co-located interferer [151, 88, 12]. Analogously, in networks with high node density, multiple paths exist, that nodes can harness to send packets to alternative paths that are less affected by interference [58, 75].

To sum up, the core idea of interference avoidance is to bring mechanisms that can allow communicating over interference-free links.

These links can be pre-established via coordinated multiple access in time, frequency, code, and space domains or temporally established based on the interferer behavior. However, the lack of interference-free channels (i.e., due to the rapidly surging number of wireless devices sharing the scarce spectrum), and the fast and unpredictable changes in the occupancy state of frequency bands make the overhead of these approaches high, particularly for resource-constrained devices.

Recovering from Interference. There is a body of research that focuses on recovering from interference by treating other transmitter's signals as noise, or decoding interfering signals (i.e., interference cancellation). This direction instead of avoiding interference copes with it by increasing communication resilience. In this approach, PHY and/or MAC layers are braced with a set of recovery mechanisms that help to restore data segments corrupted with interference. Examples of such mechanisms include: (i) Adding redundancy by applying a resilience forward error coding scheme to interference [96, 61, 133], (ii) partial packet recovery mechanisms that identify packet segments that suffered corruption from interference and selectively retransmit only the interfered segments [71, 66, 85, 115, 159, 63, 99], and (iii) interference cancellation schemes. Here the receiver, with minimal or no coordination from the sender, attempts to recover the signal of interest from interference. Typically, the receiver decodes the interfering signal first, i.e., the signal with larger power. Afterward, the interference signal is stripped away from the aggregately received signal to get the target signal [54, 166, 62, 56, 36]. Typically, these solutions are customized for a particular interference type or pattern. Hence, they fall short in the presence of heterogeneous interference.

1.2.2 Cognitive Communication

The work in this dissertation builds on the core idea of cognitive communication. The concept of cognitive radio was first introduced by Mitola [104]. He presented a broad vision for wireless communication that he described as: "The point in which wireless personal digital assistants and the related networks are sufficiently computationally intelligent about radio resources and related computer-to-computer communications to detect user communications needs as a function of use context, and to provide radio resources and wireless services most appropriate to those needs". The emerging spectrum crunch raised interest in cognitive radio. Cognitive radio generally refers to radio devices that have the capability to sense their RF environments and adapt their spectrum usage accordingly [67, 17]. This broad vision has been followed by research efforts in signal detection and classification [67, 8],

spectrum sharing models [24, 23], and centralized and distributed spectrum sharing protocols [26, 168, 167]. The cognitive radio term now is largely coined with unlicensed devices access to frequency bands that are currently reserved for licensed usage. Therefore, prior work focused on spectrum sensing in the licensed bands (detecting temporarily unused bands), where it is critical for unlicensed secondary users not to interfere with the licensed primary users.

The cognitive communication's native vision is broader than what is currently the norm in cognitive radio [104]. It enfolds incorporating machine-learning techniques that can make radios trainable in a broad sense. Cognitive radio research drifted from this vision and is now primarily governed by complex decision-making processes. In the recent years, we have witnessed the rise of interest in bringing cognitive radio concepts to the unlicensed bands. This is driven by the need to cope with high throughput demands. Few examples include work in Agile Radios, Dynamic Frequency Selection (DFS), and Smart Antenna Array. The work in this dissertation explores and builds on cognitive communication concepts to enhance wireless systems coexistence.

1.2.3 Low-Power Wireless Experimentation

Researchers in wireless communication rely on two primary methodologies for performance analysis, namely network simulation and testbed experiments. Network simulation is a widely used methodology for the development and verification of new network protocols and communication constructions. To run their experiments, researchers use a network simulator, a software that allows modeling arbitrary computer networks by specifying experiment settings such as the behavior of the network nodes and the communication channels. Simulators can be either in-house built that are developed and customized by researchers for their experiments or publicly available such as ns-3 [110], OMNeT++ [112], TOSSIM [95], and Cooja [41]. Publicly available simulators are often based on the discrete event-based simulation paradigm [156, 48] and are supported by complex models and, hence, are more credible. The wireless research community is often critical of work that is supported by only simulation results. Simulation-based experimentation provides flexibility, control, and repeatability but lacks realism and is plagued with inherent inaccuracies. To overcome this, researchers have understandably adopted testbed experimentation as a standard methodology for performance evaluation in wireless networks. They enable evaluation with realistic testing grounds, channel propagation, interference conditions, timing requirements, and hardware constraints. However, evaluation with testbeds can be tedious. Researchers are

throttled by high implementation and cost barriers, and complex and time-consuming experiment set-ups. These barriers can preclude testing in real systems. To facilitate a broader adoption of testbed-based experimentation, the communication research community worked on developing a large number of testbed facilities that have simulation like usability experience (i.e., remote users upload the specifications of their experiment and collect traces from a web interface), where setup and maintenance are managed by the testbed hosting institutes. These testbeds are heterogeneous in many dimensions: supported hardware and software, the number of available nodes, resources, supported wireless technologies, and accessibility. Some of the prominent publicly available testbeds are Emulab [70] (University of Utah, U.S.), Orbit [131] (Rutgers University, U.S.), PlanetLab [32] (University of California Berkeley, U.S.), FlockLab [97] (ETH Zurich, Switzerland), Indriya [37] (National University of Singapore), MoteLab [157] (Harvard University, U.S.), IoT-Lab [6] (Inria, France), and TWIST [64] (TU Berlin, Germany). Testbeds have constantly evolved in the last years to reflect accurately on how various environmental effects can be reproduced, such as extending existing testbed infrastructures with the support of node mobility and radio interference [20, 141]. Furthermore, several testbed facilities are being augmented with hardware and software support for cognitive systems [121]. Despite recent progress in testbed capabilities, their ability to reproduce and reflect on heterogeneous radio interference conditions is still limited. Throughout this dissertation, we opt to validate our systems using customized testbeds we set up explicitly for this work. Publicly available testbeds do not yet support tools for wireless coexistence research. In Chapter 5, we elaborate more on this limitation, and present a solution that can help to overcome this limitation by extending current testbed infrastructures with the capability of reproducing heterogeneous external interference to facilitate wireless coexistence research.

1.3 Dissertation Contributions and Roadmap

In this dissertation, we take an alternative approach to deal with CTI, where instead of avoiding interference, we adopt an interdisciplinary approach combining cross-layer designs and lightweight machine-learning techniques to build cognitive low-power wireless systems that can cope with CTI. We begin by acquiring a good understanding of how interfering wireless signals interact and harness this understanding in our designs. We then introduce a family of algorithms and system architectures that improve the robustness of low-power wireless networks

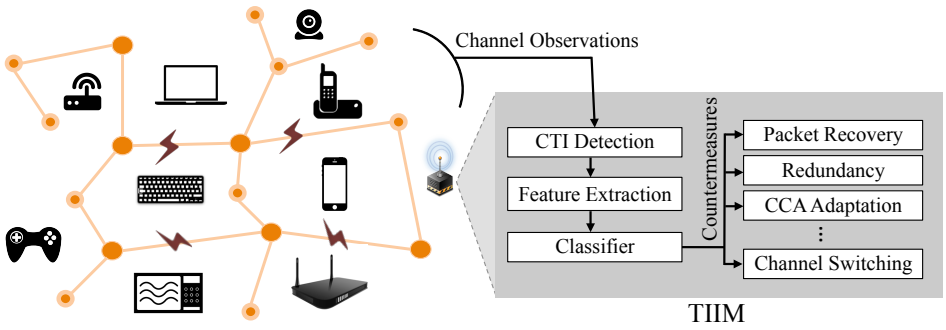

Figure 1.2: **TIIM** dynamically applies interference mitigation measures specific to channel conditions **(Chapter 3)**.

operating in interference rich environments. All systems presented in this dissertation have been deployed and evaluated in testbeds. Our evaluations reveal large performance gain achieved in each system. The detailed results are described in the individual chapters. The specific contributions of this dissertation are highlighted below.

1.3.1 Understanding the Impact of Cross-Technology Interference on Low-power Wireless Networks (Chapter 2)

Interfering radio technologies differ widely in the way they affect wireless links. Cross-Technology Interference has a strong and complex impact on wireless links that need to be taken into account when treating interference. To acquire a better understanding of the problem, we conduct a comprehensive empirical study of CTI implications on low-power wireless networks. The purpose of this study is twofold. First, acquiring a good understanding of how interfering wireless signals interact in the shared channel. Second, identify and highlight vulnerabilities in existing communication protocols that hinder them to withstand CTI. In a controlled environment, we expose a low-power wireless network to a set of prevalent interferers. The set of considered interferers is selected to represent common underlying properties adopted by most of the nowadays used wireless devices. The study analysis covers observations at multiple layers of the communication stack, namely, physical layer, MAC layer, and application layer's payload. In this study, we show that the uncertainty that CTI induces on the wireless channel is not completely stochastic; CTI exhibits distinct patterns that can be exploited by interference-aware protocols.

Chapter 2 describes in details our study and presents a detailed analysis of the CTI implications on low-power wireless networks.

1.3.2 Adaptive Cross-Technology Interference Mitigation (Chapter 3)

To date, much of the devised radio frequency interference solutions focused on resolving interference between devices of the same technology. There exists no systematic mechanism for radios to be aware of what other radio types exist in their environments and make smart decisions to adapt accordingly. Hence, coexistence between radios of different types is anarchic. To address this challenge, we resort to machine learning; we employ supervised learning to train radios to recognize interference patterns at which a particular link-layer mitigation strategy would work best, regardless of the interference type. To demonstrate the feasibility of this approach, we construct TIIM, as illustrated in Figure 1.2, a lightweight Technology Independent Interference Mitigation solution that detects, quantifies, and reacts to CTI in real-time. TIIM selects interference mitigation strategies directly based on measured medium properties. We train our system to detect interference patterns and map these to a link-layer interference mitigation strategy that works best for this particular pattern. Our evaluation shows that TIIM, while exposed to extensive and heterogeneous interference, can achieve a total packet reception rate improvement of 30% with an additional transmission overhead of 5.6%.

Chapter 3 describes the detailed architecture of TIIM and provides results of our testbed evaluation.

1.3.3 Exploiting PHY Layer Information to Combat Cross-Technology Interference (Chapter 4)

In current radio transceiver designs, the accessible information about the environment is limited. Therefore, radios coexistence policies are likely to be suboptimal. Over the last few years, researchers advocated for new wireless designs that allow better interfacing between the physical layer and higher layers, particularly, to cope with interference. Typically, information delivered from the physical layer to upper layers is restricted to decoded bits. However, physical signals convey rich information about the ambiance, which is particularly enlightening in the case of interference. In this work, we explore how physical layer information can be exploited towards wireless coexistence. We present a cross-layer wireless design, named CrossZig, which enables radios to harness fine-grained physical layer information to recover from varying sources of interference. We implement a prototype of CrossZig for the IEEE 802.15.4 in programmable radios. We show the performance gain of CrossZig through experimental evaluation considering both micro-benchmarking and system performance under various interference patterns.

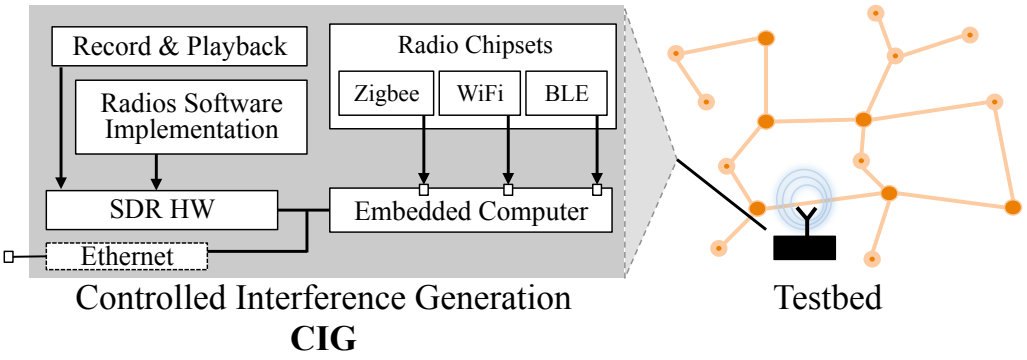

Figure 1.3: Schematic of our Controlled Interference Generation (**CIG**) framework, facilitating advanced wireless coexistence experimentation **(Chapter 5)**.

Chapter 4 expands on CrossZig architecture and components, describes its implementation on programmable radios, and details results of our CrossZig evaluation.

1.3.4 Wireless Coexistence Experimentation (Chapter 5)

As wireless networks and systems become omnipresent, there is a strong need for testing, understanding, and debugging existing wireless protocols against increasingly complex settings. Publicly available research testbeds often lack the support for repeatable generation of realistic interference patterns. Hence, studying coexistence problems between wireless devices is often a cumbersome process. In this dissertation, we address the lack of interference generation support by proposing and developing **CIG** (see Figure 1.3), a software-defined radio-based interference generator that generates customizable and repeatable interference in real-time. We consider a unified approach that incorporates a careful selection of interferer technologies implemented in programmable radios (i.e., supporting a range of digital modulation formats) and dedicated hardware that exposes the networks to realistic interference patterns. CIG provides an interface for generating interference from software implementations of devices commonly operating in the 2.4 GHz ISM band, as well as a playback interface that allows regenerating previously recorded interference patterns. With CIG researchers can rerun experiments under almost identical radio environmental conditions and, hence, CIG would largely facilitate wireless coexistence research.

Chapter 5 describes the detailed architecture of each of CIG's components and provides results of the implementation and validation of CIG on programmable radios.

1.3.5 Roadmap

The main chapters of this dissertation are organized in a chronological and didactic order and can be read as separate self-contained components.

We start this dissertation with an empirical study of Cross-Technology Interference (CTI) implications on low-power wireless networks in Chapter 2. The observations made in this chapter largely influenced the approach we adopt to address low-power wireless systems coexistence in this dissertation. Chapter 3 describes **TIIM**, a novel radio design which marries concepts of cognitive radio, machine learning, and networking into a complete system that enhances low-power wireless performance in radio rich environments. **CrossZig**, described in Chapter 4, is a cross-layer solution that enables low-power wireless nodes to make informed decisions on their coexistence strategies based on richer physical layer information. In Chapter 5, we present Controlled Interference Generation (**CIG**), a software-defined radio based solution for controlled interference generation, which can facilitate augmenting current testbeds with repeatable and realistic interference pattern generation. We conclude this dissertation in Chapter 6 with a summary of our contributions and a discussion of open challenges and future research avenues in low-power wireless systems coexistence.

All the results presented in this dissertation have been peer-reviewed, published, and presented at international conferences and workshops. The associated publications are listed at the end of the dissertation (**List of Publications**) and referenced at the corresponding chapters.

2

Cross-Technology Interference Characterization

Over the last few decades, we have witnessed a notable progress in wireless communication. This has led to a rapid emergence of heterogeneous wireless technologies that share the radio spectrum in an un-coordinated way [92]. Such a coexistence introduces uncertainty and complexity to the medium, affecting reliability and availability of wireless networks. This problem aggravates for technologies operating in the lightly regulated, yet crowded unlicensed bands. The unlicensed bands proliferate with heterogeneous devices including WiFi (IEEE 802.11), Bluetooth, 2.4 GHz cordless phones, microwave ovens, surveillance cameras, game controllers[1], and 2.4 GHz RFID. These technologies differ widely in terms of emitted power levels, wireless medium access modalities, used modulation and coding schemes, and in the width of occupied sub-bands. To address the coexistence of different technologies in the scarce radio spectrum, and to provide proper interference-aware protocols and mitigation schemes, we need to develop a clear understanding of how these technologies interact in the shared spectrum.

Contributions and Roadmap. We begin this dissertation with an empirical study of the implications of Cross-Technology Interference (CTI) on the particularly vulnerable low-power IEEE 802.15.4 wireless networks. In this study, we identify the underlying vulnerabilities that hamper 802.15.4 to withstand CTI. Furthermore, we show that the uncertainty that CTI induces on the wireless channel is not entirely

[1] For example, the Xbox 360 S wireless controller.

stochastic; CTI exhibits distinct patterns that can be exploited by interference-aware protocols. The set of interferers we consider in this study are selected to represent common underlying properties adopted by most of the nowadays used wireless devices. Our considered set consists of low/high power interferers, narrow/wideband interferers, analog/digital interferers, channel hopping/fixed frequency interferers, and CSMA and non-CSMA interferers.

The study analysis enfolds observations at multiple layers of the 802.15.4 communication stack: *(a)* Physical layer (PHY): investigation of PHY characteristics captured from off-the-shelf 802.15.4 radio chips, through fast sampling of the Received Signal Strength Indicator (RSSI) register and other channel indicators; *(b)* MAC layer: exploring CTI impact on the Clear Channel Assessment (CCA) and Carrier Sense Multiple Access (CSMA) backoffs; *(c)* Upper layers payload: analyzing corruption features such as error patterns, error bursts, and interspaces between consecutive errors. Our results show that different technologies affect 802.15.4 distinctly in aspects such as corruption rate, backoff mechanism, the location of corrupted symbols, etc. This knowledge can be exploited by interference mitigation schemes for a better resilience against CTI. Acquiring a clear understanding of CTI's footprints on low-power wireless systems is an essential step for designing reliable low-power wireless protocols that can gracefully coexist in the shared spectrum. The observations we make in this study have largely influenced the approach we adopt to address the wireless coexistence challenge throughout this dissertation.

The remainder of this chapter is structured as follows: Section 2.1 briefly reviews IEEE 802.15.4 PHY and MAC specifics. Section 2.2 describes our experiment setup and configurations. Section 2.3 discusses the impact of the interaction between high/low-power interferers and 802.15.4 networks. Section 2.5 presents related work. We conclude this study in Section 2.6. This chapter is based on the contributions made in [76].

2.1 Background

In this chapter and throughout this dissertation, we build prototypes and discuss systems that are compliant with the IEEE 802.15.4 standard. This communication standard exemplifies a low-power wireless technology which is the target of our work and is widely used in a variety of IoT and WSN applications and deployments [118, 108, 148]. Here we briefly review relevant aspects of the IEEE 802.15.4 standard to our work.

RF Technology	Abbr.	TX Power (dBm)	Bandwidth (MHz)
IEEE 802.15.4	–	0	2
Bluetooth (Class 2)	BL	4	1 (FH)
Wireless Camera	CAM	20	1.125 (FH)
Analog Phone	AN-P	n/a	0.1
FHSS Phone	FH-P	21	0.8 (FH)
Microwave Oven	MW	60	-
IEEE 802.11	WiFi	20	20

Table 2.1: Characteristics of the considered RF technologies in our study.

2.1.1 IEEE 802.15.4

The IEEE 802.15.4 standard defines both the physical (Layer 1) and data-link (Layer 2) layers of the OSI model for low-rate wireless personal area networks (LR-WPANs). The standard is maintained by the IEEE 802.15 working group and is part of the IEEE 802 standards committee responsible for specifying wireless personal area network (WPAN) standards. The focus of this working group is on defining standards for wireless communication with low data rate, very long battery life, and low complexity. The IEEE 802.15.4 standard is the basis for several standardized and proprietary network protocols, including IEEE 802.15.5, ZigBee, 6LoWPAN, Thread, WirelessHART, and ISA100.11a. Each of these protocols further extends the standard by developing upper layers protocols which are not part of the current standard.

IEEE 802.15.4 PHY. For devices operating in the 2.4 GHz band, the IEEE 802.15.4 standard [81] mandates the use of *Offset Quadrature Phase-Shift Keying* (O-QPSK) modulation scheme with a half pulse shaping. In order to increase the resistance against noise, *Direct-Sequence Spread Sequence* (DSSS) is employed. The transmitter's radio transforms binary data to modulated analog signals by adapting spreading and modulation. The data is first grouped in 4-bit symbols, which are mapped to one of the 16 *Pseudo-random Noise* (PN) sequences that are 32-bit long. Each bit in a PN sequence is referred to as a chip, which is then modulated onto the carrier signal using O-QPSK. Transmissions on the 2.4 GHz band are fixed at a rate of 250 kbps.

For demodulation, the receiver's radio converts each half-sine pulse signal into a chip. The radio performs soft decisions at the chip level, providing PN sequences with non-binary values ranging from 0 to 1 [147]. The de-spreading is performed by mapping the PN sequence to the symbol with the highest correlation. The redundancy induced by spreading allows correct decoding of the received symbol, even if few

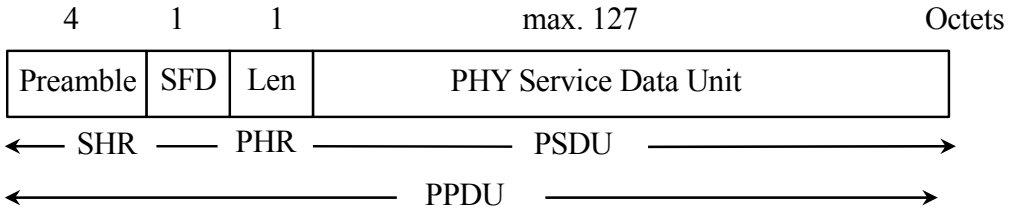

Figure 2.1: IEEE 802.15.4 PHY frame structure.

chips were not correctly decoded which increases the immunity to noise.

The PHY Protocol Data Unit (PPDU) frame format is illustrated in Figure 2.1. The PPDU consists of three fields, Synchronization header (SHR), Physical header (PHR) and PHY Service Data Unit (PSDU). The frame starts with SHR which consists of a preamble to allow clock synchronization, followed by Start Frame Delimiter (SFD). The next byte accommodates PHR which holds the frame length in bytes (coded on 7 bits) and one reserved bit. This is followed by PSDU of a maximum length of 127 bytes. In order to support larger network layer packets, adaptation layer protocols such as 6LoWPAN are typically needed to provide fragmentation schemes.

IEEE 802.15.4 MAC. IEEE 802.15.4 has several MAC-layer protocols, defined both in the original standard and its 2012 amendment 802.15.4e. In its simplest form, 802.15.4 employs contention-based CSMA/CA communication. Before a node starts transmission, it waits for a random back-off period to assure that the medium is idle. For this, it relies on *Clear Channel Assessment* (CCA). The determination of CCA considers *Energy Detection* (ED) or/and detection of 802.15.4 modulated signal in the channel. If CCA declares the channel to be free, the transmission is carried out, otherwise it defers the transmission for a random backoff time. For data verification, the receiver computes a 16-bit CRC check over the payload of a received packet. It discards packets that do not pass the check and accordingly withholds the ACK transmission. More sophisticated 802.15.4e MAC layers such as TSCH, CSL, or RIT also employ acknowledged transmissions, exponential backoff, and (optionally for TSCH) CCA.

IEEE 802.15.4 Channels. IEEE 802.15.4-conformant devices can use any of the specified frequency bands for operation (868/915/2450 MHz). 802.15.4 transmission occurs in one of the 27 non-overlapping allocated channels. Out of these, 16 (from 11 to 26) are allocated in the 2.4 GHz band, each with 2 MHz bandwidth and 5 MHz channel spacing. The Center Frequency (F_c) for the 2.4 GHz channels is defined as: $F_c = 2405 + 5(k - 11)$ MHz for $k = 11, 12, ..., 26$. The remaining 11 channels are allocated in sub-GHz bands (915 MHz and 868 MHz).

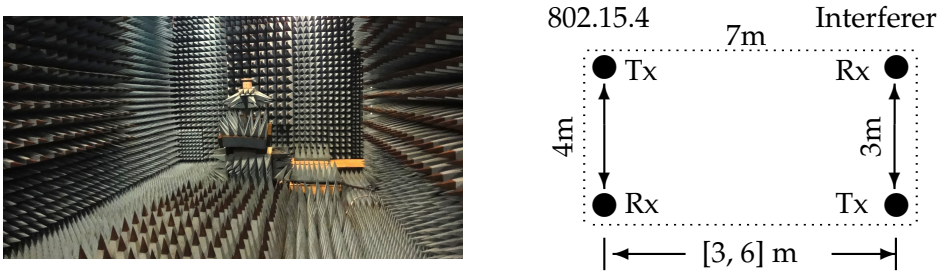

(a) Anechoic room (b) Setup

Figure 2.2: Experiment setup for the CTI impact study in an anechoic room.

2.2 Study Overview

In order to understand the implications of CTI on 802.15.4 links, we conduct a comprehensive, in-depth measurement study, where we subject 802.15.4 links to various interference sources and collect statistics on communication and channel measurements that are visible to nodes. We use the collected measurements to make observations on the temporal trends and channel effects of interfered 802.15.4 links.

Experimental Setup. We run our experiments in an anechoic chamber, with dimensions 7 m x 4 m x 4 m (length, width, height). The chamber belongs to the Millimeter-Wave Electronics Laboratory at ETH Zurich.

An anechoic chamber is a shielded room that is designed to absorb electromagnetic waves reflections. The chamber walls are typically covered with pyramidal radiation absorbent material (RAM) (see Figure 2.2(a)) that scatters and absorbs much of the incident energy. The chamber is as well insulated from exterior sources of noise. Anechoic chambers are typically used for conducting measurements of electromagnetic compatibility and antenna radiation patterns. In this study we focus our measurements in this highly controlled settings, so to have full control on the source of errors, to isolate the impact of surrounding interference sources, and to identify the mere impact of each of these interfering technologies. We consider a simple network setup, as depicted in Figure 2.2(b), which consists of one transmitter and one receiver for both 802.15.4 and the considered interfering technology, i.e., a pair of 802.15.4 nodes and a pair of interferer nodes. We alternate the transmission power of the 802.15.4 nodes to feature attenuation levels of weak signals and emulate greater distances.

Wireless Low-power Platform. We use the Tmote Sky [119] nodes primarily as our experimental platform for low-power 802.15.4 transmitters and receivers. Tmote Sky nodes feature CC2420 radios [147] which are compliant with the IEEE 802.15.4 standard and are widely used radio interfaces.

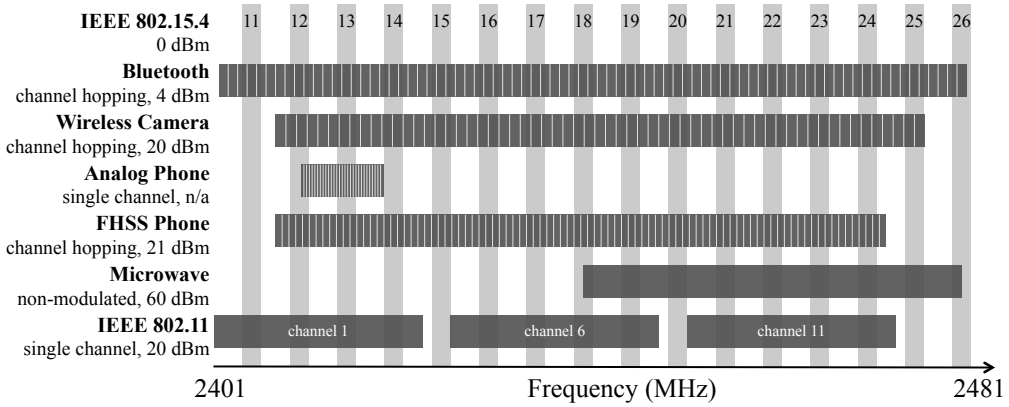

Figure 2.3: RF channels of IEEE 802.15.4 and the selected set of prevalent RF interferers in the 2.4 GHz ISM band studied in this chapter.

The CC2420 attaches two physical layer metadata to every received packet, RSSI (received signal strength indicator) and LQI (link quality indicator). These metrics are measured over eight symbols (32 bits, 125 µs) of received packets. The CC2420 calculates the LQI based on the first eight symbols of received packets but continuously calculates RSSI. Therefore, we can trigger the software to read RSSI register at any time to measure ambient RF energy. The nodes have an integrated Omni-directional inverted-F microstrip antenna.

Interfering Technologies. The set of interferers we consider in this study is selected to represent common underlying properties adopted by most of the nowadays used wireless devices. We consider technologies with low and high emitting powers consisting of the following interference sources: IEEE 802.11 (WiFi), Bluetooth, FHSS, and analog cordless phones, microwave ovens, and surveillance cameras, e.g., baby monitors (see Table 2.1). We analyze the spectral and temporal characteristics of non-standard interferer technologies considered in this study. For this, we use the software defined radio USRP N210 [44] to monitor a 25-MHz bandwidth at a given time. We round the scan in 4 tuning steps to cover 80 MHz of the 2.4 GHz band starting from 2.40 to 2.48 GHz. Figures 2.5, 2.6, 2.7, and 2.8 show the spectrograms and power-profiles for a subset of the considered RF technologies. The technologies and devices we use, are described in more detail in Section 2.3.

Communication Scenarios. As we aim at exploring low-level interference effects as precisely as possible, we eliminate all network protocol overheads by writing our receiver and sender applications to directly interface the CC2420 radio driver in the Contiki OS [39]. We use the three following communication scenarios: *(a)* **CCA-enabled:** transmissions are sent at 100 ms interval, conditioned by a CCA

(i.e., CSMA enabled) with exponential back-off, and followed by an acknowledgment (ACK) frame; *(b)* **CCA-disabled:** transmissions are sent at 100 ms interval, CSMA disabled and acknowledgments are enabled; *(c)* **Saturated:** transmissions are sent at 8 ms interval, with CSMA and acknowledgments disabled. In the first two scenarios, the transmission interval is constrained to 100 ms because of the time needed to log fast RSSI sampling information over the serial line. The third scenario disables fast RSSI sampling to reduce this interval and allows us to study the correlation among packets sent consecutively. The first two experiments run for 1600 packets, the third for 3200 packets. We fill the packet's payload with one of the 802.15.4 symbols, and periodically iterate over all 16 existing symbols to eliminate the content's influence on our results (i.e., coding can lead to some symbols being more stable than others [69, 134]).

In all three scenarios, we run a series of experiments where we vary the packet sizes among 20, 40, 100 bytes, and the power level in the range of high (0 dBm), medium (-3 dBm), and low (-10 dBm). We recognize three types of packet reception: packets that are correctly received (passed CRC check), packets that got corrupted, hence, have at least one corrupted symbol (failed the CRC check), and packets that are lost: sent but never received (corruption affected the PHY header or synchronization header). In all experiments, in case we have pre-knowledge information on the exact used frequency ranges of the interferer, the transmitter and receiver are configured to communicate over one or two channels that overlap with that of the interferers. For the technologies that affect a wide range of channels, such as for microwave oven and FHSS interferers, we loop over every second channel of 802.15.4 to broaden our scope of analysis and not to miss hopping specific channel effects.

Measurements. The transmitter logs the number of retransmission attempts (if enabled) and the noise level for each sent packet. For each received packet, the receiver logs the following information: noise level, link quality indicator (LQI), checksum value, received packet content, and the received signal strength during the packet reception associated with each received packet. We modified the CC2420 driver in Contiki to: *(a)* instruct the radio to pass packets with failed CRCs rather than discard them, to enable us processing erroneous packets; *(b)* capture RSSI values at a rate of one sample per symbol (one reading each 16 μs). Upon the detection of an incoming packet, the *start of frame delimiter* (SFD) pin is set to 1, which triggers an interrupt. In this interrupt, we capture the variations of the RSSI during the reception of a packet. The sampling is performed until the last symbol of the packet is received and the SFD pin is set back to 0.

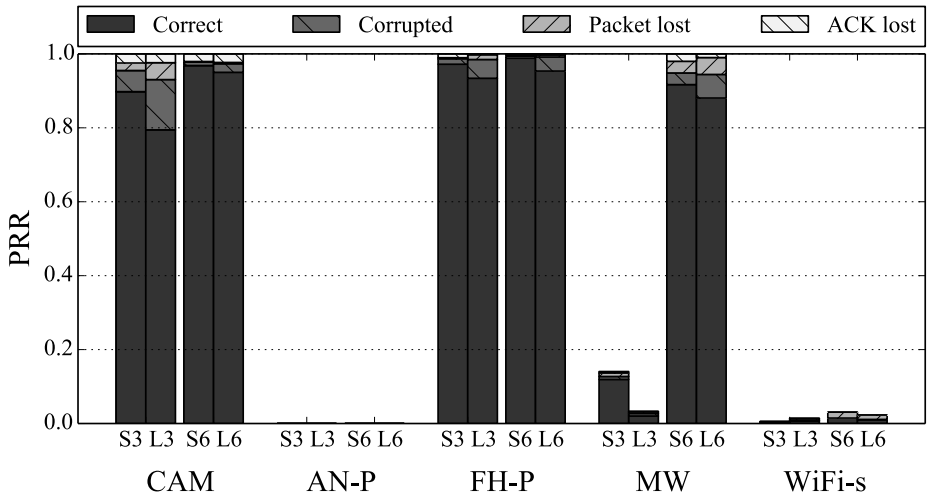

(a) CCA-enabled

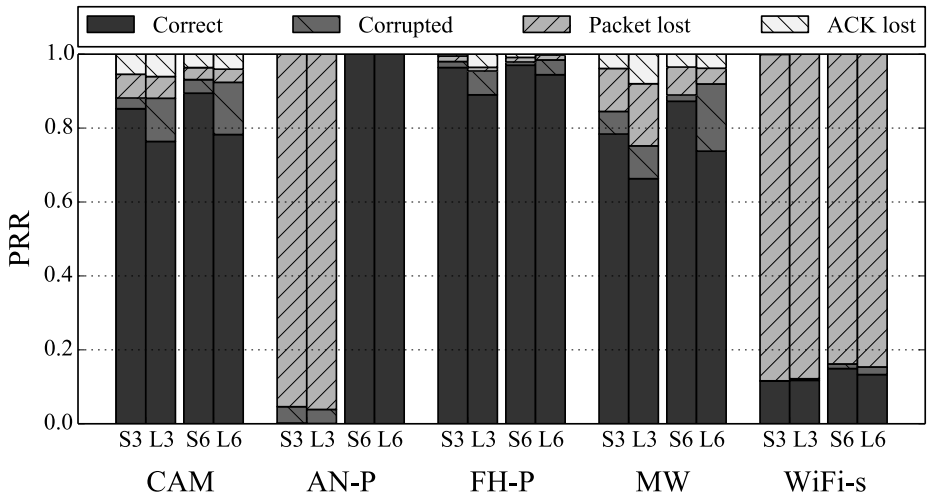

(b) CCA-disabled

Figure 2.4: Packet Reception Rate (PRR) for CCA-enabled and CCA-disabled traffic types for distances 3 and 6 m for packets with length 100 byte (L) and 20 byte (S). The saturated traffic type follows the same trends as CCA-disabled, hence, not shown here. Empty space in CCA-enabled traffic indicates no traffic due to busy medium, i.e., backoff. Bluetooth's and non-saturated WiFi's impact on the communication are almost neglectable.

2.3 Cross-Technology Interference Implications

In the following subsections, we provide an overview of the characteristics of each of the considered interferers, their overlapping spectral ranges with 802.15.4, and their direct impact on the 802.15.4 performance. The spectrum allocation of every technology considered is illustrated in Figure 2.3.

2.3.1 IEEE 802.11

Characteristics. IEEE 802.11 is the most pervasive wireless technology in indoor environments. The 802.11 b/g/n transmission occurs in one of the 14 overlapping channels spreading over the 2.4 GHz ISM band. Each channel has a width of 20 MHz, where most of these channels are overlapping with four of the 802.15.4 channels. At the physical layer, 802.11 supports a large set of modulation and coding schemes that trade performance with interference and noise tolerance. The communication signal is spread over the 20 MHz channel using DSSS or OFDM. Most 802.11 devices support power level ranges of -20 dBm to 20 dBm and commonly communicate at the highest transmission power of 20 dBm.

Setup. We evaluate the interference caused by 802.11 using a Netgear WNR3500L router and a laptop that supports IEEE 802.11 b/g/n in the 2.4 GHz ISM band. In our experiments, the router acts as an access point forwarding TCP/UDP traffic to the laptop which acts as a client. We use the network tool `iperf` [82] to generate saturated TCP traffic and non-saturated UDP traffic that resemble file download and VoIP, respectively. We configure the router to use channel 11, and study the interference impact on two 802.15.4 channels: Channel 22 as fully overlapped with the WiFi channel 11 and channel 24 which is partially overlapped with WiFi channel 11.

Observations. As shown in Figure 2.4, for all the considered configuration scenarios, the exchanged saturated TCP (WiFi-s) caused PRR to drop to below 20%. This can be attributed to the aggressive way of WiFi transmitting at 100 times higher power than the 802.15.4 nodes. Although 802.11 employs CSMA, the amount and regularity of the energy emitted by the 802.15.4 node are not sufficient to defer 802.11 communication. In the saturated TCP case the WiFi access point transmits nearly continuously, and as a result, the 802.15.4 node backs off or experiences severe packet losses. It is notable to highlight that the air time of 802.11 b/g/n packets is significantly shorter than the air time of 802.15.4 packets (about 0.54 ms for 802.11 g maximum packet length, 4.2 ms for 802.15.4 maximum packet length). The exchange of non-saturated UDP traffic, on the other hand, has a negligible impact on the performance of 802.15.4 nodes. Therefore, we only show the saturated TCP case in Figure 2.4.

2.3.2 Frequency Hopping Bluetooth

Characteristics. Bluetooth is a low-power wireless protocol standard for exchanging data over short distances in single-hop networks at ranges

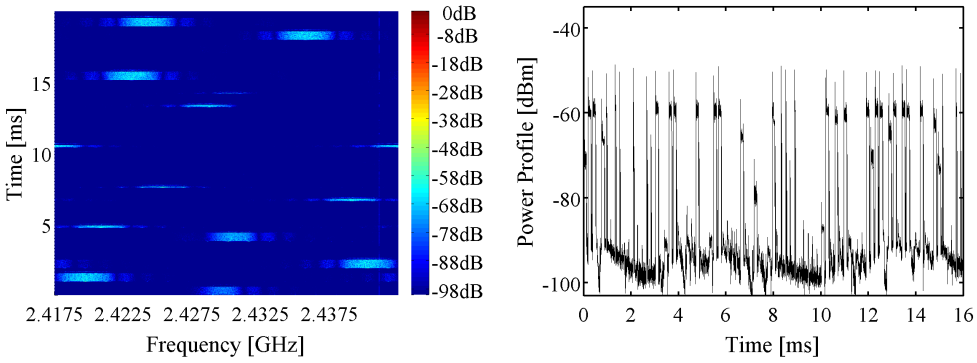

(a) Spectrogram (b) Power-profile

Figure 2.5: Spectrum characteristics of high-power Wireless Camera.

of typically less than 10 meters but can range to 100 meters. At the physical layer, Bluetooth uses the adaptive frequency hopping technique across a 79 MHz bandwidth in the 2.4 GHz ISM band, with each channel occupying a bandwidth of 1 MHz. The hopping occurs at a rate of 1600 hops/sec. Hence, it occupies a channel for 625 μs. Bluetooth defines different communication classes, which specify the transmission power, resulting into different communication ranges. However, the most common Bluetooth devices are the battery-powered Class 2, transmitting at 4 dBm which is higher than 802.15.4 devices (-25 dBm to 0 dBm) [81].

Setup. To evaluate the interference generated by Bluetooth on 802.15.4, we use two HTC Desire phones transferring a large file.

Observations. At both considered distances, Bluetooth did not have a notable impact on the performance of 802.15.4 nodes. Note, this observation cannot be generalized to other Bluetooth classes, as in a previous study [71], we observed a performance reduction of 20% caused by Bluetooth Class 1 devices.

2.3.3 Wireless Camera

Characteristics. As for a wireless camera, we use the Philips SCD 603 digital video baby monitor. It comprises a 2.4 GHz wireless camera and a wireless video receiver. The wireless camera communicates with the wireless video receiver using frequency hopping over 61 channels, where each channel has a width of 1.125 MHz.

Observations. The camera's spectrogram, as depicted in Figure 2.5, shows the frequency hopping nature of the wireless camera. Most of the hopping occurs in the frequency range [2.42-2.45] GHz. This matches our observations on the PRR, as 802.15.4 channels interleaved

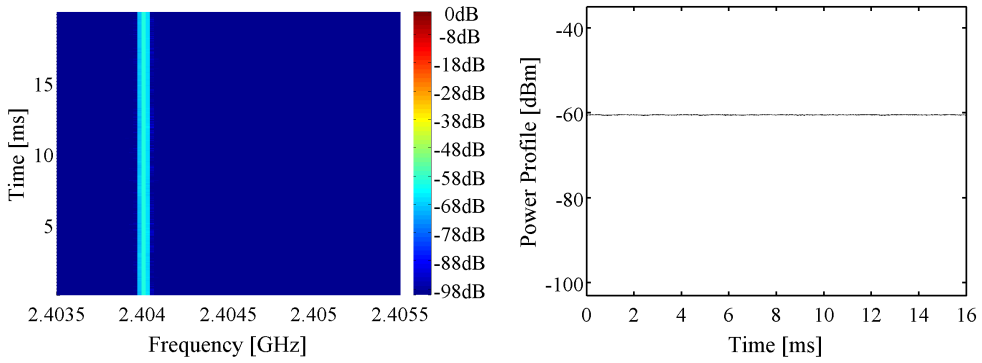

(a) Spectrogram (b) Power-profile

Figure 2.6: Spectrum characteristics of high-power Analog Phone.

in this range were affected the most. This could be due to the underlying spread sequence concentrating on this region of the spectrum. In case the camera experiences degradation in the quality, it could switch to another spread sequence that affects another region of the spectrum. In our analysis, we consider 802.15.4's channel 16 which falls in the above-mentioned frequency range. At both considered distances between interferers and the 802.15.4 nodes, we measure the performance of the 802.15.4 nodes with the camera being ON and OFF. For the CCA-enabled traffic, as shown in Figure 2.4(a), we observe more than 20% corrupted or lost packets for long data packets at distance 3 m, resulting into retransmissions. The frequency hopping nature of the interfering signal makes its effect less pronounced, specifically due to the relatively narrow band of 802.15.4, which makes it less impaired by frequency hopping.

2.3.4 Analog Cordless Phone

Characteristics. We use the Vtech GZ2456 cordless handset system in our experiments. The phone base, according to the device manual [152], transmits in the frequency range [2410.2 - 2418.9] MHz and receives in the frequency range [912.75 - 917.10] MHz. However, our experiments show that the phone base transmits in the 900 MHz band and receives in the 2.4 GHz band, which contradicts the manual description. The phone handset accordingly transmits and receives using the reverse order of frequency ranges. The phone picks a default channel out of 30 supported channels in the specified frequency range. It does not support automatic channel selection. However, channel switching can be configured manually by the user.

Observations. The spectrogram and power-profile of the analog cordless phone are illustrated in Figure 2.6. The frequency profile shows that the

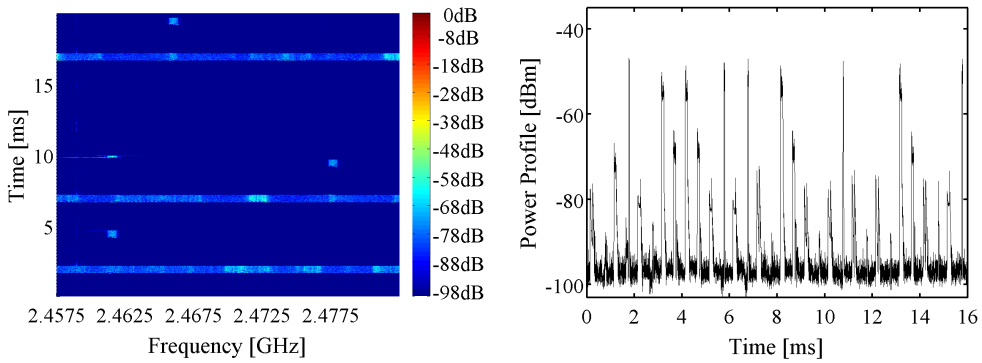

(a) Spectrogram (b) Power-profile

Figure 2.7: Spectrum characteristics of high-power FHSS Phone.

analog phone occupies a narrow channel (about 0.1 MHz) at a time. Based on the phone frequency profile, we select the 802.15.4 channel 13 centered at 2.415 GHz which overlaps with the analog phone's communication. As shown in Figure 2.4(a), 802.15.4 nodes while employing CCA could not communicate when subjected to analog phone interference, for both considered locations. This is due to the phone continuously transmitting, as seen in the corresponding power profile, depicted in Figure 2.6(b). As a result, the 802.15.4 transmitter backs off continuously due to the channel being occupied. In our experiment with CCA disabled, we force 802.15.4 transmission to occur regardless of ambient noise. Interestingly, at a distance of 6 m, as shown in Figure 2.4(b), most of the packets are received correctly. In this particular case, the default CCA-threshold based backoff cancels all transmissions, although communication is obviously still possible. We elaborate more on this behavior and possible workarounds in Section 2.4.

2.3.5 Digital FHSS Cordless Phone

Characteristics. We experiment with the Uniden DCT6485-3HS cordless handset system. The phone base and handset communicate using frequency hopping over 90 channels of 800 kHz width in the range [2407.5 - 2472] MHz.

Observations. As shown in Figure 2.4, the FHSS phone affects 802.15.4 similarly as the wireless camera, however, less destructive. This is attributed to the fact that both technologies employ the same underlying signal spreading scheme, i.e., frequency hopping, only with slight changes in channel width (cf. Table 2.1) and hopping rates.

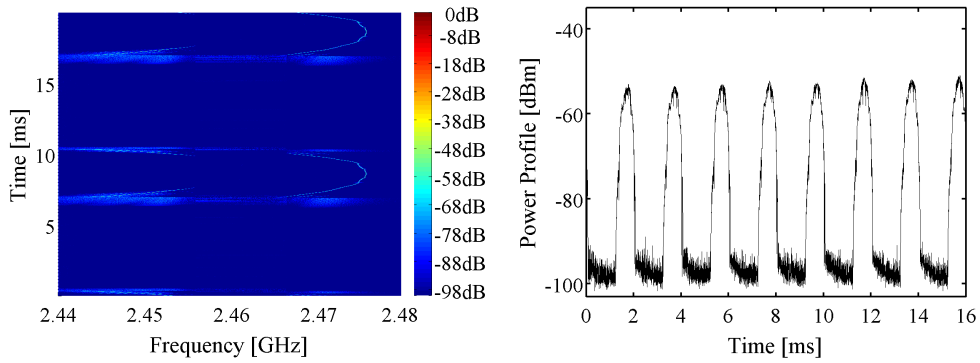

(a) Spectrogram (b) Power-profile

Figure 2.8: Spectrum characteristics of high-power Microwave Oven.

2.3.6 Microwave Oven

Observations. We use a residential microwave oven, the Clatronic MWG 758. As depicted in the spectrogram and power profile in Figure 2.8, the oven radiation distinctly affects the second half of the 2.4 GHz band, and the generated noise exhibits a temporal periodic ON-OFF pattern (~ 5 ms ON, ~ 15 ms OFF). This confirms the observations in [20, 54]. Note that there is still a level of emitted noise in the OFF period that can cause harm to communication parties in close proximity.

In the CCA enabled case, as shown in Figure 2.4(a), short packets at distance 6 m experience slightly fewer losses. This can be attributed to the ON and OFF temporal characteristics of the microwave oven. For distance 3 m, the communication is reduced down to below 20%. As we move the microwave away from 802.15.4 nodes, the PRR improves to reach 90%. For the CCA disabled case, we observe about 20% to 35% corrupted or lost packets, as shown in Figure 2.4(b). More interestingly, for distance 3 m a severe performance reduction is not observed, as with CCA enabled.

2.4 Analysis and Observations

This section provides a detailed analysis of packet error patterns and the temporal channel impairments induced by CTI at the level of symbol granularity in our collected traces. In this analysis, we distinguish between two forms of wireless medium access modalities adopted by radio interferers:

(a) Persistent form: technologies adopting this form consistently emit energy, thus monopolizing the medium completely. Analog cordless phones, as considered in our study, but also analog wireless cameras,

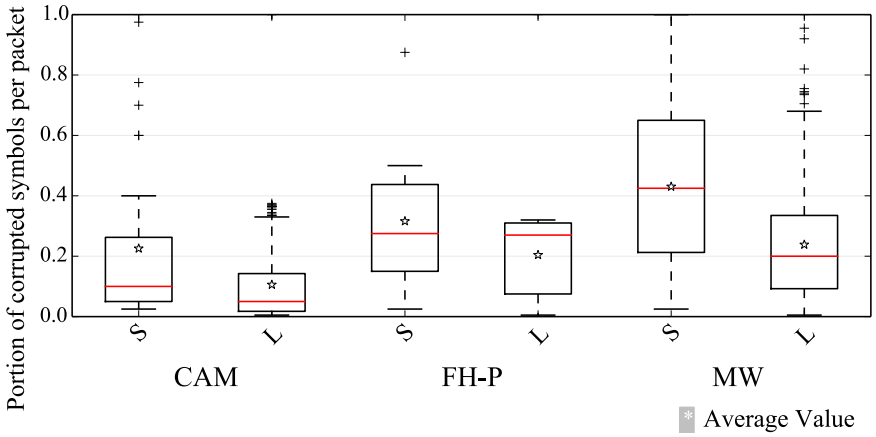

Figure 2.9: Portion of corrupted symbols in a packet, for the CCA-disabled traffic at distance 6 m for packets with length 100 byte (L) and 20 byte (S).

and DSSS cordless phones, adopt such behavior [54]. This form of interference can cause a complete loss of connectivity to the affected nodes, as the medium is constantly detected as busy.

(b) Non-persistent form: the majority of wireless devices operating in the unlicensed bands are non-persistent. They exhibit a time-variant of ON and OFF patterns of energy emission. This is attributed to the underlying adopted access mechanisms by these technologies, such as frequency hopping, continual inter-frame spacing (e.g., SIFS, DIFS), and back-off slots, or periodic ON and OFF cycles of noise radiation, as for the microwave oven. This implies that the occupied wireless channel is idle recurrently. It translates into exchanged packets from active competing transmission being either correctly received (the shorter the transmission time, the higher the chances) or being partially corrupted, where the interfering signal overlaps a segment of the target packet.

In the following, we analyze corrupted packets in our traces with a focus on key features that can potentially aid link-layer recovery mechanisms.

The Rate of Corruption in a Packet. *To what extent do non-persistent interferers corrupt a packet?* The air time of an 802.15.4 packet is in the order of a few milliseconds (max. 4.2 ms), which is often a sufficient time interval to overlap non-persistent interfering signals. This results in having portions of the packets being corrupted, in a way that varies with the time characteristics of the interferer. This insight is potentially helpful for error coding and packet recovery mechanisms.

We explore this aspect further by processing corrupted packets in our traces. The box plot shown in Figure 2.9 depicts the portion of

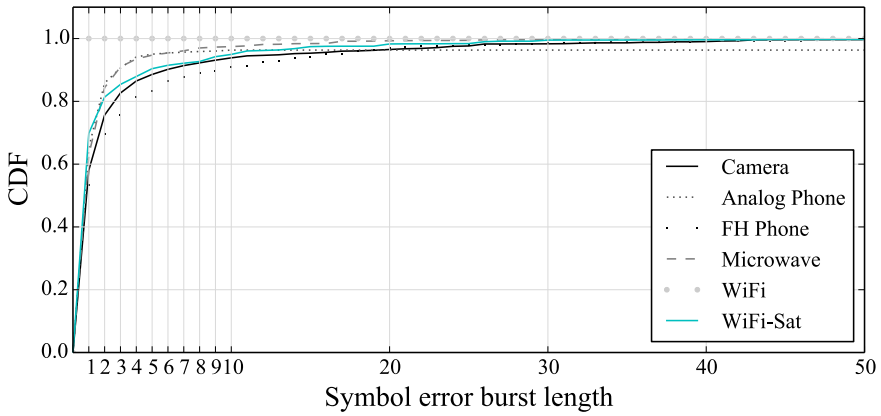

Figure 2.10: CDF of symbol error burst lengths considering all corrupted packets.

corrupted symbols in every packet for the wireless camera, the FHSS phone, and microwave oven. As a result of the technologies being non-persistent, many packets experience corruption over only a minority of their symbols. This is particularly pronounced for long packets: For 100-byte packets, on average less than 25% of the received symbols are corrupted. Such packets could potentially benefit from link-layer mechanisms that rely on physical layer hints to support identifying and recovering corrupted symbols.

Error Burstiness. *To what extent do errors occur in groups, affecting consecutive symbols and consecutive packets?* There is a common assumption that interference errors occur in bursts, thus localized in short intervals, while corrupted bits due to channel variation are nonuniformly scattered. To identify the level of error burstiness due to CTI, we process our traces and count the frequency of symbol error bursts of length n ($n \in [1 \ldots 50]$) with respect to each interferer technology across all packet lengths, power levels, and distances. Note, we make our observation at the symbol level, losing information on the error burstiness in the underlying 32-bit sequence (PN) and making our notion of burst to correspond to a symbol time of 16 μs.

Figure 2.10 shows the distribution of intra-packet burst lengths. The majority of the error burst lengths we processed in our traces are of length 1 for all considered technologies. This can be attributed to the fact that we perform our analysis at the symbol level and burntness is more pronounced at the bit level. We observe that 20-30% of error bursts range in length from 2 to 10 symbols. The wireless camera and the FH phone show a higher tendency of having error bursts of varied lengths.

Error Location. *Where in the packet do most of the symbol errors occur?* For this, we look at the distribution of corrupted symbols over the received 802.15.4 packets. We count for each symbol position how often

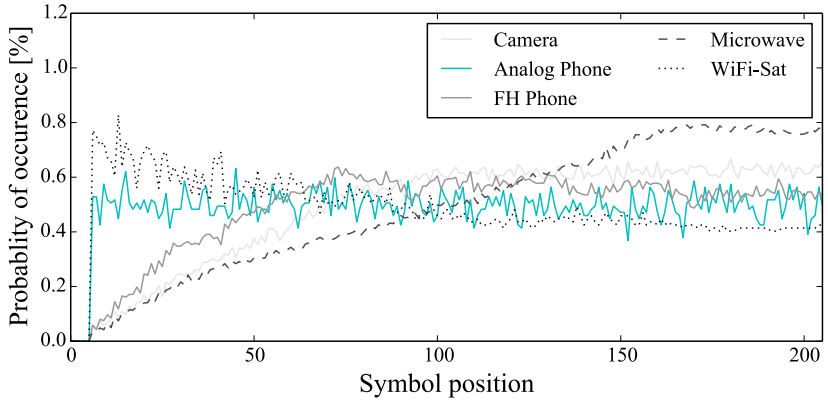

(a) CCA-disabled

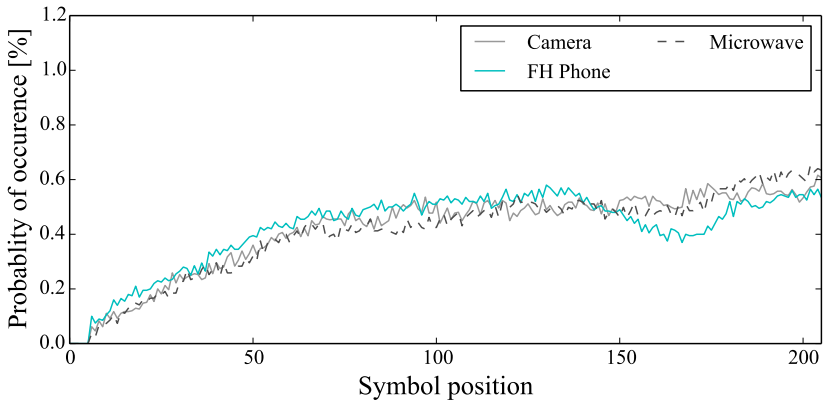

(b) CCA-enabled

Figure 2.11: Symbol error distribution for corrupted 802.15.4 packets (aggregated for packet length 100 byte) interfering with wireless camera, analog phone, FHSS phone, microwave oven, and saturated WiFi.

it was corrupted. We run this over aggregated data of both of the considered distances and transmission powers for the packet size 100 byte. Figures 2.11(a) and 2.11(b) show the probabilities of symbol corruption at different positions in a packet for both communication scenarios with CCA disabled and CCA enabled, respectively. For CCA enabled, we observe fewer corruptions in the header information. For saturated WiFi, we observe a higher chance of corruption in the beginning of a packet, which aligns with the observations of Liang et al. [96]. For persistent interferers such as the analog phone, all positions are affected with similar probability. This changes for channel hopping technologies, i.e., the FH phone and wireless camera. There we observe that the later positions have a higher chance to be corrupted. For microwave, we noticed that the probability increases with the symbol index until index 150 where it stabilizes. This can be attributed to the ON and OFF pattern of microwave and the fact the later positions are affected similarly by the ON state.

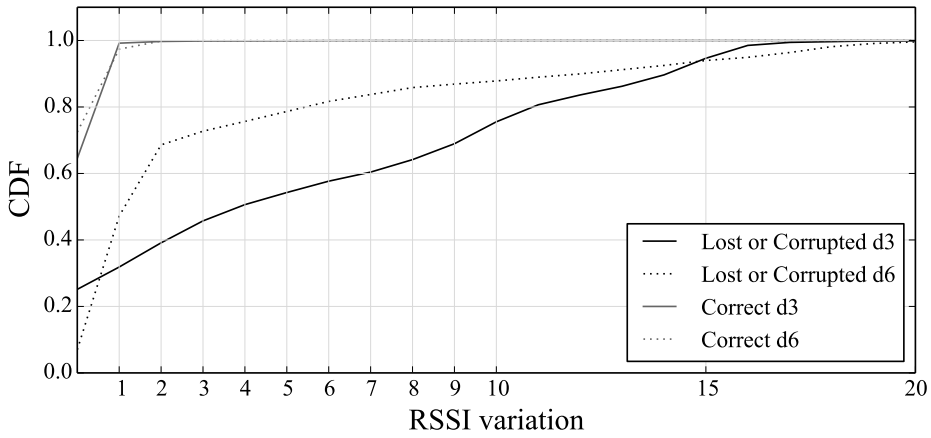

Figure 2.12: RSSI analysis based on RSSI samples during packet reception for two distances, 3 (d3) and 6 m (d6), aggregated for all technologies sending at highest transmission power.

Intra-packet Channel Variations. *How do RSSI readings vary within the span of a packet reception time?* We conduct an analysis to expose statistical differences in the level of RSSI readings during packet reception between interfered, non-interfered, and for weak signal losses. In this context, we check the level of the surge on the RSSI readings during packet reception. Figure 2.12 shows the CDF of RSSI variations for interfered and non-interfered packets. Our observations confirm that RSSI readings vary within two dBm range for the time span of one 802.15.4 frame, considering no interference during packet reception, as the coherence time is larger than one 802.15.4 packet air time [144]. This is mainly why radio chips restrict RSSI readings to few symbols (in the case of CC2420, over the eight first symbols following the SFD field). This consequently leads to missing to capture interference effects. The implication of this is that per packet RSSI and LQI readings do not reflect on the impact of interference. Indeed both LQI and RSSI provide indications of a good and stable channel in most of the interfered packets. Similar observations have been reported for 802.11 in [129]. This can be exploited to detect RF activity, diagnose packet losses, and trigger interference-aware protocols. In Chapter 3, we explain how we harness this observation for interference detection in our system. Other than detecting interference activity, the induced power level on the channel is important for other considerations, such as physical proximity to the interferer [75, 137].

CCA Deferrals and Energy Detection. *To what extent can we rely on static CCA thresholding given CTI presence?* Using our traces from the saturated experiments (CCA disabled), we investigate the relation between the RSSI sampled by the sender before transmission and

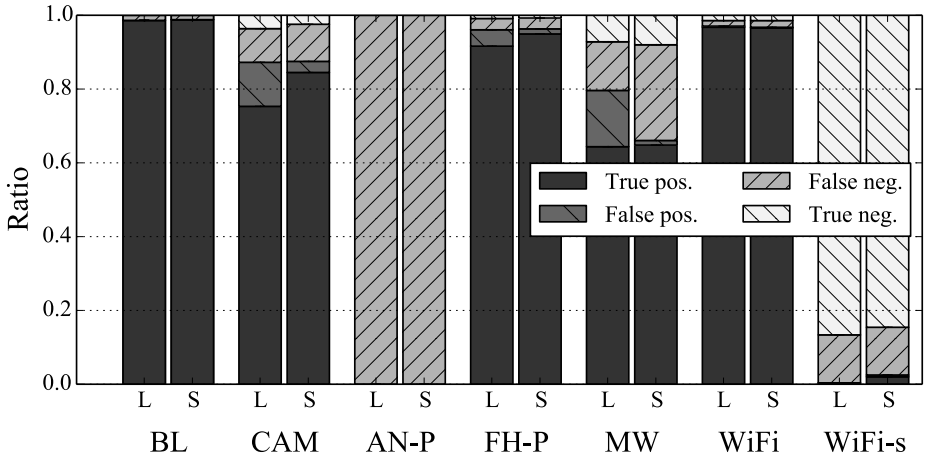

Figure 2.13: Ratio of successful or failed back-off decisions with CCA-disabled traffic at distance 6 m with the highest transmission power for 100 byte (L) and 20 byte (S) packets. Distinction between channel free and transmission ok (true positive), channel free but transmission either corrupted or lost (false positive), channel busy but transmission successful (false negative), and channel busy with either corrupted or lost transmission (true negative).

the actual success of packet transmissions. Doing so, we know, for every transmission, whether a node using CCA would have backed off or not (assuming a threshold of -45 dBm) and whether such back off would have been helpful or not. Figure 2.13 summarizes all possible 4 cases: no backoff followed by success (true positive) or failure (false positive), or backoff followed by success (false negative) or failure (true negative).

In the case of the analog phone, as indicated in Section 2.3.4, the backoff mechanism is extremely inefficient, consistently leading to false negatives (unnecessary backoff). In the saturated WiFi case, on the other hand, the backoff procedure is efficient, avoiding more than 80% of the transmissions that would have failed anyway. For the frequency hopping technologies (Bluetooth, wireless camera, FHSS phone) as well as for the non-saturated WiFi, the channel is sparsely occupied and the backoff threshold operates as intended: few backoffs, and successful transmissions. The microwave oven, with its periodic ON-OFF pattern, is more challenging and presents cases where the backoff is either too conservative or too aggressive.

This analysis shows that (1) a single CCA threshold cannot suit all setups and (2) even for a given setup, a threshold can trigger both false negatives and false positives, e.g., in the case of microwave oven. This indicates that careful CCA threshold tuning on a per-technology basis and agile thresholding would allow better utilization of the radio spectrum.

2.5 Related Work

Radio Interference has been the topic of a large body of wireless communication research. Work in this area falls under two broad categories: The first category deals with the impact of interference on wireless networks, and the second deals with interference mitigation. We focus this section on presenting recent related studies on understanding and quantifying the performance of wireless networks under interference. We further briefly survey work on interference mitigation. We comprehensively cover related work on interference mitigation in the next two chapters.

Interference Empirical Studies. Studying radio usage patterns and interference implications on wireless networks have gained large interest from the wireless research community and industry in recent years. The scarcity of available spectrum and the surge in the number of wirelessly connected devices necessitate a deeper understanding of the characteristics of current spectrum utilization. Large number of independent RF spectrum studies have been conducted by SSC [138], Farpoint Group [46], BandSpeed [15], Miercom [103], Ofcom [7], Jupiter Research [33], Cisco [18], Microsoft [102], and Google [59]. In the following, we highlight two recent studies. Cisco Meraki [18] conducted a large-scale measurement of wireless network's behavior that aims at studying how significant is the interference between 802.11 networks, and to what extent does interference arise from non-802.11 devices. Moreover, they study the occupancy patterns in the frequency bands that 802.11 channels operate in. At a larger scale, Microsoft Spectrum Observatory [102], provides a large-scale system for tracking radio spectrum usage in locations throughout the world. This system was designed to help regulators make more informed decisions on spectrum allocation, push for opening up more frequencies, and help with spectrum-sharing efforts.

With regards to Cross-Technology Interference impact, Srinivasan et al. [143], Petrova et al. [116], Pollin et al. [120], and Sikora et al. [140] have performed experimental studies to quantify the impact of interference from 802.11, Bluetooth, and microwave oven on the performance of 802.15.4 networks. These studies focused on reporting the impact on performance metrics such as throughput and packet reception ratios, however, without exploring low-level effects of interference. Liang et al. [96] studied the interplay between 802.11 and 802.15.4 networks and their patterns at bit-level granularity focusing on bit-error positions. They recognize symmetric and asymmetric interference regions. Boano et el. [20] studied interference patterns with the focus on

the coarse samples of the RSSI for the purpose of emulating interference patterns in testbeds. To the best of our knowledge, our work is the first Cross-Technology Interference characterization study that aims at providing a detailed understanding of the interaction between 802.15.4 devices and a set of prevalent RF interferers and recognizing key factors to the harmful coexistence of these technologies.

Interference Mitigation. The recognizable impact of RF interference on the performance of wireless networks has motivated researchers to look at solutions to mitigate interference. The most widely adopted mitigation solution is to avert interferer frequencies by employing spectrum sensing to identify interference-free channels [126, 162, 34]. Such approaches are resource hungry for 802.15.4 networks. Moreover, the spectrum is crowded with wireless devices which makes it hard to find interference-free channels. Another direction of research focuses on the recovery from symbol corruption, by utilizing resilience coding schemes that are robust to bursty errors. For instance, Reed-Solomon coding can be employed to mitigate the 802.11 impact on 802.15.4 networks, as suggested by Liang et al. [96]. Furthermore, partial packet recovery mechanisms are used to exploit the temporal effects of interference induced on the PHY hints, such as variations in soft errors (softPHY) [85] or RSSI variations [71, 66] to determine boundaries of the interfered fractions on the received corrupted packets.

2.6 Summary

This chapter reports and discusses results of our empirical study of the Cross-Technology Interference impacts on 802.15.4 links. We examine the interaction patterns between 802.15.4 links and a set of prevalent high and low-power radio interferers at symbol level granularity with the focus on protocol aspects, error patterns of bits transmitted over the air, and the wireless link variations as perceived by the transmitter and receiver. All our observations in this chapter are based on metrics that are exposed to off-the-shelf nodes and can easily measure network characteristics.

We show that radio interferer technologies differ widely in the way they affect 802.15.4 networks. They form a strong and complex impact on the performance of wireless links that need to be addressed with novel solutions that exploit channel information and physical layer hints. One important conclusion of this study is that there is no one-fits-all solution to mitigate the impact of Cross-Technology Interference. We need to address this by designing novel measures that take into account the properties of the interferers to adaptively select a proper mitigation mechanism.

The observations made in this chapter largely influenced the work blocks of this dissertation. In the rest of this dissertation, we harness the insights and observations collected in this chapter to develop and design adaptive lightweight systems that are apprehensive of the uncoordinated wireless coexistence problem. Hence, overcome the limitations of the static suboptimal mitigation solutions and interference source classification approaches, that we can not yet utilize in a systematic manner for mitigation.

3

TIIM: Technology-Independent Interference Mitigation

The ubiquitous and tetherless access to information enabled by the wireless medium, and recent advances in wireless communication, have led to a plethora of heterogeneous wireless devices congesting the unlicensed bands. This raises a unique set of communication challenges, notably coexistence, Cross-Technology Interference (CTI), and fairness amidst high uncertainty and scarcity of interference-free channels (see Figure 3.1).

The lack of interference-free channels led researchers to work on developing novel classification approaches that provide information about the interference source [130, 68, 78, 34]. It has been shown that when the interference source is known, specialized mitigation approaches can improve the network performance. Wireless technologies employ different physical and MAC layer schemes leading to hidden, distinct, and repeating patterns that form a signature for that particular wireless protocol. Researchers explored these properties to build interference classification tools that can report on the root source of the interference problem. This approach yields interesting results but is bound to a fixed set of interfering technologies that are known at design time. Moreover, these approaches are anticipated to be integrated into spectrum analysis tools that can help network administrators to visualize and identify interference problems. However, they are inadequate if they do not resolve interference instantaneously in a systematic manner. These approaches provide a compelling evident that miniaturized intelligence is necessary to address the complexity of the CTI problem.

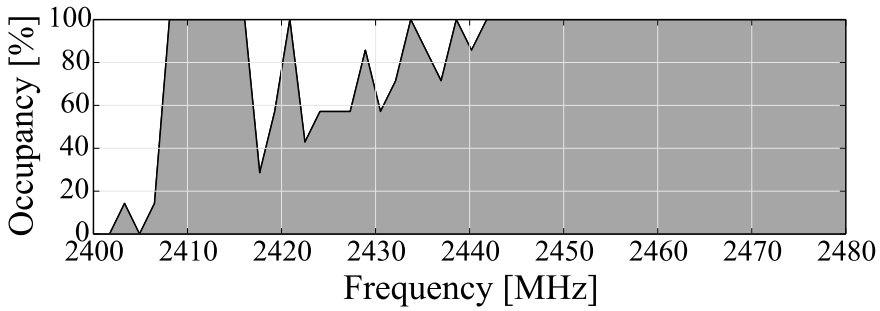

Figure 3.1: Averaged channel occupancy in the 2.4 GHz band over one week (26.-31. August 2013). Data from Microsoft Spectrum Observatory in an enterprise building in Brussels, Belgium.

However, the benefit of such solutions would be realized only if they resolve CTI without manual tuning or human intervention. How to build such smart radios remains an open research problem. In this chapter, we investigate the feasibility of integrating miniaturized intelligence in radios to combat CTI.

Looking at the design space of spectrum coexistence solutions, and based on the observations we made in Chapter 2, where we empirically studied the impact and the interaction patterns of CTI on low-power wireless networks, we conclude that to address the high uncertainty of CTI we need to design agile methods that assess the channel conditions and apply actions maximizing communication success. To achieve this one should consider the following aspects when addressing the CTI problem: *(i)* PHY-aware protocols: physical signals do not only encode bits, but also carry rich information about the ambiance, which is particularly enlightening in the case of interference. *(ii)* The presence of interference is not always harmful, metrics such as energy detection can falsely trigger the communication to back off and introduce unnecessary deferrals (see Section 2.4 for a detailed discussion). Thus, it is important to consider measures that can better quantify the harm of CTI. *(iii)* Given the scarcity of the frequencies allocated to wireless networks, it is desirable to allow concurrent transmissions that potentially can be correctly recovered to realize a better utilization of the spectrum. *(iv)* There is no one-size-fits-all solution. The high degree of diversity in radio technologies results in different implications on the wireless link that need to be addressed with different strategies. *(v)* The impact of the same source of interference can quickly change due to mobility (e.g., interferer moves away) or due to change in the configuration (e.g., WiFi bit-rate, or application traffic pattern). Thus, frequent adaptation is required.

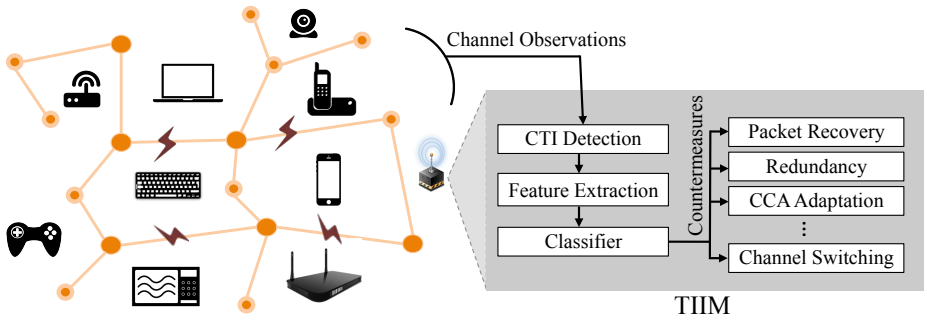

Figure 3.2: Low-power communication links suffer from crowded ISM bands. TIIM dynamically applies interference mitigation measures specific to channel conditions.

Contributions and Roadmap. In this chapter, we present TIIM, as illustrated in Figure 3.2, an adaptive interference mitigation system, which selects interference mitigation strategies directly based on measured medium properties, skipping the interference type classification step. Hence, TIIM is independent of the interference technology it is combating. To this end, we train the system to detect interference patterns and map these to a link-layer interference mitigation strategy that works best for this particular pattern, regardless of the interference type. In TIIM's design, we consider: *(i)* exploring the feature space of Cross-Technology Interference, *(ii)* constructing a lightweight decision tree classifier that learns the conditions where particular countermeasures perform best and uses this knowledge to select countermeasures for unseen channel instances at runtime.

In the rest of the chapter, we elaborate on TIIM's key intuition, empirically characterize and analyze CTI patterns on 802.15.4 links, and present TIIM's design, followed by performance evaluation. We conclude this chapter with a discussion of the limitations and opportunities of designing CTI-aware and adaptive link-layer protocols. This chapter is based on the contributions made in [77].

3.1 Background

In this section, we cover some background on wireless interference in the unlicensed bands.

3.1.1 Cross-Technology Interference in a Nutshell

The broadcast nature of the wireless medium makes it inherently vulnerable to interference from spatially close concurrent transmissions

that overlap in time and frequency. This can consequently reduce or even prevent completely the ability of receivers to decode information from signals. Wireless communication can be subject to disturbance by interference from intra-technology, cross-technology, or noise sources. Wireless technologies strive to avoid interference and typically apply a set of mechanisms to achieve fairness and reduce interference within the same technology (e.g., reserve the medium, allocate channels, and probe for idleness). However, most protocols are not designed with coexistence in mind. This is mainly because of the infeasibility of interference coordination due to the absence of communication means between these diverse technologies (i.e., speak different PHY protocols). Consequently, CTI is emerging as a major problem in the unlicensed bands [7, 76, 54, 103].

Unlicensed Spectrum. The unlicensed bands are small segments of the radio spectrum that were reserved internationally for the use of RF energy for *Industrial, Scientific, and Medical* (ISM) purposes and have been widely utilized for unlicensed short-range wireless radios. The *Electromagnetic Compatibility* (EMC) regulators, such as the Federal Communications Commission (FCC) in the United States and the European Conformity (CE), generally require a license for the use of airwaves for communication, except for some frequency bands that they leave open. Although there is no permission necessary for devices to operate in the open frequency bands, they have to comply with few technical requirements, including power limits. There are several available unlicensed bands, but most devices primarily operate in the 2.4 GHz, the 900 Mhz, or the 5 GHz frequencies. Currently, the 2.4 GHz band is the most crowded band within the unlicensed radio spectrum.

In the following, we spot the light on certain communication properties that are adopted by many radios, and make it particularly challenging for low-power technologies such as 802.15.4 to coexist in the shared spectrum: *(i)* Wide-band: many devices transmit in frequency bands significantly wider than 802.15.4. For example, to cope with the high demand of high throughput over WiFi, the recent amendments of 802.11 allow the configuration of 40 MHz-wide channels in the 2.4 GHz band. Wireless systems are in general increasingly moving to wider frequency bands to cope with the high throughput demands. Another case of wide-band occupancy is microwave ovens; they typically affect 50% of the available 2.4 GHz band. *(ii)* High-power: today's high-power interferers in the unlicensed bands pose a serious threat to 802.15.4 networks, as they can cause 802.15.4 links to experience a complete loss of connectivity. Although the EMC regulators lightly control this aspect by setting an upper limit of 30 dBm for transmit-power in the unlicensed bands, energy leaks from microwave ovens can reach up to 60 dBm.

This is significantly higher than the typical output power of 802.15.4 radios which is 0 dBm.

Communication Primer. We briefly recall how signals are transmitted and received over the wireless channel. The following assumes *Minimum-Shift Keying* (MSK) signal. Note that we omitted unnecessary details to simplify this communication primer. Radios convert binary data into modulated signals. These signals are generally represented as a discrete and complex function:

$$s[n] = A_s e^{i\theta_s[n]} \tag{3.1}$$

where A_s is the amplitude of the transmitted sample n, $\theta_s[n]$ is its phase. Note since MSK embeds all the information in the phase, A_s is constant for all samples. Hence, the signal carries constant energy.

After the signal traverses the channel, the receiver receives:

$$y[n] = Hs[n] + w[n] \tag{3.2}$$

where H is a complex number that approximates the effect of the wireless channel (attenuation and phase) from the transmitter's antenna to the receiver's antenna [149], $w[n]$ is the channel noise. In the presence of an unknown interferer, i.e., the desired signal interfered with unknown signal, the signal at the receiver is represented as follows:

$$y[n] = Hs[n] + i[n] + w[n] \tag{3.3}$$

where $i[n]$ is the interfering signal.When these two signals interfere, their energies add up (i.e., signals can add up constructively or destructively based on their phase alignment), Hence, causing perceptible variations in the received signal strength. This insight on the additive energy of interfering signals highlights the ambient information the signal carries along, and that can assist in detecting interference and localizing interfered symbols within interfered packets.

3.2 TIIM Overview

We now present a high-level overview of our interference mitigation system TIIM. The intuition underlying TIIM's design is that each interference mitigation approach works well under specific channel assumptions of error patterns, such as error rate, signal to interference ratio, or occupancy level. Each interference instance, independent of the technology, leaves a particular signature in the channel that shapes the channel properties in a unique way. The goal of TIIM is to automatically

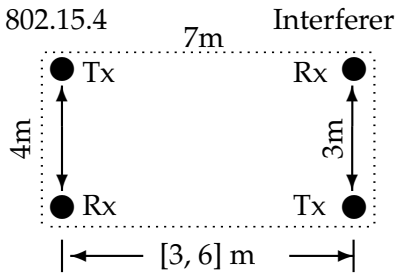

Figure 3.3: Controlled experiments setup for CTI characterization in an anechoic chamber.

select at runtime the most effective mitigation strategy for the current interfered channel and get the best out of the interfered communication link. The core component of TIIM is a lightweight decision tree classifier that is trained to learn under which channel conditions (i.e., signatures) a particular mitigation scheme empirically works the best. TIIM uses the decision tree online to predict the best channel mitigation strategy for yet unseen instances of an interfered channel. TIIM's design consists of the following steps:

CTI Characterization. The first step in the design of TIIM is to characterize *Cross-Technology Interference* (CTI). We run experiments where we expose an 802.15.4 communication link to various types of interferers, both in an anechoic room and in an office environment. We collect channel properties and communication statistics at high frequency. Section 3.3 presents the results of this characterization in details.

Learning Phase. We simulate every considered mitigation strategy against the traces collected in the characterization step, and compute both their gain and cost. In this phase of supervised learning, the decision tree classifier learns for each channel feature which particular mitigation strategy scores highest.

Runtime. At runtime, nodes monitor their current channel condition mostly through signal strength sampling at high frequency during packet reception. Whenever interference is detected, they feed the decision tree with channel statistics as input and obtain a decision about the mitigation strategy to employ.

TIIM is inspired by the core idea behind interference classification approaches, such as SoNIC [68] and Airshark [130], which use measurement samples drawn from commodity hardware to detect the type of interference source. While these approaches can provide useful information on how to potentially mitigate CTI, they are bound to a fixed set of interfering technologies that are known at design time and cannot combat CTI autonomously. They require either user intervention or querying a central entity that maintains the mapping of an interference

source to the corresponding countermeasure. It is yet not clear how such approaches can be utilized in an automated way. TIIM departs from the above in that it skips the interference classification step. Instead, it infers the best mitigation strategy from channel properties directly and independent of the technology causing the interference.

3.3 Characterizing Cross-Technology Interference

In this section, we extend the analysis of characterizing how arbitrary interfering signals interact with 802.15.4 communication, which we presented in Chapter 2. We focus here on identifying distinct features of interfered channels and packets that are relevant to the effectiveness of potential CTI countermeasures and, hence, could assist in: *(i)* detecting and quantifying interference, *(ii)* pinpointing the viability of opportunistic transmission in interfered channels, *(iii)* selecting the countermeasure that works best for the current underlying interference patterns.

3.3.1 Controlled Experiments Setup

We run our experiments in an anechoic chamber, in order to have full control on the sources of errors, type of channel distortions, and to isolate the impact of surrounding interference sources. We consider a simple network setup, as depicted in Figure 3.3, which consists of one transmitter and one receiver for both 802.15.4 and the considered interfering technology, i.e., a pair of 802.15.4 nodes and a pair of interferer nodes. We base our sender and receiver applications on Contiki OS [39] and directly interface them to the node's radio driver. We consider different traffic patterns and different configurations of packet length and transmission power.

Interfering Technologies. We focus on a set of interferer technologies that are prevalent in today's environments. Our considered set consists of low/high power, narrow/wide band, analog/digital, channel hopping/fixed frequency, and CSMA/non-CSMA interferers. This represents common underlying properties adopted by most radio technologies. Figure 2.3 summarizes the features of the considered RF technologies in our study. In the following, we briefly highlight some of their properties.

• *IEEE 802.11.* We create WiFi interference using a Netgear WNR3500L router and a laptop that supports IEEE 802.11 b/g/n in the 2.4 GHz ISM band. We use the network tool `iperf` [82] to generate saturated TCP

traffic and non-saturated UDP traffic that resemble file download and VoIP, respectively.

• *Bluetooth.* To evaluate the interference generated by Bluetooth on 802.15.4, we use two HTC Desire phones transferring a large file. Bluetooth uses the adaptive frequency hopping technique across 79 MHz of bandwidth in the 2.4 GHz ISM band. The hopping occurs at a rate of 1600 hops/s. Hence, it occupies a 1 MHz channel for 625 μs.

• *Digital Cordless Phone (FHSS).* We experiment with the Uniden DCT6485-3HS cordless handset system. The phone base and handset communicate using frequency hopping over 90 channels of 800 kHz width in the range [2407.5 - 2472] MHz.

• *Analog Cordless Phone.* We experiment with the Vtech GZ2456 cordless handset system. The phone base transmits in the 900 MHz band and receives in the 2.4 GHz band. The phone handset accordingly transmits and receives using the reverse order of frequency ranges.

• *Wireless Camera.* We use the Philips SCD 603 digital video baby monitor. It comprises a 2.4 GHz wireless camera and a wireless video receiver. The wireless camera uses frequency hopping over 61 channels, where each channel has a width of 1.125 MHz.

• *Microwave Oven.* We use a residential microwave oven, the Clatronic MWG 758. We heat a cup of water in the microwave to emulate an interference typical to that emitted by these appliances.

3.3.2 Interference Detection

Interference detection is a key component for addressing interference. Performance degradation in wireless systems can be due to *interference* or *insufficient signal strength*. Determining the cause of performance degradation is an essential input for addressing the problem as this defines the corresponding mitigation action to be considered. In the following, we investigate how off-the-shelf 802.15.4 radios can detect whether a received packet has been subject to interference or not by exposing differences in channel metrics between packets corrupted by interference and packets corrupted due to insufficient signal strength (i.e., weak signal).

The 802.15.4 transmitted signal encodes information in phase rather than amplitude. Therefore, its amplitude (i.e., energy) is constant within the coherence time. When two signals interfere, their energies add up, consequently impacting the energy level of the interfered segments in the received signal. Confined with the PHY and link layers of the OSI stack, the standard design of off-the-shelf radios treats the PHY layer as a black box that provides decoded bits (i.e., MAC layer PDU) and a limited PHY information and deprives the access to signal level information.

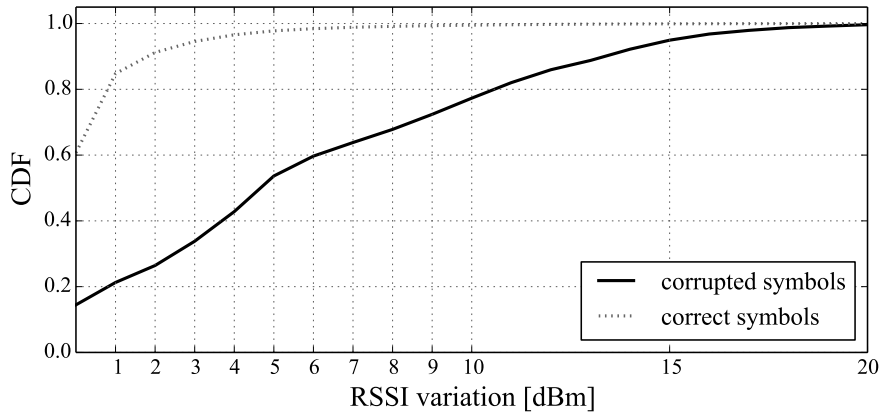

(a) CDF of RSSI variations for correct and corrupted symbols due to interference.

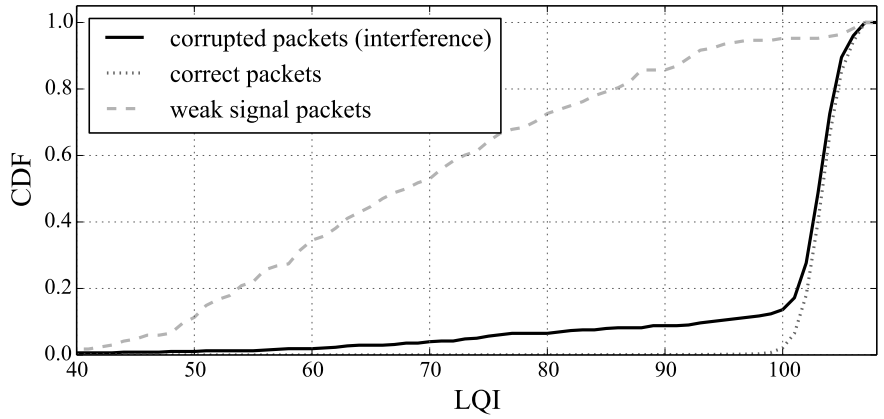

(b) CDF of per-packet Link Quality Indicator (LQI).

Figure 3.4: Observations from our traces on interference detection. (a) shows that interfered packets often experience high RSSI variations. (b) shows the clear distinction between LQI of corrupted packets due to interference and those corrupted due to weak signal.

This leaves us to work with the limited available PHY information to design an interference detection mechanism. Note that in Chapter 4, we explore this direction further using software radios (i.e., richer PHY information) to realize a better detection.

For detecting interference using off-the-shelf 802.15.4 radios, we explore two different possibilities:

I. Capturing energy variations during packet reception by sampling the radio's RSSI register. We modified the CC2420 driver in Contiki [39] to capture RSSI values at a rate of one sample per symbol (i.e., one reading each 16 μs). The sampling is performed from the *Start of Frame Delimiter* (SFD) to the last symbol of the packet. Figure 3.4(a) depicts the CDF of the RSSI variations of

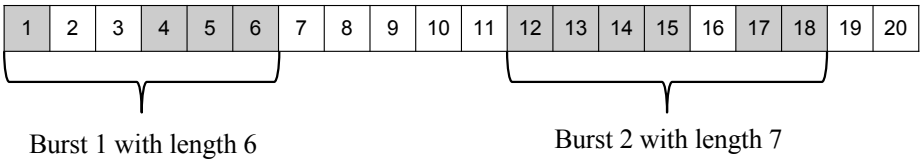

Figure 3.5: An error burst consists of co-located corrupted symbols with less than 5 correct symbols in between. Gray and white symbols represent corrupted and correct symbols, respectively. The first error burst starts at symbol 1 and ends at the 6th symbol and it has a length of 6. The second burst starts at the 12th symbol and ends at the 18th symbols.

interfered erroneous packets and correct packets as observed in our traces. We see a clear correlation between RSSI variation and packet corruption. For instance, 90% of the non-interfered packets have a variation under 2 dBm, while more than 70% of the interfered packets experienced variations higher than 2 dBm. In case the variance of RSSI is greater than a threshold (default to 2 dBm), our system recognizes the received packet as an interfered packet.

II. The second alternative is to exploit how per-packet channel metrics are computed in off-the-shelf radios. The Link Quality Indicator (LQI) is confined to average the readings of few symbols (in the case of CC2420 [147], over the first eight symbols following the SFD field). Consequently, it cannot capture the spike in the received power due to interference, as long as the interference spike does not fall within these few symbols. As a result, per packet LQI readings do not reflect the impact of interference. In fact, LQI provides indications of a good and stable channel in most of the interfered packets. Similar observations have been reported for 802.11 in [129]. Figure 3.4(b) highlights this insight on LQI considering weak signals and interfered signals. It shows that 85% of corrupted packets due to weak signal have an LQI of about 90 or less, whereas only 10% of packets suffering interference have an LQI of 85 or less. To detect interference, the system can monitor the LQI of received corrupted packets. Frequent erroneous packets with good link metric (LQI > 90) could be used to detect interfered packets.

Interference mitigation schemes come with an overhead that should be avoided in the absence of interference. Hence, we exploit the LQI approach of interference detection to trigger the recovery phase in TIIM, since it comes at a lower overhead. In the recovery phase, we utilize the more reliable RSSI variation mechanism to pinpoint the interfered packet.

3.3.3 Exploring the Feature Space

We explore the feature space of potentially relevant attributes to the CTI problem that we later use for constructing the classifier in TIIM. We focus on features that can reflect on the occupancy and error patterns in interfered channels. Table 3.1 presents the list of features we utilize in TIIM's classifier. In the following, we highlight the relevance of these features to CTI.

Persistency. The primary spectrum usage modalities that exist in the unlicensed bands are comprised of: *(i)* Persistent. Technologies adopting this form operate in dedicated frequency bands and generate static energy, thus monopolizing the medium completely. Legacy analog devices adopt this modality. This form of interference causes a complete loss of connectivity for the interfered low-power nodes [76], primarily because the continual energy emission prevents the carrier sense from declaring the channel to be free.

(ii) Non-persistent. Technologies adopting this form operate either in dedicated frequency bands and generate traffic with time varying load, or exploit frequency diversity and hop across the spectrum. They exhibit a time-variant ON and OFF pattern of energy emission due to underlying communication patterns, such as frequency hopping, continual inter-frame spacing (e.g., SIFS, DIFS), back-off slots, and varying load, or periodic ON and OFF cycles of noise radiation, as for the microwave oven. This translates to exchanged packets being either correctly received (i.e., the shorter the transmission time, the higher the chances) or partially overlap with the interfering signal leading to packet loss or corruption. It is clear that even at this level, realizing whether the interference is persistent or non-persistent should be followed by adopting different mitigation schemes.

Properties of Corrupted Packets. Features that can represent an estimate of the error bit-rate and error patterns in packets can serve as an important meta-information for error recovery mechanisms. We define a nominal feature that can take four values representing different classes of error patterns. C_0: error rate > 33%, C_1: few corrupted symbols, between 1 and 12 corrupted symbols, C_2: one error burst not larger than 33% of the payload; we define an error burst as co-located corrupted symbols with less than five correct symbols in between (see Figure 3.5), and *others*. Figure 3.6 depicts the ratio of each of these error classes aggregated over the traces from the controlled experiment.

To compute this feature, we can either rely on retransmissions to identify the corrupted parts of a packet or, as illustrated in Figure 3.7, by analyzing energy surge bursts in the received sampled RSSI as these

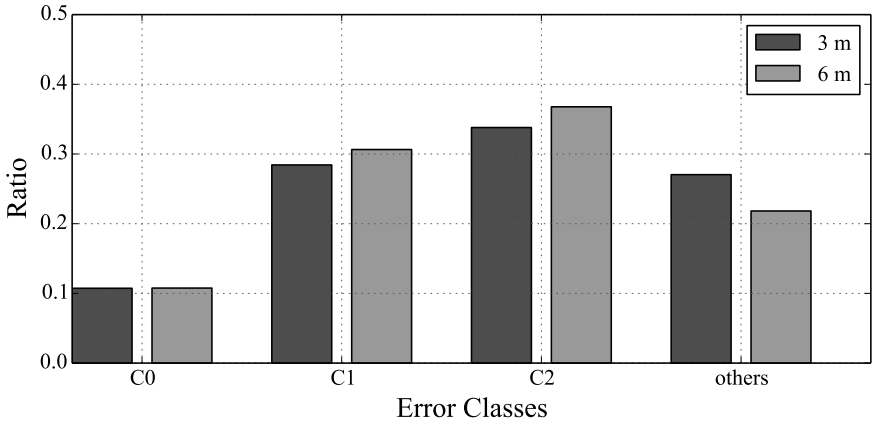

Figure 3.6: Error classes from CCA-enabled traces averaged over all sources of interference at distance 3 m, and respectively at 6 m. C0: corrupted symbols larger than 1/3 of the payload's length, C1: few (1-12) corrupted symbols (suitable for FEC), C2: single error burst not longer than 1/3 of the payload length (suitable for packet merging).

surges correlate to corrupted parts of a packet in case of missing or corrupted retransmissions.

Interference Quantification. The mere presence of interference is not always harmful. Hence, finding metrics that can better quantify the actual impact of interference can largely influence the way we address interference and potentially increase spectral efficiency. We define a metric that considers the reception status of packets during an observation window. Assuming a total number n of packets transmitted during the observation window, n_i is the number of interfered packets, n_s is the number of corrupted packets due to other channel impairments (e.g., weak signal), and n_l is the number of lost packets. The estimated CTI impact is: `estimated_interference` $= (n_i + n_l)/n$. As we cannot clarify the source of lost packets, i.e., packets that the receiver failed to detect the preamble of, we take a conservative approach and account them as impacted by interference.

3.3.4 Countermeasures

In the following, we briefly cover a set of link-layer mitigation schemes that we consider in the design of TIIM. These mitigation schemes have been proposed and evaluated in the literature in the context of increasing the resilience of 802.15.4 against interference. TIIM is not bound to this set of countermeasures, and can be trained and extended with further countermeasures. In the current prototype of TIIM, the classifier is trained to select one of the following mechanisms or a combination of them:

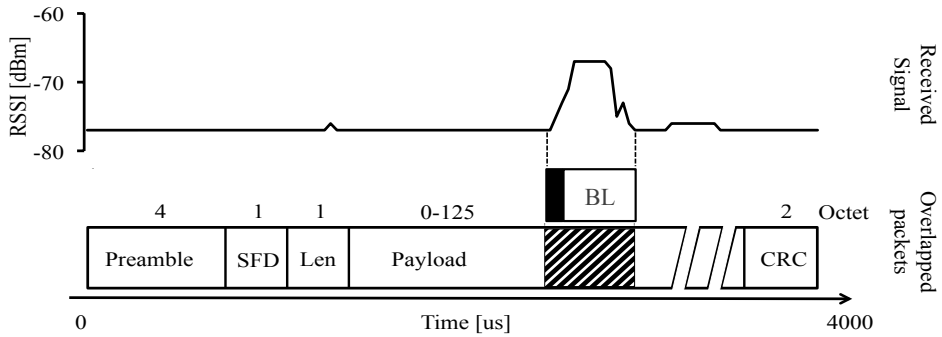

Figure 3.7: An example of the correlation of signal variations with the symbol errors within a received frame. TIIM makes use of this knowledge to request the retransmission of only the co-located corrupted symbols.

- *Reed-Solomon Forward Error Correction (FEC):* The Reed-Solomon (RS) code is a block-based error correcting code that is particularly effective at correcting burst errors. RS code divides a message m into n blocks of defined size and adds an extra redundant parity of t blocks to the message. RS code can correct up to $t/2$ and detect up to t block errors. The overhead cost of RS code is constant, both correct and corrupted packets bear the redundancy overhead. RS code works well for error patterns that fall under the recovery capacity of the parity check. In TIIM, we use 12 Bytes of parity.

- *RSSI-based Packet Merging (PM):* In the presence of interference, a sender often has to retransmit a packet several times until the receiver decodes a correct copy. Partial packet recovery [85, 66, 63] and packet merging [56, 71] aim at reducing the amount of redundantly received data in such cases, by reconstructing packets from already received corrupted instances. For instances, with long error bursts (beyond RS recovery capability) or with sparse erroneous packets, PM is a good countermeasure candidate. As depicted in Figure 3.6, about 40% of the corrupted packets we witnessed in the controlled study traces, are of class C_2 (it varies for different interferers). This class of errors stands for a single error burst in the corrupted packet, where the rest of the packet is error-free and would potentially gain from packet recovery mechanisms.

  In PM, the receiver selectively requests the segments of a packet where symbols are likely to be corrupted [66]. The segments are identified based on the surges in RSSI readings. Figure 3.7 illustrates an example of RSSI variations (from our controlled experiment) during the reception of an 802.15.4 packet. It shows how the surge

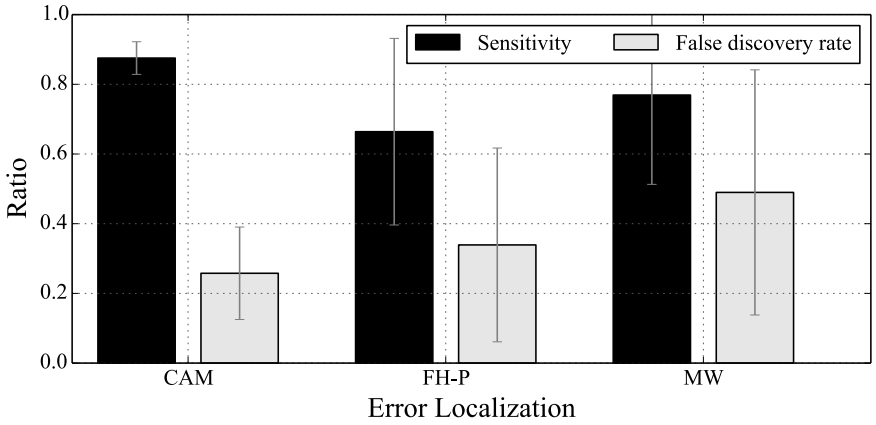

Figure 3.8: Error localization for CCA-disabled traces from the anechoic chamber. TIIM uses fine-grained RSSI sampling to localize the area of corrupted symbols in the payload. Sensitivity is a metric indicating the ratio of correctly detected corrupted symbols, whereas the false discovery ratio shows the ratio of correct symbols among the detected symbols. Different technologies experience varying bit-error localization performance.

in RSSI corresponds to the error location due to interference by Bluetooth. Figure 3.8 depicts the accuracy we achieved in localizing corrupted symbols by utilizing sampled RSSI in our traces.

- *Adaptive CCA Thresholding (no-CCA):* 802.15.4 networks generally use carrier sensing before transmission, to reduce collisions. We evaluate the efficiency of carrier sense in 802.15.4 radios under CTI, from our saturated CCA disabled experiment traces. We investigate the relation between the sampled energy at the sender before transmission and the actual success or failure of packet transmissions. Figure 3.9 summarizes the results of carrier sense efficiency experiments. The carrier sense works well in many scenarios, but may lead to false positives (i.e., channel free but transmission either corrupted or lost) and false negatives (i.e., channel busy but transmission successful). The following scenarios are worth looking at: *(i) Frequency Hopping* (FH) interferers: In the presence of frequency hopping interferers, carrier sense is not effective. FH interferers do not react to 802.15.4 transmissions. Accordingly, the chances a packet encounters corruption given channel is sensed free or occupied is the same. *(ii)* Analog interferers: In the case of the analog phone at distance 6 m, although the energy level in the channel is high, PRR is barely affected by interference. Hence, carrier sense causes unnecessary deferrals. In Figure 3.9, we observe that for the analog

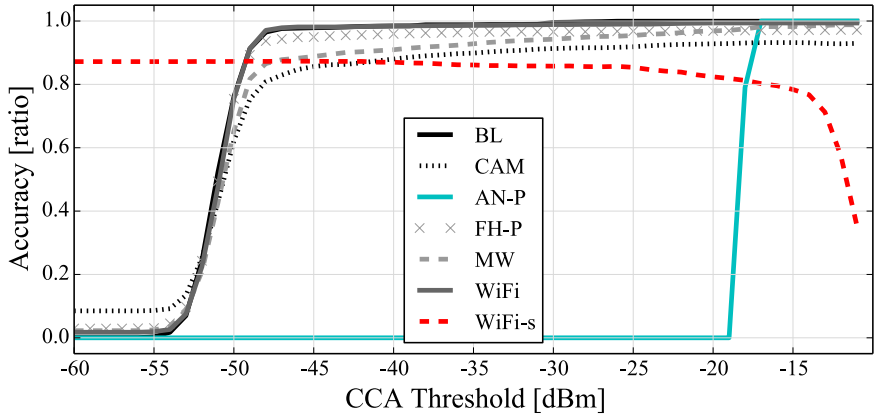

Figure 3.9: Accuracy of different CCA thresholds for anechoic room traces with interferers at distance 6 m. Accuracy is defined as $\frac{\text{TP+TN}}{\text{\#transmissions}}$. True Positives (TP) are the cases where the channel was free and the transmission successful. True Negatives (TN) indicate cases where the channel was busy and the transmission was lost or corrupted. A fixed CCA threshold does not serve well under all channel conditions.

interferer (AN-P) which transmits with high power, any threshold below -18 dBm results in 100% false negatives. Thus, an adaptive CCA scheme that can assess and prevent harmful concurrent transmissions, while allowing safe concurrent transmissions, could largely enhance spectral efficiency. TIIM follows rather a radical strategy and recommends *no-CCA* when it detects that using CCA is causing high false negatives.

- *Channel Switching* or *No Action:* TIIM can infer that the interference in the channel is not harmful thus no action is required or that the interference in the channel is severe thus communication over this channel is not viable even with the assistance of link-layer interference mitigation mechanisms. In such situations, it gives the recommendation of channel switching.

3.4 TIIM Architecture

So far we have concentrated on describing TIIM at a high-level and discussing our design decisions and some essential empirical observations on the 802.15.4 interfered channels. We now provide an overview of TIIM's components and its operation modes before detailing the classification algorithm and discussing TIIM's integration into the system. All aspects of TIIM have been carefully chosen and designed with runtime and memory efficiency in mind. We focus on

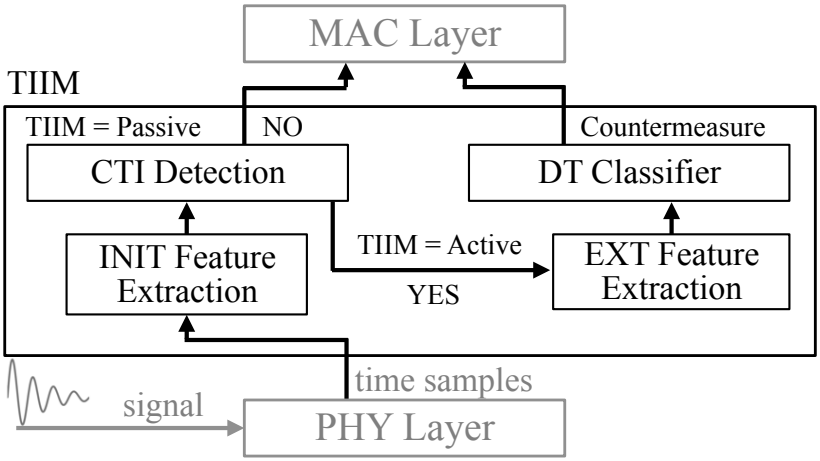

Figure 3.10: TIIM's Design. TIIM remains passive while observing channel conditions, i.e., initial (INIT) features. Upon detection of interference, it turns active and collects further channel metrics, i.e., extended (EXT) features, for a given time window and inquires the *Decision Tree* (DT) classifier for a countermeasure.

realizing the following four primary goals: *(i)* Improving *spectral efficiency* and *packet reception ratio* in the presence of CTI. *(ii) Compatibility*: we design TIIM such that it can be implemented as software modifications on top of commodity hardware. *(iii) IEEE 802.15.4 PHY compatibility:* compliance to the existing standards which allows seamless integration into existing systems. *(iv)* Supporting *heterogeneous CTI patterns*, oblivious to interference source type, distance, or configuration.

3.4.1 Modes of Operation

TIIM operates in two modes: passive and active. Having two modes allows avoiding imposing additional computational overhead to interference-free communication. The system runs mainly in the passive mode which monitors a number of channel metrics, namely, the *Link Quality Indication* (LQI) value of corrupted frames at the receiver side, and packet losses and CSMA deferrals at the sender side. TIIM detects harmful interference using a simple threshold mechanism (τ_{active}) (see Section 3.3.2). Upon detecting harmful interference, TIIM switches to the active mode. An overview of TIIM is shown in Figure 3.11.

TIIM's active mode operates on a window of communication events W_{active}. In our experiments, we use a window length of 5 seconds. We found this to be a good tradeoff between time to react to interference and confidence of selection. Note that the time window length can be adapted to the level of activity in the channel, and its optimization is out of the scope of this work.

During active mode, TIIM continuously processes time samples collected during packet reception and communication statistics for the time of the observation window and computes a set of channel features. Then, the node feeds the decision tree classifier with the features and triggers the mitigation strategy inferred by the classifier. Finally, TIIM switches to the passive mode to monitor the performance of the activated mitigation strategy.

3.4.2 Inferring Countermeasures

To infer the best (set of) countermeasure(s) for a given input feature set (see Section 3.3.3 for the feature space discussion), we use a supervised learning approach for classification. At runtime, the trained classifier assigns unseen instances of interfered channel, i.e., observation window W_{active}, to one of the six output classes. The classes represent the set of mitigation strategies we consider in this prototype of TIIM (see Section 3.3.4): *(i) no-CCA:* disables the carrier sense. *(ii) FEC:* applies forward error correction with fixed block of 12 Bytes of redundancy. *(iii) nC-FEC:* applies forward error correction and disables the carrier sense simultaneously. *(iv) PM:* applies RSSI-based packet merging. *(v) nC-PM:* applies RSSI-based packet merging and disables the carrier sense simultaneously. *(vi) no-action:* takes no action as the potential gain of countermeasures is not significant or interference is not harmful.

3.4.2.1 Decision Tree Classifier

Our proposed algorithm for inferring CTI countermeasures is based on a machine learning technique named decision tree learning (DT) [35, 38]. This learning technique is popular among the inductive inference algorithms and has proven successful in a broad range of tasks.

3.4.2.2 Feature Selection

In Section 3.3.3, we empirically explored the feature space of the CTI classification problem and highlighted the set of features that best describe the problem. We consider packet specific features, spectrum specific features, and communication link features.

The initial set of features consisted of 25 features derived from domain specific knowledge. Training the classifier using all avalible features is not a good practice as this can result in overfitting. To reduce the feature set, we use the following selection techniques: First, we evaluate the level of intercorrelation among features and exclude redundant features. We identify a subset of five features that are uncorrelated with each other,

Feature	Description	Purpose
(1) Occupancy level	number of busy samples / total number of samples	Temporal behaviour of interferers
(2) Duty cycle	{T,F}, patterns of x consecutive idle samples	
(3) Energy span during packet reception	range($RSSI_{normalized}$)	Capture RSSI temporal characteristics
(4) Energy level during packet reception	median(range($RSSI_{normalized}$))	
(5) RSSI regularity during packet reception	weighted-average(mode($RSSI_{normalized}$))	
(6) Packet corruption rate	ratio(P_{crc})	Detect and quantify harmful interference
(7) Packet loss rate	ratio(P_l)	
(8) Packet length	average(packet_length)	Capture error characteristics
(9) Error rate	ratio(err_typ == C1)	
(10) Error burstiness	ratio(err_typ == C2)	
(11) Energy perception per packet	range(CCA)	Detect unnecessary deferrals
(12) Energy perception level per packet	median(CCA)	
(13) Backoffs	ratio(CCA_deferral)	

Table 3.1: Features utilized by TIIM. These features are calculated over an observation window W_{active}. To remain environment-agnostic, each packets's fine-grained RSSI series are normalized. Features 1-2 can either be induced from RSSI series of multiple packets, or as stand-alone sampling of the medium for a short period of time. Features 3-5, and 12-13 are collected during packet receptions, and aggregated in a representative way. Features 6-8, and 13 are communication statistics. Features 9-10 are based on per packet RSSI computations.

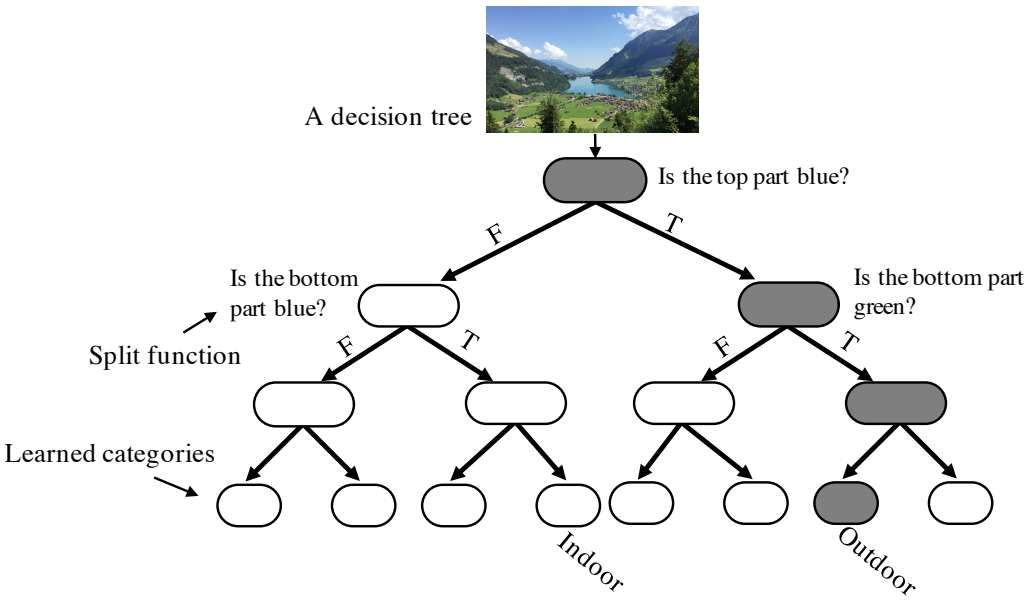

Figure 3.11: A decision tree that illustrate how to find out if a given picture is taken in indoor or outdoor. Each internal node in the tree stores a split function to be applied to incoming data. Leaf nodes store the final answers. Figure adapted from [35].

yet correlated in predicting the same class. Then, to increase the accuracy of the initial subset, we apply exhaustive search to evaluate all possible remaining subsets in combination with these five fixed selected features. We benchmark all features that contribute to decision trees with high accuracy and select those with the highest rank. Table 3.1 summarizes the final set of features we use to train our classifier and which are used by TIIM during runtime.

3.4.2.3 Data Labeling and Ground Truth

We develop an automatic annotating software for extracting the countermeasure labels for our dataset.

We divide our dataset into training and test sets. The training set is the set of data points for which the corresponding countermeasure we seek is known. We use this set to compute the tree parameters. Each sample in the dataset (x_i, y_i) represents the feature vector x_i that describes a window of communication events in an instance of interfered link, and y_i the corresponding label which indicates the best mitigation strategy for this instance. Using the training set, the supervised learning algorithm aims at constructing a good model that learns the complex pattern in the training set to predict class y' for an unseen feature vector x'.

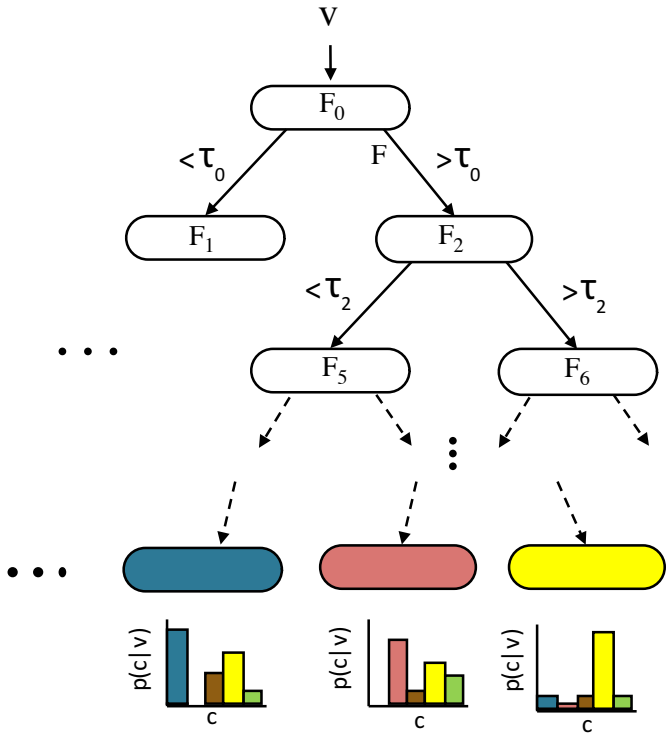

Figure 3.12: Illustration of a testing decision tree. Starting at the root, each split node, evaluates its function F_i on the input data v. Input data is forwarded to left or the right child node based on the output of F_i.

In Section 3.4.2.2, we elaborated on how to construct and calculate the feature vectors $\{x_1, \ldots, x_N\}$ (see Table 3.1 for complete list of features). Now we discuss how we automatically label the samples in our dataset. To label our dataset for each window of observations, we simulate the outcome of each of the considered mitigation strategies. This yields for each instance a corresponding gain (i.e., PRR increase), and cost (i.e., communication overhead). The labeling algorithm quantifies the benefit of each countermeasure and selects an optimal countermeasure $A \in \{no\text{-}CCA, FEC, nC\text{-}FEC, PM, nC\text{-}PM, no\text{-}action\}$ that achieves the highest gain-cost balance, as defined by application requirements.

The application expresses its requirement through a gain $g()$ maximization and cost $c()$ minimization equation: $f(A) = g(A) - c(A) \times \alpha$, where the best countermeasure is the one with the highest $f(A)$. The configuration parameter α defines the weight of the cost. Note that $g(A)$ and $c(A)$ are normalized, and both are in the range $[0, 1]$. With $\alpha = 0$ the cost is not at all considered, which results in the highest possible total gain. The higher α, the more the algorithm emphasizes on minimizing the total cost. In our experiments, we consider $\alpha = 0.5$ which is a good tradeoff between total cost and total gain.

Decision Tree (DT). DT classification model can be seen as sequential binary decisions, corresponding to traversing binary trees. Trees here consist of inner nodes (i.e., split nodes) and leaf nodes (i.e., classes). The split functions embody questions that add incrementally to the certainty of the correct class. The DT classification model estimates an unknown property of a given object by consecutively querying about its known properties (i.e., features). Each node evaluates a computationally inexpensive function on one of the input features and as a result, forwards the currently evaluated data recursively down in the tree until the data reaches one of the leaf nodes. Determining which question to ask next depends on the answer to the preceding question. The final decision corresponds to the leaf (terminal) node reached by the input object. Figure 3.11 presents an illustrative example of a decision tree algorithm that checks if a given picture represents an indoor or outdoor scene.

Tree Training. This phase is in charge of the optimization of split functions attributes. Each intermediate node in the decision tree has associated test function $F(v, \tau_i)$ with a binary output (T, F) and split attribute τ_i associated with the split node. These functions are vital to the classification performance and their attributes are learned automatically from statistics of the training data in the training phase. The tree is typically constructed from top to bottom, where the attributes that maximize the information gain about the classification are selected first. At each node in the tree, the attribute that best splits (i.e., the attribute with the highest normalized information gain) the incoming training set E_i into E_i^L, E_i^R is selected. The symbols E_i, E_i^L, E_i^R represent the training point before and after the split. Figure 3.12 illustrates a general learned decision tree for testing.

Albeit decision tree classifiers are not necessarily the best classifiers in terms of accuracy, they are relatively efficient regarding computational and memory overheads; with careful optimization, they can run on severely constrained devices. We use the C5.0 algorithm [132] to generate our decision tree. Our tree consists of 200 leaves in case trained by anechoic chamber traces and 300 in the case of office traces.

3.5 Experimental Evaluation

In this section, we present the experimental evaluation of TIIM. We begin in Section 3.5.1 by describing the experimental setup and briefly describe the trace collection methodology. Then, in Section 3.5.2, we elaborate on the accuracy of the decision tree classifier in inferring the correct countermeasures. In Section 3.5.3, we perform a trace-driven simulation to demonstrate the prospective gain of applying TIIM compared with the

gain of applying fixed mitigation strategies, followed by evaluating the system gain online.

3.5.1 Experimental Settings

Up to this point, the analysis has been carried out in an anechoic chamber, an environment that is shielded from external radio interference and diminishes multipath propagation effects. This allows us to recognize patterns and identify the impact of each of the considered interfering technologies in isolation. In the following, we address the evaluation of TIIM in a typical environment that incorporates the impact of external uncontrolled interferers, other channel impairments (e.g., multipath effect), and multiple sources of interference. We have performed all our experiments in ETH Zurich's computer science building. Figure 3.13 shows the layout of the experimental setup. There are two stationary sensor nodes located in an office room with a line-of-sight link of 4.5 m.

We consider two types of experiments: *(i)* Single active interferer: in this run, we use each of the considered interferers to generate interference individually. The interferers are located in this run first at location L1 (3.5 m), then L2 (6 m), and L3 (10 m). Both locations L2 and L3 are in non-line-of-sight to the sensor nodes, while L1 is within the line-of-sight. *(ii)* Multiple active interferers: in this run, we consider interference generated from multiple sources running simultaneously. The positions of the interferers are highlighted as circles. During the experiments, the nodes were exposed to interference from various uncontrolled sources existing in the building. To mention some, the university's WiFi network which is present on 802.11 channels 1, 6, and 11, Bluetooth mice and keyboards, and a small 802.15.4 heating control system deployed in the same floor.

Methodology. Our focus is to capture channel statistics over the 802.15.4 link between node A and node B. Node A sends short packets (20 byte) and long packets (100 byte)[1] at 100 ms intervals to node B.

We first perform experiments without controlled interference, then with a single interferer activated at a time, and finally with activating multiple interferers. For the single interferer run, we consider all of the interferers mentioned in Section 3.3.1. For the multiple interferers scenario, we consider a subset of these interferers, as highlighted in Figure 3.13. In all office environment experiments, the sender uses its maximum transmission power (0 dBm).

We instruct the sender to disable carrier sense and log the CCA

[1]Note that, the experiments carried out in the anechoic chamber considered more configurations of different packet lengths, transmission powers, and traffic patterns.

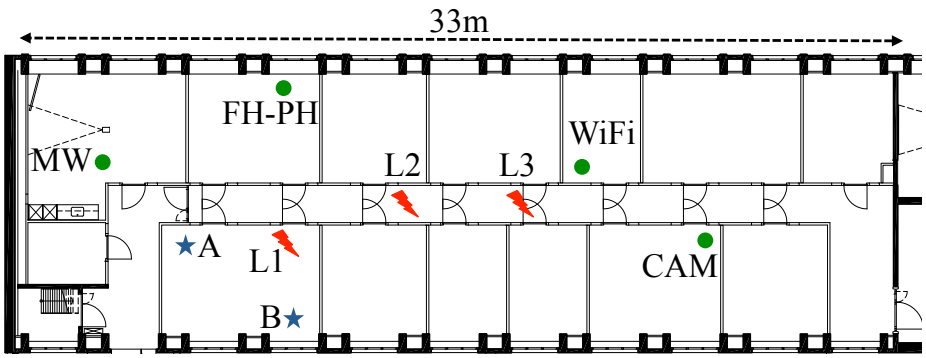

Figure 3.13: Layout of the office experiment setup. A and B are TelosB motes located in an office with a line-of-sight link of 4.5 m. A is the sender and B the receiver. The interferers are located at locations L1 (3.5 m), L2 (6 m), and L3 (10 m), where L1 is within line-of-sight to A and B and the other two locations within non-line-of-sight. Circles indicate the location of our multiple interferer scenario.

value at the time of transmission. This allows us to perform trace-driven simulation for countermeasures involving carrier sense enabled and disabled. We instruct the receiver's radio to pass packets with failed CRCs rather than discarding them to enable us processing erroneous packets. Moreover, the modified radio driver samples RSSI at a rate of 62.5 kHz during packet reception along logging other relevant PHY and link-layer metrics. Overall we capture 64 hours of extensive CTI experiments. We collect fine-grained channel and communication measurements to allow systematic evaluation and comparison of TIIM and the considered countermeasures under the exact adverse link dynamics.

Rather than detailing on recovery results per technology, we focus on discussing the adaptability of TIIM under dynamic and various interference implications with the goal of maximizing the overall performance gain and minimizing the overall overhead cost.

3.5.2 TIIM's Classifier Accuracy

We first discuss the performance of TIIM's core component, the decision tree classifier, in inferring the correct countermeasure. As mentioned earlier, we use half of the data points in the office environment for training the classifier. The data points include various types of interference instances, as described above. The other half of the data points are used to evaluate the prediction accuracy. Table 3.2 shows the confusion matrix for DT classification. The decision tree achieves a mean classification accuracy of 92.9%. The reason why TIIM's classifier achieves higher accuracy than traditional interference classification

Predicted as	(a)	(b)	(c)	(d)	(e)	(f)
(a) no-action	**95.0**	3.9	0.0	0.2	0.5	0.4
(b) no-CCA	0.0	**97.2**	0.0	1.7	0.0	1.0
(c) FEC	11.4	4.5	**68.2**	11.4	2.3	2.3
(d) nC-FEC	0.1	16.8	0.2	**80.8**	0.0	2.2
(e) PM	31.1	12.2	1.1	0.0	**41.1**	14.4
(f) nC-PM	0.5	36.6	0.2	5.2	0.1	**57.5**

Table 3.2: Confusion matrix of the decision tree on the traces collected in office environment.

approaches [68], is that our classifier does not need to differentiate between radios causing similar channel signatures.

While TIIM achieves high accuracy in inferring the correct countermeasure for most of the classes, it performs poorly for *Packet Merging* (PM). This is mainly due to the low occurrence of incidents that were labeled as PM in our dataset, e.g., only 0.18% are labeled as PM in our office environment traces. Since PM is under-represented in our dataset, the DT could not learn well its characteristics. The low number of PM instances in our traces is mainly due to some hardware-based inaccuracies affecting the localizing of corrupted symbols within a packet. Consequently, we were limited from achieving the full potential of PM.

In general, achieving high accuracy in localizing the positions of errors by solely relying on off-the-shelf radios is hard. For instance, we encountered in some cases of RSSI sampling non-consistent delays that are reflected in a slight drift between the RSSI surge position and the actual location of the error burst which affected the accuracy of our scheme of error localization. We believe that designing radios that allow better interfacing between PHY and upper communication layers can yield better performance for RSSI-based packet recovery schemes.

The consequence of prediction inaccuracies in our system is not necessarily high. TIIM has the chance to re-adjust its suggested countermeasure after the time window of observing the channel. Hence, the worst case scenario is that for the next window time, TIIM introduces a cost that does not yield any gain. To verify whether the accuracy achieved by the classifier is not tied to the training environment, we train the classifier on our dataset from the anechoic chamber and evaluate it on the dataset from the office environment. The mean classification accuracy of 94.1% is even slightly higher than the accuracy achieved when the DT was trained and evaluated on data points from the same environment.

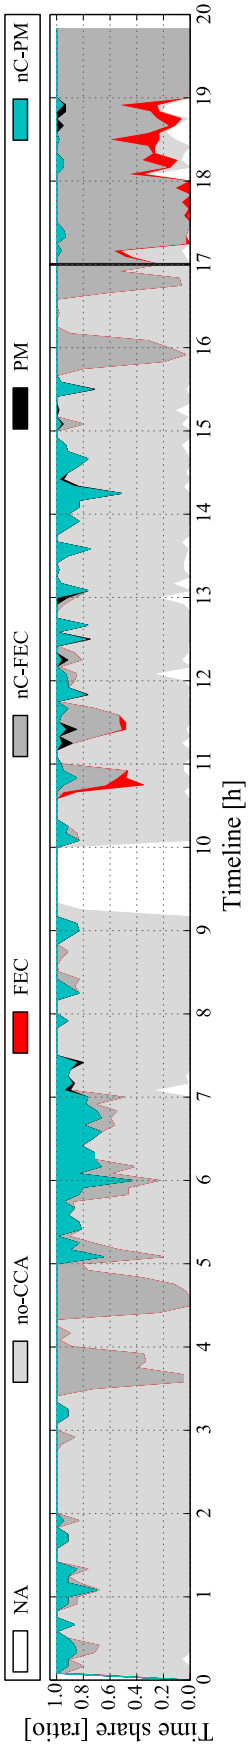

Figure 3.14: Offline performance of TIIM (resolution 5 min).

3.5.3 Results

We discuss the evaluation results of TIIM from two aspects. First, we discuss the potential benefits of leveraging the detection capability of TIIM, as opposed to employing a fixed countermeasure through trace-driven evaluation over 64 hours of CTI extensive runs. Second, we present the overall system performance gain achieved while running TIIM.

Evaluation Metrics: We employ the following metrics to evaluate the performance of TIIM.

• *Packet Reception Ratio (PRR):* This is the ratio of successfully received packets over the total number of transmitted packets during a specific time period.

• *Gain:* compares the achieved PRR to the baseline PRR (default 802.15.4 PRR under interference).

• *Cost:* the ratio of transmission overhead introduced by the countermeasure to the base transmission. For instance, for FEC a fixed transmission overhead of 12 Byte per packet is considered. For no-CCA, the transmission of positive deferrals is considered as the cost. Positive deferrals are those transmissions that were lost or corrupted, but would have been deferred with CSMA. PM requires retransmission of the localized corrupted symbols. Hence, we consider the overhead cost of the new frame.

• *Adaptability:* to detail the ability of TIIM to adapt quickly to unanticipated changes in the interfered channel, we illustrate the degree of the system adaptability by showing its dynamic behavior in a timeline plot.

3.5.3.1 Adaptive Interference Mitigation

We now evaluate the overall prospective performance gain for an adaptive interference mitigation system as compared first to the performance of standard 802.15.4 and then to the gain of applying a fixed interference countermeasure. We perform a trace-based simulation using the 64 hours of CTI traces. We run each of the countermeasures and calculate its corresponding cost and gain. As depicted in Table 3.3, the PRR of our traces under interference lies around 56%. Applying a static countermeasure could potentially be the best solution in case the channel conditions remain static and best for that countermeasure. However, due to changes of interference patterns in our traces, the static countermeasures can cause high-cost overheads. For instance, FEC in combination with noCCA causes about 30% additional transmission. TIIM achieves almost the highest PRR gain with a cost of 5.6%.

The adaptability of TIIM enables it to perform best in a dynamic

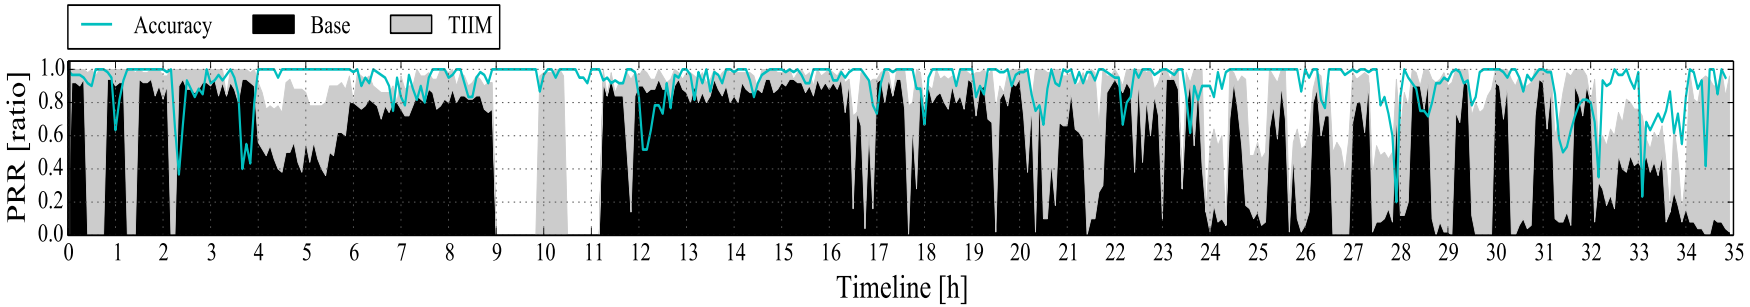

(a) Performance of TIIM while applying recommended countermeasures.

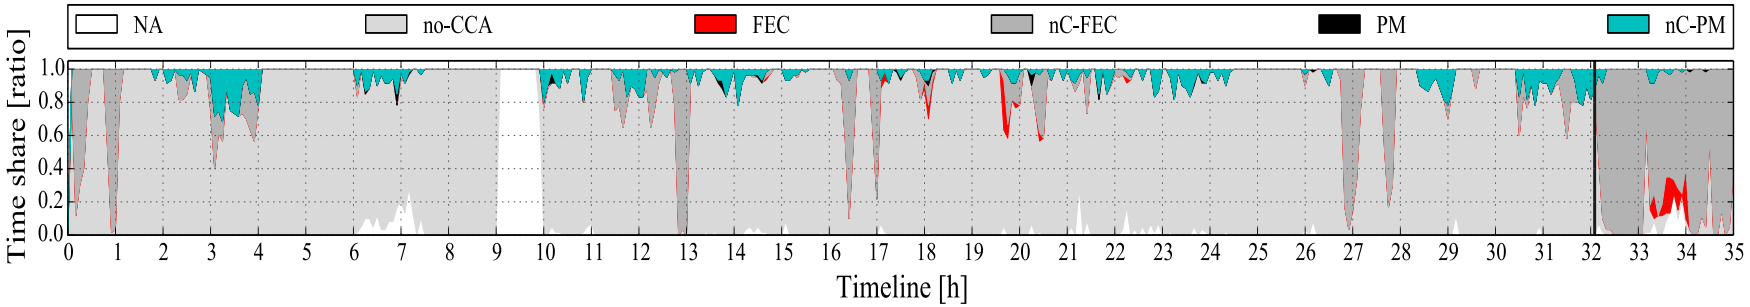

(b) Time share of each countermeasure during the online phase.

Figure 3.15: Online evaluation of TIIM. Inquiring TIIM for countermeasures while providing the channel conditions as input (resolution 5 min).

Countermeasure	Gain (PRR)	Cost
Base	0.566	0.000
no-CCA	0.863	0.106
FEC	0.572	0.198
nC-FEC	0.882	0.291
PM	0.567	0.005
nC-PM	0.723	0.037
TIIM	**0.873**	**0.056**

Table 3.3: The performance of TIIM as compared to static mitigation assignment. The dynamic countermeasure selection of TIIM allows it to reach an interesting trade-off between PRR and cost.

channel. Figure 3.14 shows the timeline of selected 20 hours of our traces and how TIIM adapts the selection of countermeasures according to its assessment of the channel.

In the following, we spot some interesting observations. Starting from hour 0 to hour 3, we notice that no-CCA is dominated with transitions to nC-PM and nC-FEC. From hour 3.5 to hour 5, nC-FEC becomes the dominating countermeasure. We notice the first use of PM at hour 7. Particularly, the no action between hour 9 and 10 is interesting, it happens while the high-power analog phone is active at distance 3 m, causing severe interference. In this period, communication is not viable over the interfered channel, and thus TIIM recommends channel switching. Identifying such instance is essential for saving energy. We encounter the first uses of FEC from hour 10.5 to hour 11.5. During the same period, PM has become more active. It is worth noting that the microwave oven is the dominant interferer during this period. Notice that starting from hour 17 (indicated by the vertical black line), we expose TIIM to condensed corrupted traces. Consequently, TIIM's countermeasures are dominated by nC-FEC, FEC, and partly PM.

3.5.3.2 Online Performance

Now, we evaluate the online performance of TIIM. To this end, we run TIIM on extracted features computed over the time window of 5 s. Whenever TIIM recommends a countermeasure, our system applies it.

Figure 3.15(a) illustrates the base PRR of the traces in black and the additional gain achieved by TIIM in gray. The accuracy of our system in selecting the right countermeasure, at any given time, is plotted as a line. To visualize TIIM's dynamic behavior, we show the time share of each countermeasure in Figure 3.15(b), time-synchronized with the PRR in Figure 3.15(a). Within the first hour, TIIM achieves a gain of 100%,

several times enabling PRR to go from 0% to almost 100%. At hour 2, we observe a short drop in the accuracy, but the resulting gain remains high. There are further sharp drops in the accuracy of the system that do not lead to a PRR decrease. These are the cases where the recommended countermeasure still could yield an acceptable gain, but possibly at a higher cost. However, not all inaccuracies remain unpunished by the PRR. For instance, around hours 7, 17, 28, 33, and 24 the PRR drops as a consequence of temporarily decreased accuracy.

TIIM successfully detects the first occurrence of persistent harmful interference (analog phone at location L1) between hours 9 to 10. It detects persistent non-harmful interference (analog phone at location L2) and consequently recommends no-CCA, PM, and nC-PM which yields 100% of gain. However, it fails from hour 10.5 to hour 11 to detect another occurrence of persistent harmful interference, where the system should have recommended channel switching. Starting from hour 19, TIIM shows its potential in recovering severe performance degradation. TIIM is stressed with a heavy load of concentrated corrupted packets after hour 32. It applies dominantly nC-FEC, FEC, and partly nC-PM to improve a close to 0% PRR to almost 95%.

Reaction Time. Traditional classification approaches need a few seconds to detect the type of interference source, e.g., 18.14 s for WiFi [68]. This is relatively a long time for highly dynamic channels. One advantage of TIIM is that it reacts to interference shortly after detecting the degradation caused by interference. Currently, TIIM provides a recommendation 5 s after detecting interference. This allows TIIM to react timely to time-variant interference patterns.

Coexistence with other Radios. We verified empirically that 802.15.4 does not cause harmful interference to high-power wireless devices such as the wireless cameras. Even when disabling CCA, we did not observe any effect on wireless cameras' operation. On the other hand, 802.15.4 can cause harmful interference with coexisting low-power radios, resulting in 802.11 deferrals and packets losses for Bluetooth. TIIM can potentially be trained to make a tradeoff between its performance and the harmful interference it may cause to the low-power networks or be trained to apply no-CCA only for high-power interferers.

3.6 Discussion

This chapter provides a proof of concept on the potentials of a CTI-aware and smart adaptive link-layer solutions. However, more research and experimentation are needed to generalize and realize the full potential

of TIIM. Here, we address some practical challenges and research points that assist in evolving TIIM further.

(a) Porting TIIM to other Radios. Although this work focuses on low-power networks, most of the observations can be projected to analogous RF technologies, such as 802.11 radios. The core concept of integrating reasoning to combat interference diversity in the unlicensed bands has not been explored before. We believe that TIIM's core concept can be beneficial to analogous wireless technologies, and leave further investigations for future work.

(b) Extending TIIM with New Countermeasures. In this work, we explored the feasibility of addressing the CTI heterogeneity problem by focusing on aspects that are relevant to the set of few mitigation approaches considered in this prototype of TIIM. To extend the system with new mitigation approaches, relevant features need to be redefined and the classifier needs to be retrained. Possible examples of interesting countermeasures that can benefit from a learning module are: Detecting systematically the duration and interspace of interference pulses could be a useful metadata for an adaptive FEC, to select the right level of redundancy required by FEC, or capturing tendencies in duty cycles can be exploited for a better MAC scheduling.

(c) TIIM's Limitations. TIIM has a narrow view of the RF spectrum that is limited to the 802.15.4 channel width. It focuses on increasing the spectral efficiency over interfered channels, with lack of cognition about the state of the rest of the spectrum. Thus, it lacks a comprehensive view of the RF spectrum to decide whether communication over an interfered channel is preferred over channel switching.

(d) Interference and PHY Layer Information. Over the last few years, researchers advocated a design of wireless systems that allows a better interfacing of physical layer information for higher layers, particularly to cope with interference. We developed this work with legacy systems in mind, thus we were limited to the PHY space provided by these systems. One PHY aspect that can be integrated into TIIM to overcome its limitation of narrow spectrum perception, is the use of cyclostationary analysis for bandwidth estimation of interfering signals as suggested by DOF [78]. DOF estimates the bandwidth of interfering signals solely based on the PHY information retrieved from the channel frequency in use. This allows a better reaction in severe channel conditions, where a less affected channel outside the interferer's bandwidth could be selected. In Chapter 4, we continue in this direction; we harness physical layer information to develop solutions that enhance wireless systems coexistence.

(e) Machine Learning in Wireless Communication. The concept of bringing intelligence to radios is not new. Communication concepts such as cognitive radio, promise integration of intelligence into radios, such that they can sense, learn from, and adapt to their environment. To date, most of the cognitive radio research focuses on licensed bands and has been restricted to policy-based radios that are hard-coded with rules on how to react in certain scenarios. Devising radios that utilize machine learning techniques, i.e., a learning-based cognitive radio, is a relatively unchartered research area. Ideally, the learning algorithm uses past observations to form a hypothesis about the nature of the channel, by exploring the relation between attributes and inferring patterns in channel statistics. Then use this hypothesis for predicting measures of interest in the future. In these settings, the learning engine is responsible for augmenting the list of actions available to the radio that allows it to adapt to a changing environment. Hence, radios can remember lessons learned in the past and act upon in the future. This chapter shed the light on the potential of using machine learning in one wireless communication application, namely automating countermeasure selection in the presence of interference. Interference in the unlicensed band is primarily from communication systems that follow systematic protocols which can be learned and exploited for better coexistence. This chapter is devoted to exploring this direction. The scope of machine learning applications in wireless communication is wide and needs further research exploration. The need of learning components in radios is more evident now, as the rise of active wireless devices implies that more RF optimizations and tuning is needed.

3.7 Related Work

We distinguish three major directions adopted to combat interference in the unlicensed bands and related to the work presented in this chapter. The first direction aims at **detecting and avoiding interfered frequencies** by employing spectrum sensing to identify interference-free channels [126, 162]. Musaloiu et al. [107] propose a distributed algorithm for channel selection and interference estimation using RSSI sampling for 802.15.4 networks. The lack of interference-free channels, and the fast and unpredictable changes in the occupancy state of frequency bands make the sampling overhead of these approaches high, particularly for resource-constrained devices.

The second direction aims at **increasing resilience against interference**, by bracing PHY and MAC layers with auxiliary mechanisms.

For instance, Liang et al. [96] studied the interplay between 802.11 and 802.15.4 and applied a resilience forward error coding scheme to interference [96]. Analogously, some solutions focused on exploiting the temporal effects of interference induced on the PHY hints, such as variations in soft errors (softPHY) [85] or RSSI variations [71, 66] to recover interfered packets. Others focused on increasing the robustness of existing MAC protocols against interference [20, 21], or considered utilizing multiple radio channels for communication [9, 60] to exploit frequency diversity. Moreover, further PHY solutions have been considered, such as utilizing advancements in MIMO for interference cancellation [54, 166].

The third direction aims at **identifying the type of interference** by employing signal classification techniques [34, 93, 78] or featuring distinct interferer's patterns on corrupted packets [68]. It is, however, not yet clear how the interference classifiers can be utilized in an automated way to mitigate interference, given the diversity of interference technologies. Our work aims at bridging the second and third directions, by featuring classification to recognize interference patterns which can provide useful meta-information about the applicability of a certain mitigation strategy.

3.8 Summary

Wireless interference has been a long sought but still a crucial problem in wireless communication, notably for systems operating in the unlicensed bands. While most existing solutions focus on the careful tuning of signals to realize frequency isolation, less work has focused on thoroughly utilizing the interfered links under heterogeneous interference patterns.

In this work we investigate how to enable wireless nodes to make optimal decisions in situations involving high uncertainty, consequently, dynamically adapt based on their surrounding wireless environment. We have been inspired by the broad vision of cognitive communication to design wireless systems that are computationally intelligent about radio resources and the surrounding wireless environment. It is essential to realize this comprehensive, intelligent system for managing wireless resources before the wireless medium becomes so unreliable as to be unusable.

More concretely in this chapter, we introduced TIIM, an interference mitigation system that proposes countermeasures that work best under the current interference patterns, independent of the particular technology causing it. We leverage previously unconsidered channel attributes and employ a lightweight machine learning classifier to *(i)* decide whether the communication is viable over the interfered link,

(ii) learn contending signal patterns, and *(iii)* find the best underlying link-layer coexistence scheme. Doing so, TIIM realizes the full potential of interfered wireless links and consequently enhances spectrum efficiency. Our evaluation shows that TIIM improves the packet reception ratio under interference by about 30% with only 5.6% additional transmission overhead.

In the following chapter, we further pursue research in adaptive recovery but focus on cross-layer solutions. We move way from bringing solutions for off-the-shelf radios and instead utilize software radios to explore the space of richer physical layer information that can be used to enhance wireless coexistence.

4

Cross-Layer Optimization for Wireless Coexistence

Wireless systems use a variety of physical layer techniques to combat channel impairments. For instance, 802.15.4 employs spread-spectrum modulation and error control coding. However, these techniques alone fall short in mitigating the effects of Cross-Technology Interference (CTI). CTI severely reduces the *Signal-to-Interference-plus-Noise Ratio* (SINR) of the intended transmission, which results in high bit-error rates and limits the effectiveness of these techniques. Energy and complexity constraints in low-power networks prohibit the use of complex interference suppression and cancellation techniques that are finding their ways into unconstrained wireless systems [79, 54]. In this chapter, we shift our focus to the physical layer. Particularly, we investigate how to exploit physical layer properties of low-power signals to address CTI better, while still maintaining a low complexity. Consequently, we explore the space of physical layer information in the 802.15.4 radios that can be used to enhance low-power wireless systems coexistence.

Contributions and Roadmap. To tackle the issues above, we present CrossZig, a cross-layer solution that enables low-power wireless nodes to make informed decisions on their coexistence strategies based on richer physical layer information, thus adapt autonomously to the current interference patterns in the channel. CrossZig achieves this by leveraging novel building blocks and system designs.

This chapter makes the following contributions:

- We introduce a *Physical layer Hints Interface.* This interface allows higher layers to access richer physical layer information and consequently use it in a variety of algorithms to boost performance under CTI (Section 4.1).

- We develop a novel lightweight technique that allows low-power nodes to recognize the type of interference in interfered packets (Section 4.3). Our design achieves high accuracy in detecting CTI at all SINR ranges where the target signal cannot be decoded correctly and it does not require any prior synchronization between nodes (unlike [135, 56]).

- We present CrossZig, an adaptive recovery mechanism that exploits both block-based error correction and packet merging through diversity combining at the signal level (Section 4.4). We show that both of these schemes come at low complexity costs while - if carefully performed - they can effectively alleviate the damage due to *Cross-Technology Interference*.

In light of our contributions, Section 4.5 and Section 4.6, cover the implementation and evaluation of CrossZig. We implement and evaluate a prototype of our system in SDR using GNURadio [53] with USRP-N210 [44]. Section 4.7 surveys related work, and Section 4.8 provides brief concluding remarks. This chapter is based on the contributions made in [74, 73].

4.1 Physical Layer Hints Interface

In this section, we present the Physical layer Hints Interface (PHY-hints). This interface is an extension of the standard physical layer. It allows higher layers to access a richer physical layer information. When physical layer signals are received, besides standard processing, such as demodulation, chip-to-symbol mapping, and delivering decoded symbols to the data-link layer, the physical layer also accommodates further hints that can be exploited to boost the performance of wireless systems [159, 85, 153, 115]. In this work, we exploit such hints to detect *Cross-Technology Interference* and to estimate the confidence of received symbols in interfered packets. We begin this section by giving an overview of the PHY-hints interface design. Then we discuss the specific details of the PHY-hints interface for the IEEE 802.15.4 physical layer.

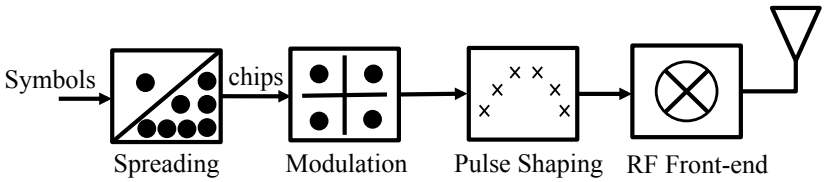

Figure 4.1: Simplified block diagram of the 802.15.4 transmitter.

4.1.1 Wireless Cross-Layer Design

Many communication systems still strictly confine with the conventional network layering principle. In such systems the physical layer at most output frame units consisting of streams of bits after demodulation and decoding to upper layers. Layers on top compute the checksums for the received frames and discard erroneous frames regardless of the level of damage in received frames. There has been a compelling evident that in wireless systems the conventional layering abstraction is limiting and has led to suboptimal wireless operations [65, 139, 142, 145, 91]. This is because such a design paradigm obscures upper layers from reasoning about channel dynamics. Alternatively, a cross-layer design seeks to enhance system performance as it allows upper layers to adapt better to channel conditions by exposing more information at the physical layer. This has led to a move towards designing wireless systems with a higher level of physical layer flexibility. The resulting flexibility helps to enhance the network performance but can increase the system complexity. To bound this complexity, cross-layer designs should maintain the classical digital abstraction between the physical layer and the upper layers. This implies that the interface between the physical layer and the upper layers should evolve to allow access to physical layer information. This would facilitate interaction between layers without eliminating the layering principle.

PHY-hints Interface. We design the PHY-hints Interface as an extension to the physical layer. This extension will allow upper layers to access further physical layer information without modifying the standard physical layer interface, with this we maintain the digital abstraction. In our design, the physical layer, similar to the conventional systems, assembles received bits into frames and passes them up to the link layer. For each group of bits b the physical layer additionally can pass a block of meta information. Higher layers can retrieve and process the side information passed for each block of bits b without further involvement of the physical layer. In our design the flow of information can be triggered on demand and upper layers can specify the type of information to be delivered as side information. The granularity of the data delivered through the PHY-hints

interface can be optimized to lower the complexity and overhead of the interaction between layers.

A key challenge in such a cross-layer design resides in identifying the essential information to be exchanged between layers. Communication systems often add redundancy to cope with the noise in the channel. This can be achieved by using a modulation with a large separation between constellation points or by applying channel coding. This redundancy can be utilized as a source of physical layer hints. We will highlight how we can leverage this inherent redundancy in communication systems to extract information that can be used by upper layers to boost performance. Jamieson [91] has extensively researched the topic cross-layer abstractions in his dissertation. Our PHY-hints interface is influenced by the findings in this dissertation. However, our design is evolved to cater hints relevant to the CTI problem and deliver an interface customized for the 802.15.4 technology.

4.1.1.1 Physical Layer Background

We briefly recall some background material on physical layer encoding/decoding and modulation schemes. Note that we omit details unrelated to the context of our discussion to simplify this communication primer. Along this presentation, we discuss how signal distortions due to channel impairments are reflected on decoders.

Symbol Constellations. In basic wireless communication systems, the physical layer encodes bits in one of few M symbols $s_i(t), ..., s_m(t)$ that are sent at singling intervals. The PHY symbols are typically represented as complex valued signals in a 2D complex plane referred to as the constellation diagram. Figure 4.2(a) illustrates the ideal constellation diagram for an example modulation scheme, namely, the 4-ary quadrature amplitude modulation (4-QAM). The receiver recovers the *In-Phase* (I) and *Quadrature* (Q) values of received symbols and places them on the constellation diagram. As the signal get distorts by channel impairments, the received symbol positions get dispersed from the ideal position. This is illustrated in Figure 4.2(b), where α as physical layer hint reflects on the distance in signal space between received constellation point (received signal) and the decoded constellation points (expected symbol). Communication systems typically use modulation schemes with the separation between constellation points relative to the noise in the channel to accommodate for symbol dispersions.

Coded Communication. Channel coding is typically utilized in communication systems to control errors in noisy communication channels. This is done by introducing redundancy in the channel encoder

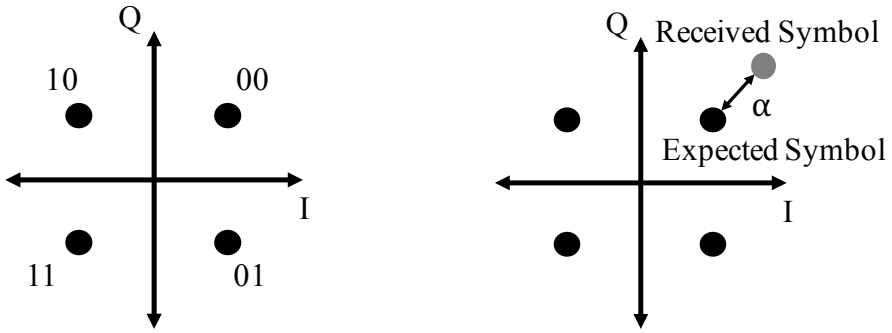

(a) Ideal 4-QAM. (b) Distorted Signal.

Figure 4.2: Representation of 4-QAM Constellation Diagram.

so as to enable the decoder to reconstruct the original source sequence as accurately as possible. We confine this discussion to *block coding*, a popular coding scheme used in communication systems. The core idea of block coding is to use finite-lengths vectors (referred to as code-words) to transmit blocks (of length k) of segmented incoming binary data. For block code with a codeword of length n (where $n > k$), there are 2^n possible codewords. A subset of these codewords is used to make the codebook $C_1, ..., C_M$, which the transmitter is restricted to send over the channel. The encoding process maps the incoming sequence of source data blocks to one of code blocks specified in the codebook. The resulting code is referred to as an (n, k) binary block code of rate $R = k/n$. The transmitter then groups the coded data into channel symbols, modulate it and send it over the channel. At the receiver, the decoder computes the Hamming distance between the received symbol y_s and each codeword in the codebook. Then decides on the codeword r with minimum Hamming distance d_H.

$$r = \arg\min_{r} d_H(y_s, C_r) \tag{4.1}$$

The Hamming distance $d_H(y_s, C_r)$ between the decoded C_r and received symbol y_s, reflects on the decoding confidence.

4.1.2 IEEE 802.15.4 PHY-hints Interface

The discussion carried before provided a general overview. Now we dive into the specifics of our cross-layer interface. The scope of physical layer hints is typically tied to the details of the physical layer, namely the modulation and coding schemes employed by the target wireless technology. Therefore, we briefly review relevant aspects of the IEEE 802.15.4 physical layer and then elaborate on the subset of physical layer hints we consider in our design.

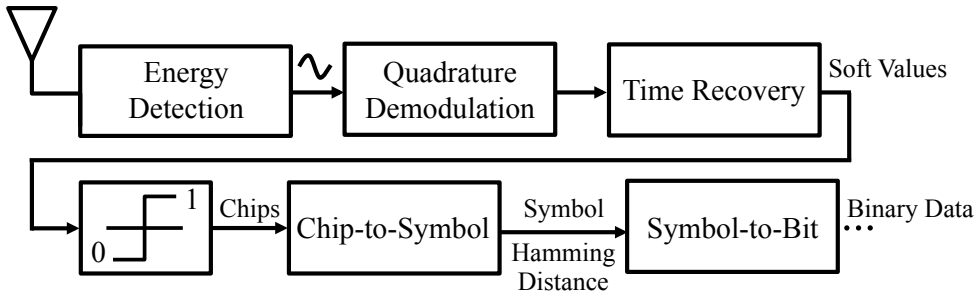

Figure 4.3: Block diagram of the receiver and corresponding PHY hints.

IEEE 802.15.4 Physical Layer. The IEEE 802.15.4 standard [81] for wireless devices operating in the 2.4 GHz employs Offset Quadrature Phase-Shift Keying (O-QPSK) for modulation and Direct-Sequence Spread Spectrum (DSSS) for spreading. Figure 4.1 shows a block diagram of the physical layer components of an 802.15.4 transmitter. At the physical layer, data is first grouped into 4-bit symbols and then spread to a specified 32-bit long *Pseudo-random Noise* (PN) sequence ($b_0b_1b_2b_3 \rightarrow c_0c_1c_2 \ldots c_{31}$). Each bit ($c_i$) in a PN sequence is then modulated using Offset Quadrature Phase-Shift Keying (O-QPSK). As shown in Figure 4.4, the even chips $c_0c_2c_4 \ldots$ are modulated as *In-phase* (I) component of the carrier and the odd indexed chips $c_1c_3c_5 \ldots$ are modulated as *Quadrature* (Q) component of the carrier. The time duration of each chip is 1 μs and there exists a half chip time (T_s = 0.5 μs) offset between the Q-phase chips and I-phase chips, which results in a continuous phase change and constant envelope.

For demodulation, the receiver's radio converts each half-sine pulse signal into a chip. Then these chips are grouped to provide PN sequences. The de-spreading is performed by mapping the PN sequence to the symbol with the highest correlation. Unlike modulation schemes such as QAM or ASK, which operate by varying the amplitude of the carrier wave, 802.15.4 adopts O-QPSK modulation. Hence, the carrier wave amplitude of all chips within one packet is constant and depends on the selected transmission power (i.e., constant envelope). The 802.15.4 chips are shaped by half-sine pulses at the transmitter. While the signal's shape will be distorted by noise in the wireless channel, its basic shape is maintained. The demodulator's output provides an indicator of how close the received signal shape is to the expected shape. We leverage these two features (i.e., constant amplitude and signal shape) of the 802.15.4 PHY in the design of CrossZig. We now elaborate on the physical layer hints we harness in CrossZig.

Signal Power. When two signals interfere, their energies add up (i.e., signals can add up constructively or destructively based on their

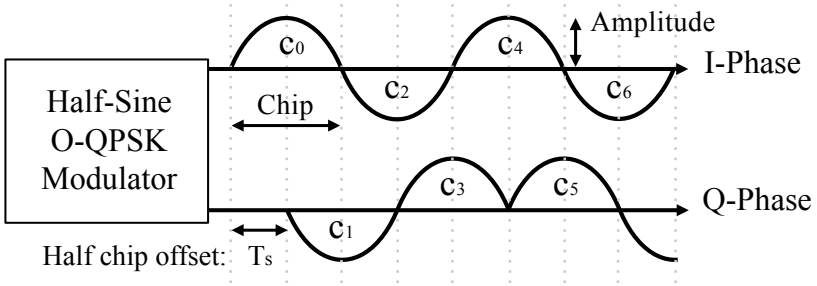

Figure 4.4: IEEE 802.15.4 modulation (O-QPSK).

phase alignment). Interfered segments of the received signal generally experience larger power than the rest of the signal, assuming the signals add up constructively. The interfered segment of the signal exhibits lower SINR, thus experiences a higher error rate. This insight on the additive energy of interfering signals highlights the ambient information the signal carries along and can assist in detecting interference and localizing interfered symbols within interfered packets. Figure 4.5(a) plots the signal power of a partially interfered packet. Once exposed to interference, the signal experiences a sudden sharp increase in the signal power.

Hamming Distance. In the 802.15.4 PHY, symbols are spread to a 32-chip codeword before transmission (one of 16 PN codewords). The de-spreading is performed by mapping the received codeword to the symbol with the highest correlation. For an erroneous mapping of a received codeword, many chips have to be flipped. The distance between the input and output codewords of the chip-to-symbol mapper can serve as an indicator for the confidence of symbol decoding. Figure 4.5(b) plots the Hamming distance within an interfered packet. Large and low Hamming distance values provide a good indicator of corrupted or correct symbols, respectively.

Demodulation Soft Values. *Soft Values* (SV) of demodulated bits are real numbers output by the demodulator. These values are approximations of the transmitted symbols. The receiver's demodulator maps the SV to the closest ideal symbol. For instance, in the case of Binary Phase-Shift Keying, the binary demodulated bits are retrieved after passing the soft values through a binary slicer. The bit is set to 1, if the SV is a positive number, otherwise it is set to 0. However, besides the bit value, SV also reflects on the demodulation confidence [159].

The confidence information of the SV can be interpreted based on the type of the demodulator. In case the receiver adopts a matched filter-based coherent demodulator, the soft demodulated values indicate the similarity between the received signal and ideal signal shape. Thus

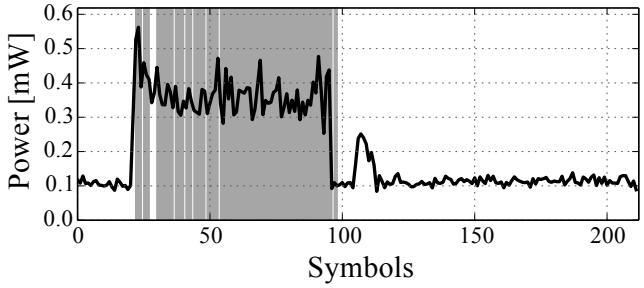

(a) Power.

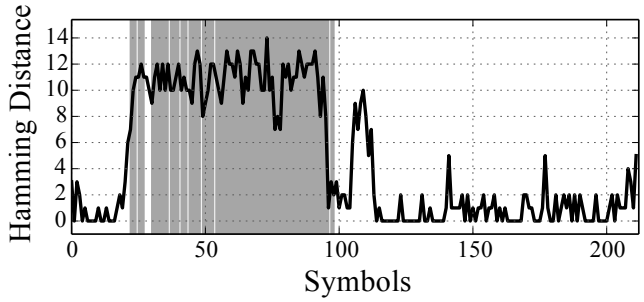

(b) Hamming distance.

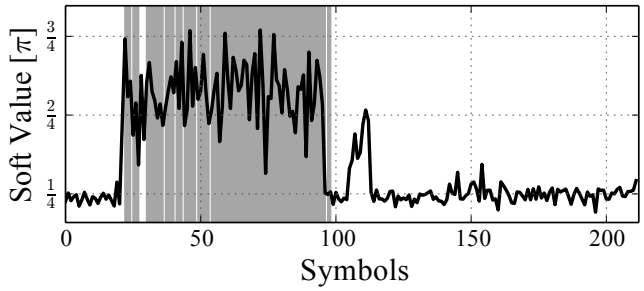

(c) Absolute Soft values.

Figure 4.5: Physical layer hints of a corrupted packet by Interference. The gray area indicates erroneous symbols.

the larger SVs, the higher the confidence for the corresponding bit to be correctly demodulated. However, if a non-coherent demodulator is used, the SV carries different information. For example, in our case the receiver uses a quadrature demodulator, which outputs the phase differences between two successive signal samples as SVs and can be computed as:

$$SV(i) = \measuredangle(s(i) \times s^*(i-1)) = \pm\frac{\pi}{4} + \delta, \tag{4.2}$$

where $\pm\frac{\pi}{4}$ is the ideal value of SV and δ is the error caused by interference and noise. Each chip is modulated by a single half-sine pulse in the transmitted signal and is represented by a sequence of four complex samples at the receiver. This implies a total phase change of π for one chip. Hence, the expected phase change between two signal samples

is $\pm\frac{\pi}{4}$. The demodulation confidence does not depend on the absolute value of SV, but the difference between $|SV|$ and $\frac{\pi}{4}$. Chips (i.e., bits of a codeword) with $|SV|$ closer to $\frac{\pi}{4}$ have a higher probability to be correctly demodulated (see Figure 4.5(c)).

4.2 Symbol Error Localization (in interfered packets)

As our recovery mechanisms involve processing of incomplete packets (partially interfered), it requires the receiver to be able to discern with a high accuracy and without additional feedback from the sender which symbols in a packet are correct and which are not. The physical layer hints described above expose statistical differences between interfered and non-interfered symbols, which render them suitable candidates to detect erroneous symbols. However, designing practical error detection algorithms based on these PHY hints with acceptable false positive and false negative rates is challenging. As we are interested in per symbol error estimation, we do not utilize SVs in error estimation (which would introduce more overhead with one SV per chip, i.e., 32 per symbol).

A direct method to estimate the symbol error is setting a threshold on the number of unmatched bits (reflected in Hamming distance) of the decoding results. This indicates the disparity between the chip sequence derived from the received signal and ideal symbol sequence. However, finding a good threshold is not trivial, as discussed before and illustrated in Figure 4.5(b). We propose an error estimation algorithm that jointly uses the number of unmatched bits of decoding results and the received signal power. Figure 4.6 shows the power mean, power variance, and Hamming distance for correct and corrupted symbols for one of our traces (we detail in Section 4.6 our experiment setup). This plot captures the intuition behind our algorithm; for low and high Hamming distances, we can classify a symbol with high confidence as successfully decoded or corrupted, respectively. For intermediate values, the Hamming distance alone is not enough (as dark and light gray dots indicating correct and corrupted symbols, respectively, overlap in this Hamming distance range in Figure 4.6). The input power though can assist to detect corruption for these cases.

Our symbol error detection algorithm works as follow: A symbol is classified as correct if its Hamming distance is lower than τ_l, whereas those with Hamming distance $\geq \tau_h$ are classified as erroneous symbols. For symbols with a Hamming distance between the decision boundaries τ_l and τ_h, we check the channel SINR. In case SINR is lower than τ_s, we

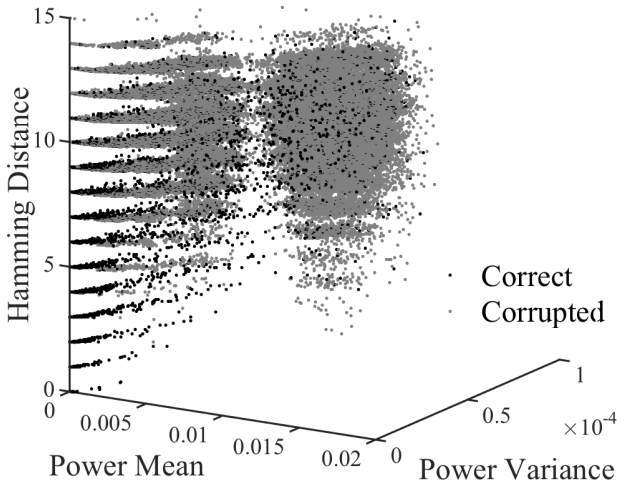

Figure 4.6: Physical layer hints for correct and corrupted symbols under wireless camera interference. Dark grey dots indicate correct symbols and light grey dots indicate corrupted symbols. The majority of received symbols with Hamming distance under 4 were received correctly and above 10 were corrupted. Correct and corrupted symbols overlap in the Hamming distance range in between. Therefore, the Hamming distance alone is not sufficient for the determination of the symbol fate. Hence, the power information can help for these symbols.

mark the symbol as erroneous. The SINR measures the channel noise and interference. It, therefore, reflects to what extent the channel preserves the correlation between transmitted and received symbols. We find this joint estimation method to be slightly better and much stabler than just setting a threshold on the number of unmatched bits. The threshold values (τ_l, τ_h, τ_s) are configurable system parameters which we derive empirically.

4.3 Cross-Technology Interference Detection

Performance degradation in wireless systems can be due to *Intra-Technology Interference*, *Cross-Technology Interference*, or *insufficient signal strength*. Determining the cause of performance degradation is essential for the coexistence problem as this defines the corresponding mitigation action to be considered. Current receivers deliver only a binary feedback on the reception status of received packets (i.e., packet passed or failed the checksum), consequently leaving receivers with suboptimal information to perform adaptations.

In Chapter 3, we addressed the interference detection problem. However, there we focused on differentiating loss causes between weak signal and interference without any consideration of the interference type. Recent works [129, 56] tackled interference detection problem as well in

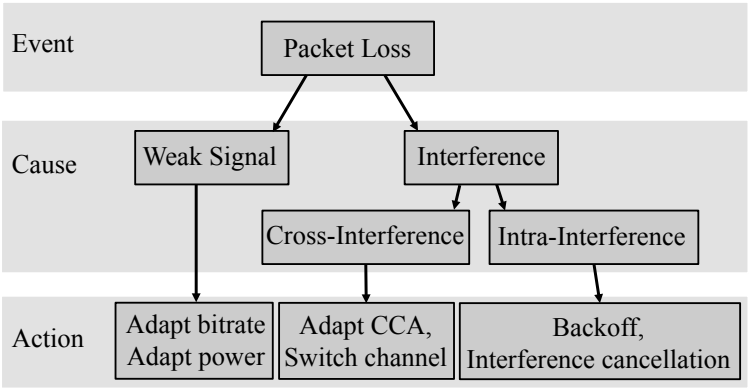

Figure 4.7: Determining the cause of packet loss can help guide better link parameter adaptations.

such binary settings.

Differentiating *Intra-* and *Cross-Technology Interference* has not yet been addressed with practical mechanisms. The inability to distinguish the type of interference leads to rather conservative approaches that blindly treat packet losses as collisions (i.e., overlapping transmissions of the same technology). Thereby exponential backoffs are invoked which can lead to starvation of low-power radios competing with high-power interferers. Moreover, interference cancellation based solutions (e.g., SIC [62]) would impose undesired overhead and worsen the chances of decodability for the target signal, if applied in the presence of CTI. Generally speaking, bracing wireless nodes with mechanisms that increase their ability to reason about the channel state will allow better adaptation and recovery (Figure 4.7). Hence, beyond detecting the presence of interference, we are interested in detecting the presence of CTI. We introduce two complementary CTI detection mechanisms that are utilized in our system: SV-based and correlation-based detections.

CTI Detection is fundamentally related to the wireless signal detection and cognitive radios spectrum sensing. Signal detection focuses on discerning signals that carry information (i.e., target signal) and random signals (i.e., noise). In CTI detection, and unlike signal detection, we are not interested in restoring the information carried by signals, but rather interested in binary feedback on the type of signal buried in a specific segment of the packet. This simplifies the detection process as we do not need to recover the signal of interest but identify its type. On the other hand, spectrum sensing in cognitive radios is focused on detecting a primary transmitter (i.e., a known signal) that is locally present in a certain spectrum. CTI detection is similar but focused on the detection of the presence of a known signal in the interfered segment of the packet, hence,

higher uncertainty in SINR. In the context of this dissertation, we present two CTI detection mechanisms that we later utilize in CrossZig. With this, we can have robust mechanisms in place that allow us to determine the cause of a packet loss between intra-technology interference, cross-technology interference and weak signal in real-time.

4.3.1 SV-based Detection

We explore the possibility of exploiting variations in demodulated soft values for interference type detection. The core idea is to inspect the modulation and signal shape of the interfered signal, which is reflected by the soft values. While experiencing *Intra-Technology Interference* (interferer is 802.15.4), the demodulator demodulates the stronger signal. Since the interference signal is of the same type (i.e., shape), the variations in the soft values remain small. In contrast, with *Cross-Technology Interference* the signal shape differs from the ideal signal. Thus, the variations in soft values are higher. We take signal samples from the interfered part and compute the variation metric V which we use to determine the received signal type:

$$V = \frac{1}{N} \sum_{i=1}^{N} \left(|SV(i)| - \frac{\pi}{4} \right)^2 \tag{4.3}$$

V measures the average distance between the received $|SV|$ and the ideal value $\frac{\pi}{4}$. The smaller V, the higher the chance that the signal is 802.15.4 (i.e., *Intra-Technology Interference*).

This SV-based detection mechanism does not require to compute complex compensation of channel distortions or signal decoding. Moreover, the soft values are readily available which makes this detection mechanism lightweight. Note that any interfering technology using O-QPSK with half-sine pulse shape and similar baseband signal bandwidth other than 802.15.4 is identified as Intra-Technology Interference using this mechanism. This technique exploits the capture effect phenomenon, in which a strong interfering signal is successfully demodulated (i.e., of the same technology). It, therefore, works well in the low SINR region, where the target signal is much weaker than the interferer. Hence, the interference signal dominates the signature shape in the received signal. If this is not the case, we resort to a more costly technique: correlation-based detection. Although we focus our discussion on 802.15.4 PHY, this approach can be adapted to other wireless technologies that provide SVs.

4.3.2 Correlation-based Detection

The receiver can exploit the fact that 802.15.4 packets start with a predefined preamble and SFD symbols for synchronization, and search for this known signal pattern within the interfered segment by computing the temporal cross-correlation between received signal and the ideal preamble plus the SFD. In case the pattern is present, the receiver can conclude that the interference is of type *Intra-Technology Interference*. Otherwise, it is a *Cross-Technology Interference*. In general, correlation is a typical functionality in standard wireless receivers [101]. To detect the interference type in the received signal, the receiver can perform the cross-correlation between the 802.15.4 preamble (p) and the start of the interfered segment. This approach yields a good performance in theory.

In practice, however, the transmitter and receiver are typically not centered on the same frequency. Hence, there is a small frequency offset (Δf) between the transmitter and the receiver that causes a linear shift in the phase of the received signal. This frequency offset can distort the correlation and needs to be compensated before the correlation process. Standard receivers typically estimate the offset and compensate for it. In the context of CTI, since we are agnostic of the transmission source and do not have access to a decodable pilot or decodable preamble in the interfered signal, it is not possible to compensate the frequency offset of the interference signal, even in the case of *Intra-Technology Interference*. This consequently limits the accuracy and usability of this approach.

An alternative approach is applying correlation in the frequency domain. Since frequency offset in the time domain will translate into sampling offset in the frequency domain, it does not affect the value of correlation, but only shift it. The frequency domain correlation with consideration of the frequency offset can be formulated as follows:

$$c(y, \tau, p) = \sum_{n=1}^{N} P^*(n) Y(n + \tau) \tag{4.4}$$

$$= \sum_{n=1}^{N} P^*(n) \mathcal{F}\left((s(k) + i(k) + w(k)) e^{j2\pi(\Delta f - \tau)kT}\right) \tag{4.5}$$

The ideal preamble is independent of transmitted data and the noise. Therefore, the correlation between the ideal preamble and the noise w is about zero.

$$c(y, \tau, p) = \sum_{n=1}^{N} P^*(n) \mathcal{F}\left(i(kT) e^{j2\pi(\Delta f - \tau)kT}\right) \tag{4.6}$$

$$= \sum_{n=1}^{N} P^*(n) I(n + \tau - \Delta f) \tag{4.7}$$

where $P(n)$ and $Y(n)$ are the ideal preamble signal and the received signal in the frequency domain, respectively. $\mathcal{F}(x)$ stands for the Fourier transform of x. Moreover, s, i, and w represent the signal, interference, and noise in the received signal, respectively. The correlation value is maximized when $\tau = \Delta f$, and the signal is the expected preamble signal ($I(n) = P(n)$). Since Δf is unknown and we cannot compensate it, we compute the correlation for a certain range of τ instead and consider its maximal value as:

$$C(y, p) = \max_{\tau} c(y, \tau, p) \tag{4.8}$$

The range of τ is not large, given that the frequency offset is typically small.

Complexity. Applying correlation in the frequency domain involves transforming a signal from its time representation to the frequency domain ahead of applying the correlation, which can be an expensive procedure for low-power receivers. CrossZig primarily runs the SV-based mechanism for detection and utilizes the correlation-based technique just for the SINR ranges where the SV-based technique does not yield a good accuracy. The SV-based mechanism in its core examines variations in the SVs which makes it a lightweight mechanism that is practical for low-power radios.

4.4 CrossZig Architecture and Design

In this section, we present the detailed design of our approach that extends the 802.15.4 stack to improve medium access efficiency under CTI and recover partially interfered packets. We start by presenting a high-level overview of CrossZig, then present its core components, and finally describe the system integration.

4.4.1 Overview

CrossZig is an extension to the standard 802.15.4 that allows low-power wireless nodes to communicate better in interfered environments. Upon the detection of CTI, CrossZig triggers an adaptive recovery scheme. Our extension is accompanied by a CTI-aware medium access mechanism that

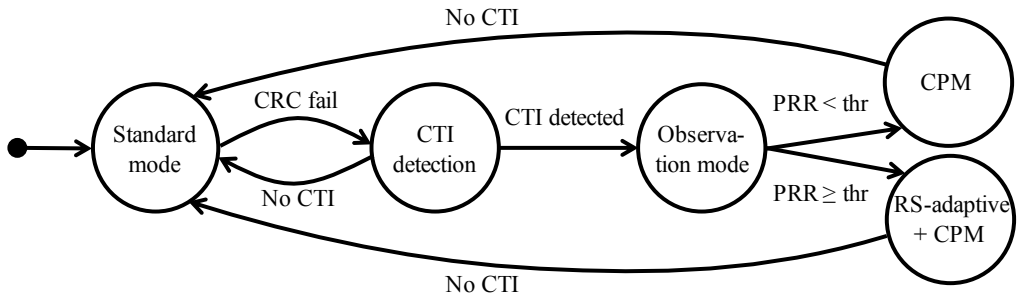

Figure 4.8: Overview of CrossZig transitions.

opportunistically leverages the silence duration in interfered channels. In particular, our extension consists of the following components:

PHY-hints Interface. This interface allows higher layers to access richer physical layer information and consequently use it in a variety of algorithms to boost performance under CTI. The details of this interface are discussed in Section 4.1.

CTI Detection. This component resides in the physical layer. We show (in Section 4.3) how we exploit variations in the PHY hints to detect interference in incoming corrupted packets. More importantly, how we can differentiate the interference type between *Intra-* and *Cross-Technology Interference*. The recovery mechanism presented next are enabled by CrossZig only when CTI is detected; otherwise nodes operate in the normal mode, as depicted in Figure 4.8.

CTI-aware Packet Recovery. Our system recovers from a variety of interference patterns. The receiver estimates errors in interfered packets by relying on physical layer hints. Error information is used to choose a suitable recovery mechanism; currently selecting or combining two recovery mechanisms: *(a)* Cross-layer based packet merging to recover long error bursts, and *(b)* Adaptive error-correction coding to deal with transient interference with low BER rates.

4.4.2 CTI-aware Packet Recovery

CrossZig mitigates CTI through an adaptive packet recovery scheme. It observes error characteristics and adjusts the recovery mechanism settings accordingly. The recovery scheme integrates two recovery mechanisms, namely cross-layer based packet merging and adaptive Reed-Solomon (RS) coding. Cross-layer packet merging tackles long burst errors which are beyond coding recovery capabilities. Adaptive RS coding targets packets that can be recovered with moderate code redundancy, i.e., low bit error ratio. We first explain how these

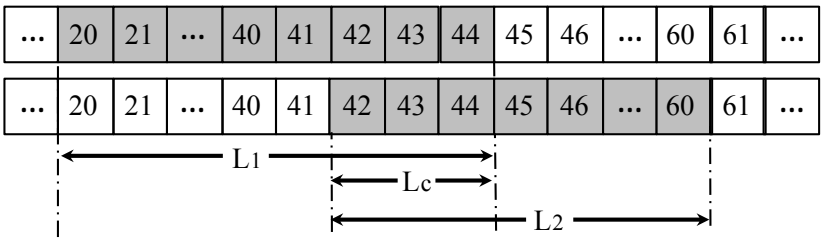

Figure 4.9: Two consecutive corrupted transmissions of the same packet. L_c highlights the overlapped interfered segment in the two packets which cannot be recovered with basic packet merging. We exploit MRC to recover L_c.

mechanisms work independently and later we describe how they are integrated into CrossZig.

Cross-layer based Packet Merging (CPM). CPM is a packet recovery mechanism that aims at improving retransmission efficiency by exploiting partially correctly received bits within the interfered frames, and additionally, combine incorrect symbols from multiple transmissions to construct correct ones.

Packet retransmission is a fundamental mechanism used in communication systems to recover lost packets, typically referred to as Automatic Repeat reQuest (ARQ). Once a received packet is erroneous (i.e., failed the CRC check), the receiver abstains from sending an ACK frame. Transmitter retransmits the same packet until an ACK frame is received for that particular packet. Note that even if the majority of bytes within an interfered packet are correct, packets are discarded due to the bit-by-bit correct transmission enforcement by CRC. In the presence of an active interference source, an immediate retransmission is susceptible to be disturbed by the same source of interference which can degrade the performance of standard ARQ mechanism. Instead of repeatedly retransmitting such packets, we investigate how nodes can accept and buffer corrupted packets, and combine multiple, possibly erroneous, copies of a given packet in an attempt to recover the original packet from the corrupted copies.

Our CPM is realized at two stages; symbol level and signal level. The symbol-level packet merging reconstructs the target packet by combining correct symbols from two packet instances. CrossZig identifies correct symbols using our error localization mechanism. As long as we receive one correct instance of every symbol, this technique allows us to reconstruct the original packet with high confidence (see Figure 4.9). For symbols that are corrupted on all received instances, as depicted in Figure. 4.9, we combine them at the signal level by employing *Maximum Ratio Combining* (MRC).

In MRC [22], each signal branch is multiplied by a weight factor that

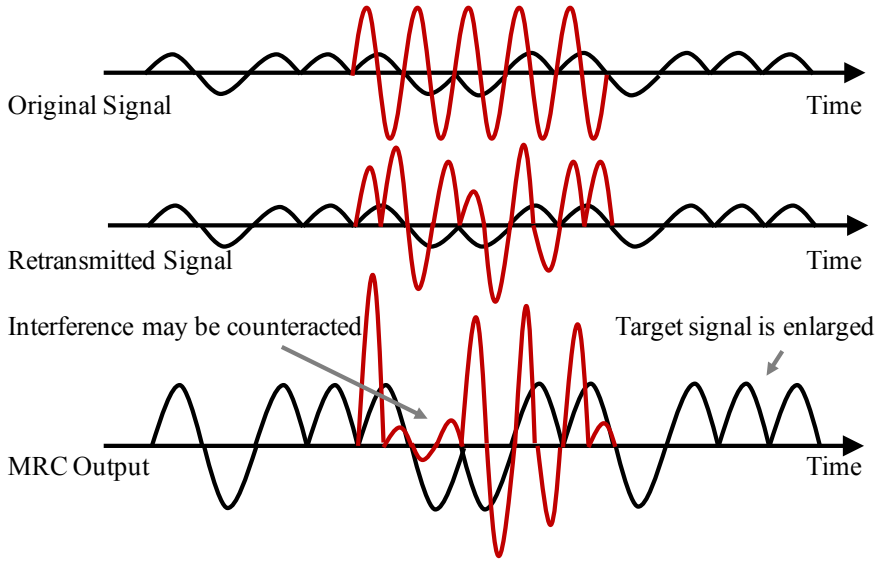

Figure 4.10: Diversity combining of two identical signals under interference.

is proportional to the branch SINR. That is, branches with strong SINR have a larger weighted factor and are further amplified. The signal and noise power are computed over interference-free PHY header and SFD symbols. Thus, interference signal power is derived as the difference between the approximated target signal and noise power from interfered signal power. Signals of the first transmission and the corresponding retransmission can be represented as:

$$y_1(t) = s_1(t) + i_1(t) + n_1(t) \tag{4.9}$$

$$y_2(t) = s_2(t) + i_2(t) + n_2(t) \tag{4.10}$$

The maximal combined signal can be represented as:

$$y(t) = w_1 y_1(t) + w_2 y_2(t) \tag{4.11}$$

where the weighted coefficients (w_1, w_2) are computed by the SINR over the sum of signals of commonly corrupted symbols:

$$w_j = \frac{SINR_j}{\sum SINR_i} \tag{4.12}$$

Time diversity involves transmitting the same information at two distinct times. In Equations 4.9 and 4.10, s_j, i_j, n_j represent the signal, interference, and noise components of the received signal at time instant j. The noise in each time instance of the channel is independent of the signal. The signals s_1 and s_2 are essentially identical. In contrast, i_1 and i_2 are not identical and most probably completely uncorrelated.

After combining, target signal s components are amplified. This yields an increase in the power of the target signal, thus increases its decodability chances. Interference i and noise n components can be either canceled, attenuated, or amplified. However, on average, the SINR is increased (see Figure 4.10).

For the MRC-based combining, co-phasing of all signals is necessary to avoid target signal cancellation. Instead of performing computationally complex frequency and initial phase offset compensations for each signal, we estimate the relative phase offset between two signals. Since they are transmitted and received by the same sender and receiver, their frequency offsets are the same. Thus, after compensating the relative phase offset by utilizing the preamble signals in each packet, we can correctly align them.

RS-Adaptive Coding. When the system observes a high ratio of corrupted packets, FEC, which adds redundant information to the payload, is used to potentially recover the errors and possibly avoid retransmissions. Although FEC codes are widely used in communications systems, selecting the right coding scheme and setting the right level of redundancy for constrained devices is not trivial. We investigate how to derive an adaptive encoding strategy for low-power devices under various interference patterns, where transmitted redundancy is bounded to the inferred error patterns.

We choose *Reed-Solomon (RS) codes*, which are practical for constrained devices [96, 99]. RS codes are systematic codes, i.e., redundancy data is appended to unaltered source data. This results in no decoding overhead when no error is present. The RS codes are block-based error correcting. The length of the redundant parity (t) defines the maximum number of corrupted blocks a receiver can successfully recover within a partially corrupted packet. RS coding can correct up to $t/2$ and detect up to t block errors. It works well for error patterns that fall under the recovery capacity of the parity check.

The primary goal of our adaptive strategy is to increase packet recovery rates, yet minimize the redundancy overhead on the channel to meet the energy constraints of low-power radios. To realize this, we rely on physical layer hints to infer error information, namely, identify and locate corrupted symbols as discussed in Section 4.2. With this in hand, we can calculate further error meta data information such as per packet error rate. This allows us to adaptively derive a redundancy level for our adaptive RS-code based on the symbol error rate in the window of received packets. In case the number of corrupted packets is low (i.e., reasonable *Packet Reception Ratio* (PRR)), CrossZig retain from activating the RS-coding to avoid introducing redundancy overhead in good links.

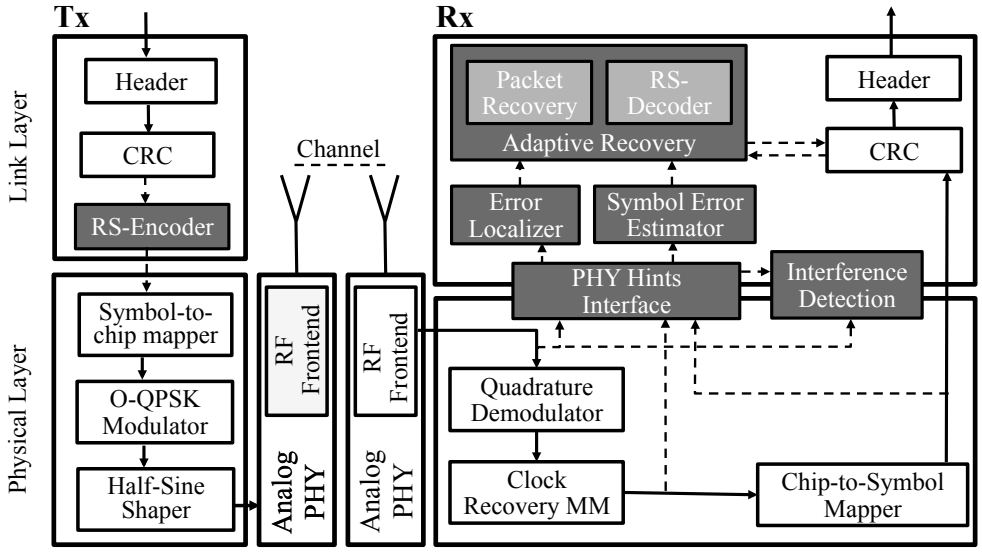

Figure 4.11: Cross-Layer architecture design of CrossZig. Dark gray shaded boxes depict our components added to standard 802.15.4 stack.

Our adaptive algorithm can be described as follow: For the window w_i of observations, we calculate the redundancy level R_i based on the observed degree of corruptions (i.e., R_i is derived from the average number of erroneous symbols per corrupted packet in w_i and is bounded by an upper bound ($R_i \leq \frac{1-PRR_i}{PRR_i} \times$ packet_length)). CrossZig triggers adaptive coding only if R_i is lower than the calculated upper bound. This allows us to ensure that the introduced redundancy is not significantly higher than potential symbol errors to be recovered. We assume that the corruptions in the upcoming packets follow the trend of our current observations. Hence, the window size should be selected carefully. In our evaluation, we noticed that window sizes of 300 ms to 1 s result in a good and stable performance.

4.4.3 System Integration

CrossZig extends the basic 802.15.4 PHY and MAC layers as illustrated in Figure 4.11. It provides a single hop reliable delivery mechanism that can counter the CTI effects. The receiver performs packet detection and decoding in a manner similar to the standard 802.15.4. In case a jump in the signal strength is observed during packet reception and the received packet fails the CRC, the receiver initiates the interference detection algorithm discussed in Section 4.3. If CTI is detected, the transmitter adapts the *Channel Clear Assessment* (CCA) threshold to allow an opportunistic access to the channel. Upon the reception of few partially interfered packets, the system adjusts the initial recovery settings for the

next observation window w. Any notable changes in the observed error characteristics trigger the system to adjust the recovery settings.

CrossZig performs the following logic while adjusting the recovery settings. Packet retransmission always carries a fixed, low level of redundancy code. This is used to boost the performance of our CPM: the packets are merged, lowering the BER to a level recoverable with FEC. In the case of high packet corruption levels (low PRR, in our settings lower than 75%), RS-adaptive coding is triggered. Here, the redundancy R is derived adaptively based on the observed degree of corruptions. R is derived from the average number of erroneous symbols per corrupted packet in the observation window. Hence, the coding is adapted according to the dynamic interference patterns in the channel. When a packet cannot be recovered with the current redundancy level, CPM is applied.

4.5 Implementation

We build a prototype of CrossZig using SDR. For the SDR hardware, we rely on the USRP-N210 [44], equipped with an SBX radio daughterboard [43] as radio front-end. The SBX board incorporates a wide band transceiver that operates from 400 MHz to 4400 MHz, i.e., covers the 2.4 GHz band. For development, we use the GNURadio [53], an open source software toolkit for building software radios.

The transmitter and the non-coherent receiver nodes run 802.15.4 PHY and MAC layers [81]. We modified the receiver PHY to incorporate interference detection logic, error estimation, and channel estimation in our codebase, as described in Section 4.4. Moreover, we implement an RS-decoder and the CPM scheme at the receiver side. At the transmitter side, we incorporate the RS-encoder. We implement a virtual feedback channel at host software to carry the receiver feedback to the sender. Note that in our SDR prototype implementation of CrossZig we do not use carrier sense. USRP radios introduce inevitable delays into the processing path of packets, which makes confining with carrier sense strict timing requirements hard to realize [111]. This constraint, however, does not hinder us, as the opportunistic access to the medium in interfered channels is possible without carrier sense. While this is not an optimal solution, it is sufficient to manifest empirically the concepts covered in this chapter.

Cross-layer Packet Merging. The standard MRC is carried out on complex signal samples and requires coherent combining at the receiver. In the micro-evaluation of CPM covered in Section 4.6.2, we perform the signal alignment offline ahead of the MRC step (trace-based evaluation).

This is necessary as our prototype implementation is based on non-coherent receivers. Thus the receiver does not require signals to be synchronized in phase and frequency. To cope with the lack of phase offset compensation in our prototype (i.e., non-coherent receiver), we carry out MRC on the demodulated SVs instead of the complex signal samples. This only applies for the system performance covered in Section 4.6.1. Given that the quadrature demodulator measures the phase difference between two successive input signal samples, the initial phase offset is no longer an issue. The demodulated soft values of the first transmission and the corresponding retransmission can be represented as follow:

$$y_1[n] = SV_{ideal} + \delta_1[n] \tag{4.13}$$

$$y_2[n] = SV_{ideal} + \delta_2[n] \tag{4.14}$$

where $y_1[n]$ and $y_2[n]$ are the soft demodulated values for the n-th symbol in two transmissions, $SV_{ideal} = \pm\frac{\pi}{4}$ is the ideal soft demodulated value for our target signal and $\delta_1[n]$ and $\delta_2[n]$ are the errors caused by interference and noise. Since interference and noise in different transmissions are independent and identically distributed (i.i.d.) with zero means, by weighted averaging of the soft value we increase the chances of successful demodulation.

4.6 Experimental Evaluation

Now we present the experimental evaluation of our prototype implementation on the USRP-N210. In the following, we first define our evaluation objectives and describe the experimental methodology, the considered interferers, the evaluation setup, and the metrics. We continue with a discussion on the system's online performance, followed by a detailed evaluation of the system components, namely CTI detection, error localization algorithm, and the MRC-based packet merging.

Methodology. The ideal experiment setup would evaluate the end-to-end performance of CrossZig using real traffic models with different prominent low-power MAC protocols. However, due to inevitable processing latencies in current software radio platforms, the realization of such an evaluation setup is hard or not feasible with regard to strict time constraint components. Instead, we focus on link performance, by measuring the packet reception rate for various communication links that we subject to external interference sources. Note that the performance degradation under CTI is primarily attributed to starvation

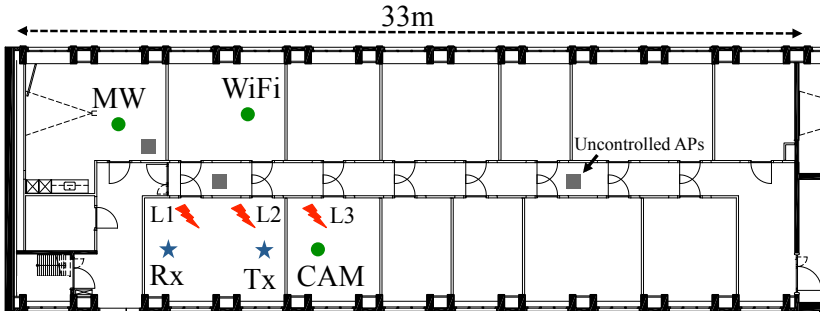

Figure 4.12: Layout of the online evaluation experiment setup. SDR 802.15.4 Rx-Tx located in an office with a line-of-sight link of 5 m. The interferers are located at locations L1 (1 m), L2 (4 m), and L3 (7 m), where L1 and L2 are in line-of-sight to Rx-Tx and L3 is in non-line-of-sight. Green circles indicate the location of our multiple interferer scenario. Gray squares indicate the location of uncontrolled access points placed on the floor.

or/and discarded corrupted packets. Packet losses, where packets are not successfully detected, account less to the overall performance degradation, and are not directly addressed in this work [76]. Such losses can be resolved by considering better packet detection mechanisms as suggested by [96, 85]. In the second part of this section, we cover the evaluation of individual components of CrossZig. Note that all system components are evaluated empirically. Additionally, in the micro-evaluation, we support part of the empirical results with Matlab simulations, e.g., to show the algorithm's behavior under SINR ranges beyond the empirically-captured ranges.

Interferers. Our set of interference sources includes low/high power, narrow/wide band, channel hopping/fixed frequency, and CSMA/non-CSMA. This represents common underlying properties adopted by most radio technologies. More specifically, as CTI we consider 802.11 (heavy and light UDP traffic), digital wireless camera, and microwave oven. 802.15.4 is considered as *Intra-Technology Interference*.

Evaluation Setup. The system evaluation is performed in ETH Zurich's computer science building. Figure 4.12 shows the layout of the experimental setup. Experiments are carried out with controlled single active interferers mentioned above and multiple active interferers. Multiple active interferers are different combinations of single interferers running simultaneously and defined as *Multipe-1*: microwave oven and wireless camera running simultaneously, *Multipe-2*: microwave oven, wireless camera, and 802.11 with light UDP traffic, and *Multipe-3*: microwave oven, wireless camera, and 802.11 with heavy UDP traffic.

The 802.15.4 transmitter-receiver pair was represented by our prototype implementation on USRPs. During the experiments, the 802.15.4 communication link was also exposed to interference from various uncontrolled sources existing in the building. In each experiment, we transmit 6000 packets consecutively with 60 Byte payload, at a 10 ms interval.

Metrics. Within our evaluation we use the following metrics: *(a) Goodput ratio:* defines the ratio of useful received data over total received data. It quantifies the system's efficiency as it reflects both the gain and the transmission overhead together. This metric allows us to observe how efficiently transmitted bytes are utilized. *(b) FEC overhead:* indicates the added transmission overhead which is directly related to energy efficiency, a vital factor in low-power networks. *(c) Packet recovery ratio:* indicates how many of the corrupted packets our recovery mechanisms could recover. The recovery ratio and redundancy overhead show the performance of the considered schemes compared to the baseline where no mitigation scheme took place. Note that in our definition, the basic scheme has 0 recovery ratio and 0 cost. *(d) Precision and Recall:* values are relevant for the performance discussion of symbol error detection, where the selection of parameters has an impact on the performance. Precision indicates how many of the identified corrupted symbols are indeed corrupted. Recall indicates how many of the overall corrupted symbols are identified. *(e) Symbol Error Rate (SER):* is the number of corrupted symbols over the total number of symbols in a received packet.

4.6.1 System Performance

We expose CrossZig first to single active interferers at different distances. The interferers are located first at location L1, then L2, and L3 (see Figure 4.12). Second, we consider interference generated from multiple simultaneous sources. Figure 4.13 shows the evaluation results achieved by the following recovery schemes: RS-coding with fixed redundancy of 30 Byte, our adaptive coding scheme which selects a redundancy between 0 and 30 Byte based on the average observed SER in the 500 ms window of observations (irrespective of the PRR in the channel), packet merging, and finally CrossZig which combines our cross-layer based packet merging and our adaptive RS-coding scheme.

The error patterns caused by interferers vary as we change the interference types. Therefore, different experiment settings yield varying performance in terms of goodput ratio, packet recovery ratio, and redundancy overhead. The RS-fixed scheme achieves the highest packet recovery ratio, but this comes with a fixed 30 Byte redundancy per packet,

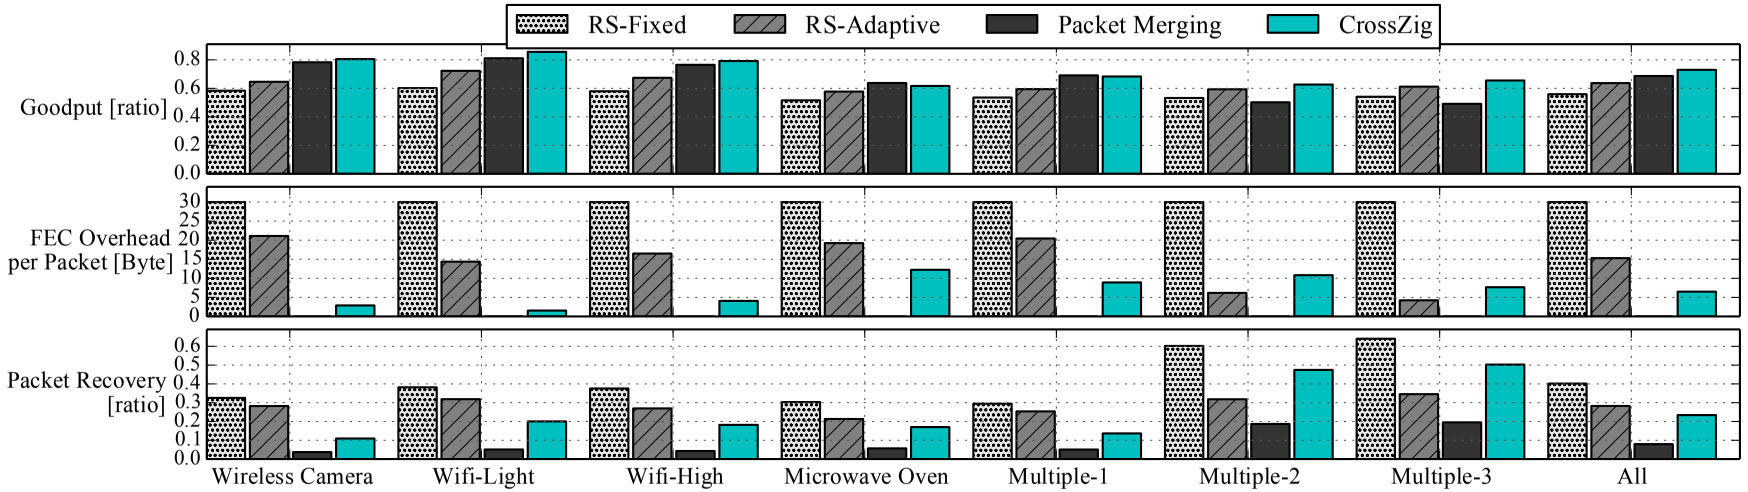

Figure 4.13: Online performance of CrossZig exposed to various types of interferers. The upper plot depicts the goodput ratio for the considered recovery mechanisms. It conveys how much of the transmitted data is useful data, and hence, reflects how well the channel is being utilized. The middle plot depicts the varying redundancy overheads, which affect the radio on time. The bottom plot shows the recovery ratio without considering the cost. CrossZig achieves the highest goodput in most cases, thanks to its adaptivity and a favorable balance between overhead and recovery ratio.

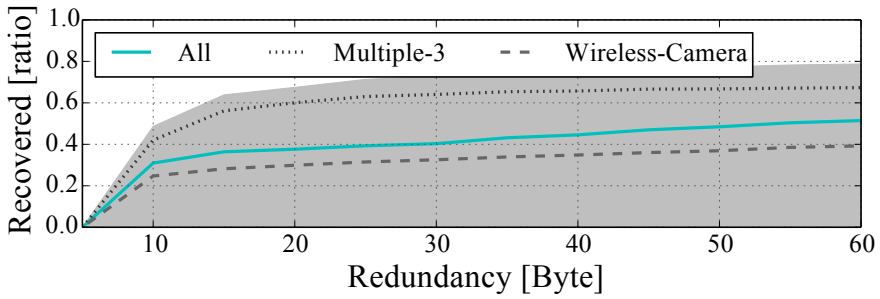

Figure 4.14: The implication of applying fixed levels of redundancy under CTI. The gray-colored area depicts the ranges for recovery ratio under varying interferers per technology.

regardless of channel conditions. This has a negative impact on goodput. This overhead exacerbates for good quality channel conditions which we did not consider in this study. This results in higher in-air time and increased processing for decoding at the receiver side, which are both undesirable for low-power devices. The RS-fixed scheme evaluated in Figure 4.13 considers 30 Byte redundancy. Figure 4.14 depicts the recovery ratio at different fixed redundancy levels. With a redundancy higher than 20 Byte we do not observe a notable improvement of the recovery ratio. Note that increasing the redundancy has the side-effect of increasing the probability of overlap with interference, hence, reducing the effectiveness. With our RS-adaptive scheme, we observe a similar packet recovery ratio as with the fixed strategy, but at a lower overhead (in average 15 Byte for each packet). This yields a higher goodput. Packet Merging comes with no FEC overhead because it simply works on the received signal of incoming packets. Its recovery ratio is modest in most cases, except in the presence of multiple interferers, because Packet Merging is particularly effective at higher SER levels.

CrossZig improves RS-adaptive coding which relies only on the observed SER rates for adaptation. In addition, CrossZig recovers long error bursts using Packet Merging and is able to keep its cost low under sparse interference. We reach an average packet recovery ratio of 23% overall and up to 50%, for the multiple-3 setup. This is about half of the average packet recovery ratio achieved with the aggressive RS-fixed (40%) over all cases. However, the overall overhead of CrossZig is by a factor of 4.6 lower than the other schemes and reaches up to a factor of about 20 for the case of WiFi-light. As a result, CrossZig achieves the highest goodput ratios in most scenarios. Note that for fairness we did not compare the performance of CrossZig to the case of no active interferer, where goodput falls drastically for RS-fixed and improves to even higher values for CrossZig.

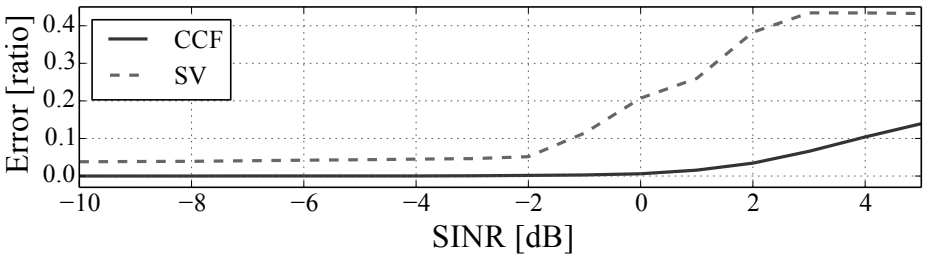

Figure 4.15: Error ratio in discerning the type of interference between CTI and 802.15.4 achieved by the SV-based method (SV) and the correlation-based method (CCF).

Conclusion. We show that performing timely adaptation to match induced error patterns from external interference is possible with the help of physical layer hints. With this timely adaptation, we can achieve better goodput and avoid excessive redundancy which comes at a high price for low-power devices.

4.6.2 Dive in CrossZig

We carry out an offline micro-benchmark analysis of CrossZig to quantify the performance of its individual components independently. Our traces for this evaluation include the complex signal of 35,875 packets corrupted by interference.

4.6.2.1 CTI Detection

We now discuss the performance of our CTI detection scheme introduced in Section 4.3. We estimate the effectiveness of our scheme in detecting the occurrence of *Intra-* and *Cross-Technology Interference*. Figure 4.15 shows the detection error ratio for both the SV-based and correlation based detection mechanisms. For low SINR ranges under -2 dB, both mechanisms perform well with error rates below 5%. SV-based detection performs well at low SINR because in the case of intra-technology interference, the interfering signal can be demodulated by the receiver and this is reflected in lower variations of the soft values. As SV-based detection is the cheaper mechanism, we rely on it for SINR under -2 dB. As the SINR increases (weaker interferer), detecting the source of interference is more challenging. The accuracy of the SV-based scheme degrades sharply, while the correlation-based detection still yields error rates below 10%. Therefore, for SINR greater or equal to -2 dB, we use correlation-based detection. Note that for SINR ranges above 3 dB, the interference signal is very weak and, hence, the target signal is decodable. Consequently, CTI detection is not required for these SINR ranges.

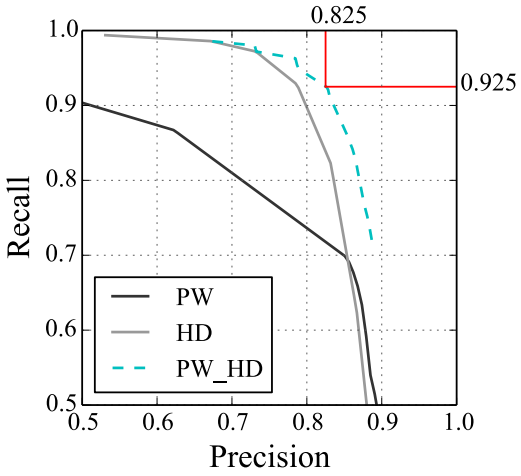

Figure 4.16: Precision-Recall analysis for symbol error estimation with Power (PW), Hamming Distance (HD), and combination of both (PW_HD). Precision indicates how many of the identified corrupted symbols are indeed corrupted. Recall indicates how many of the overall corrupted symbols are identified.

In our system, we select the threshold parameters to balance false positives and false negatives in a reasonable way. We use threshold values τ_{sv} = 0.15 for SINR lower than -15 dB and τ_{sv} = 0.3 otherwise. Figure 4.15 shows the detection performance using these thresholds. Note that misclassification of 802.15.4 communication as CTI can trigger CTI-recovery for *Intra-technology Interference*. On the other hand, detecting CTI as internal would translate to the default behavior of current systems. Thus, no further harm is introduced. Therefore, the thresholds are selected to balance the precision and recall with the goal of higher accuracy in detecting 802.15.4 communication. The average false positive ratio (802.15.4 interference mistaken as CTI) is 0.16 and the average false negative ratio (missing to detect a CTI) is 0.23. Thus, the interference type detector is sufficiently accurate for our purpose. Note that this analysis considers per packet detection.

4.6.2.2 Symbol Error Localization

We now discuss the performance of our error localization algorithm introduced in Section 4.2. Figure 4.16 shows the precision and recall of the symbol error detection mechanism using signal power only, decoding Hamming distance only, and using them jointly. This result is aggregated over all the collected traces. The line corresponds to precision and recall for various thresholds (τ_l, τ_h, and τ_s). For our system, we select the thresholds that yield a good balance between precision and recall in the micro analysis, τ_l=4, τ_h=10, and τ_s=4.

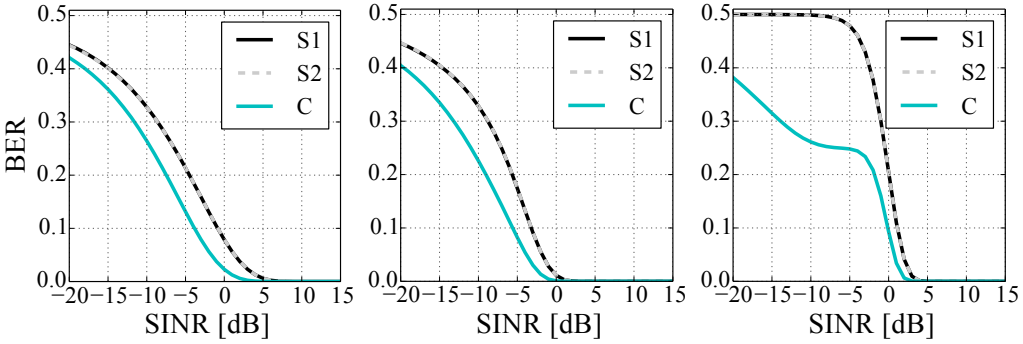

(a) GRN. (b) Inter: QPSK. (c) Intra: 802.15.4

Figure 4.17: Simulation performance of Diversity Combining (C) for 802.15.4 signals (S1 and S2) averaging over 1k independent cases subject to Gaussian random interference (GRN), QPSK interference, and 802.15.4 interference.

By combining power and Hamming distance, our symbol error detection approach yields a stable performance with a precision and recall of 82% and 92%, respectively. The average achieved accuracy is 94.3% ± 2.4.

4.6.2.3 Diversity Combining under CTI

In this section, we investigate variables that impact MRC performance under CTI which is utilized in our CPM. Moreover, we investigate to what extent MRC can increase the symbol error recovery probability and put this into the context of recovering packets with bursty errors. In the context of this work, we exploit time-diversity by combining two interfered copies of the same signal received in different instants of time. We employ the MRC technique for combining the signals. MRC amplifies the SNR of the target signal. The SNR of the combined signal y_c is by factor 2 higher. Therefore, the theoretical SNR gain is 3 dB.

To understand the impact of the interference on the performance of MRC, we first carry out simulations in Matlab. Figure 4.17 plots the Bit Error Rate (BER) vs. SINR for an interfered 802.15.4 signal before and after MRC. We consider three types of interference here: QPSK signal representing CTI, 802.15.4 representing internal interference, and Gaussian random interference. The time diversity gain from MRC is reflected in the BER drops. As we can see, the MRC gain varies with respect to the type of interference signal. Under interference, the gain can exceed the 3 dB expected gain. MRC performs better when the interference signal has an underlying modulation scheme as opposed to noise.

This observation is aligned with our empirical results carried out with the trace-based evaluation. There, the MRC gain for the wireless camera

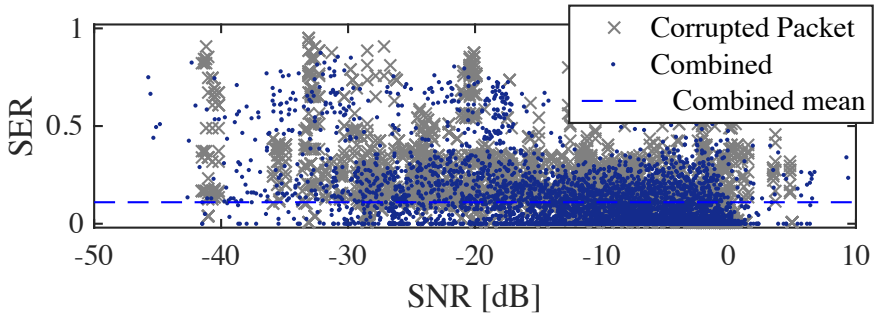

Figure 4.18: Our cross-layer based packet merging mechanism reduces the average SER per packet to 0.11.

and 802.11 is higher than that for the microwave oven (which is noise radiation). In practice, with MRC we can increase the recovery chance of a symbol by up to 15% which is 2.5x times higher than the random guess. To see how this is reflected in our cross-layer based packet merging, we extend our evaluation to packet level. Figure 4.18 shows the results of CPM applied to 2 consecutive corrupted transmissions of a packet in our traces. The outcome shows that the mean SER of combined packets is reduced to 0.11 which has good chances to be recovered by low FEC coding on top. Our cross-layer based packet merging can achieve an overall gain of up to 0.34 in SER.

4.7 Related Work

Wireless interference is (and has long been) an important topic in wireless communication research. Recent years have seen significant and fundamental contributions to the state-of-the-art interference management, for instance by techniques like interference alignment [27] or joint/coordinated transmission [57, 52]. Nevertheless, these approaches typically require significant computational complexity and/or significant coordination bandwidths, which hinder them applicable for low-power, low-complexity devices of interest in this dissertation. Hence, in the following we focus on interference mitigation in the unlicensed bands and work related to CrossZig.

Physical Layer Information. This work is not the first to use physical layer information. This information has been used for various purposes such as rate adaptation [153, 125], interference boundary detection [54], and packet recovery [56, 85], to mention few examples. To the best of our knowledge, this work is the first to utilize this information for interference type detection.

Interference Avoidance. Research in this direction aims at *detecting and avoiding interfering signals in space, time, or frequency*. The most common avoidance approach is to employ frequency-based isolation by employing spectrum sensing to identify interference-free channels [126, 162] or adaptive frequency fragmentation [29, 162]. The lack of interference-free channels and the fast and unpredictable changes in the occupancy state of frequency bands make the sampling overhead of these approaches high, particularly for resource-constrained devices. Huang et al. [80] and Boano et al. [21] proposed approaches to avoid interference in time by learning transmission characteristics and the idle cycles of interferers. Radunović et al. [123] proposed an adaptive preamble design to increase the probability of detecting low-power transmissions by high-power competing technologies.

Packet Recovery. Research in this direction aims at *increasing resilience against interference* by bracing PHY and data-link layers with auxiliary mechanisms. For instance, Liang et al. [96] studied the interplay between 802.11 and 802.15.4 and applied a resilience forward error coding scheme against interference. Analogously, some solutions focused on exploiting the temporal effects of interference induced on PHY hints, such as variations in soft errors (softPHY) [85, 159] or RSSI variations [66] to localize interfered segments, hence, adapt standard ARQ to retransmit only the interfered segments.

Interference Cancellation. Further physical layer solutions, such as Interference Cancellation, have been considered to combat interference [36, 90, 56]. Here the receiver, with minimal or no coordination from the sender, attempts to recover the signal of interest from interference. Halperin et al. [62] utilized Successive Interference Cancellation (SIC) to recover from collisions. The key idea of SIC is that interference signal and target signal are decoded successively. First, the receiver decodes the interference signal, i.e., the signal with larger power, afterward the interference signal is stripped away from the aggregately received signal to get the target signal. Note that these techniques require knowledge about the interfering signal modulation scheme, which makes them not suitable for CTI. Gollakota et al. [54] proposed TIMO, a MIMO design that enables 802.11n to communicate in the presence of CTI. TIMO exploits MIMO capabilities to cancel the interference signal. However, low-power wireless devices are typically single antenna devices, where such approaches are not applicable.

Collision vs. Fading. Several recent schemes have been proposed to recognize the type of losses in the channel between fading and collision [136, 71, 129, 153, 66]. Many of these mechanisms have been

utilized to boost the performance of loss-based rate adaptation schemes. In COLLIE [129], the transmitter distinguishes between a fading by analyzing the patterns of bit errors in received packets. SoftRate [153] utilizes SoftPHY to distinguish between collision and fading for rate adaptation. AccuRater [136] compares constellation dispersions of the preamble and postamble to detect a collision.

Interference Classification and Signal Detection. Research in this direction aims at *identifying the type of interference technology*. The lack of interference-free channels led researchers to work on novel classification approaches that make networks aware of the type of the existing interference [130, 68, 78, 34]. It has been shown that when the interference source is known, specialized mitigation approaches can improve the network performance. Researchers explored signal properties by employing signal classification techniques [93] or featuring distinct interferer's patterns on corrupted packets [68] to build interference classification tools. It is not clear though how these classifiers can be utilized in a systematic way to combat interference. In the previous chapter, we address this limitation and propose a system that employs a lightweight machine learning classifier to map the current channel signature to a coexistence strategy. However, this approach requires prior training of the adaptation algorithm which might not always be feasible. Analogously, signal detection techniques [78, 50] for spectrum sensing are important requirements in cognitive radio networks. These techniques enable detection of unused spectrum and sharing of it without causing harm to primary users. This direction has been widely explored in cognitive networks with the focus on detecting known signals in noise. On the contrary, in this work we focus on detecting the type of signal in interfered segments of the packets. Hence, the focus is on signal detection in mixed signals (i.e., interfered signals) where the target signal is mixed with an unknown signal.

Exploiting CTI in Low-power Networks. Recent research efforts focused on exploring opportunities in CTI. For instance [137, 55] harness CTI to beneficially provide security and privacy-preserving counting [98]. Others [164, 87] harness channel overlapping between 802.15.4 and 802.11, to allow cross-talk to dispense the role of a dedicated gateway to interconnect these two technologies. CTI is inevitable, hence, utilizing it to provide additional services will enhance spectrum usability. This direction of research is orthogonal to interference mitigation, which is the focus of this work.

Our Approach. Analogously, our work features physical layer hints to infer and recognize interference patterns and harness this to adapt

the recovery mechanism. We propose a solution that neither requires interactions with the interfered technology nor depends on any prior training of the adaptation algorithm and is agnostic to the interference type. Finally, our system is related to prior work on cross-layer wireless design [162, 54, 115, 153, 85]. However, our system is optimized to address CTI in low-power and low-complexity radios.

4.8 Summary

Operating in dense and diverse spectral environments demand wireless systems to attain a high degree of flexibility and be perceptive to the changes in their environment. By their nature, wireless signals encode rich information that can be harnessed to better understand the RF ambient. In this chapter, we investigate how to exploit fine-grained physical-layer information to increase receiver's cognition of the channel dynamics. Consequently, better reason and adapt parameters to recover from CTI in a low-power environment. This chapter presents CrossZig, a CTI-aware adaptive recovery mechanism. Our system combines interference detection, error localization, and an adaptive error recovery mechanism. We do not restrain ourselves with off-the-shelf radios and resort to SDR for our prototype implementation. Experimental results show that our approach can substantially improve the goodput of 802.15.4 links under various CTI patterns. Moreover, we anticipate that the analysis, insights, and discussions carried out in this chapter can inspire further work to address low-power coexistence unconstrained by current chip designs.

5

Wireless Coexistence Experimentation

Wireless networks deployed in indoor environments are susceptible to link quality deterioration due to changes in their environments. One major hazard affecting wireless link quality today is interference. In the recent years, we have witnessed a rapid surge in wireless data traffic congesting the unlicensed bands. This traffic is generated from heterogeneous radios that follow different protocols and communication primitives. To date, the 2.4 GHz ISM band is by far the most congested segment of the radio spectrum. Networks operating in this band have to compete with co-located transmissions from WiFi (IEEE 802.11), Bluetooth, IEEE 802.15.4, 2.4 GHz cordless phones, surveillance cameras, game controllers, and 2.4 GHz RFID, as well as with noise generated by the microwave ovens. Several independent academic and industrial studies [7, 76, 54, 103, 155] show that wireless networks and RF-based systems operating in this band experience serious performance degradation due to interference.

As an increasing number of devices share the unlicensed bands [92], it is crucial to understand how interference impacts the performance of wireless networks and emerging pervasive RF-based services, such as indoor localization [165, 4, 137] and activity recognition systems [122, 154]. Particularly as we find more deployments of wireless networks and services in critical domains such as the health sector [89, 10, 5], and smart grids [25]. Developers cannot provide useable health solutions if their functionality is dependent on the state of the uncontrolled radio interference in their surroundings.

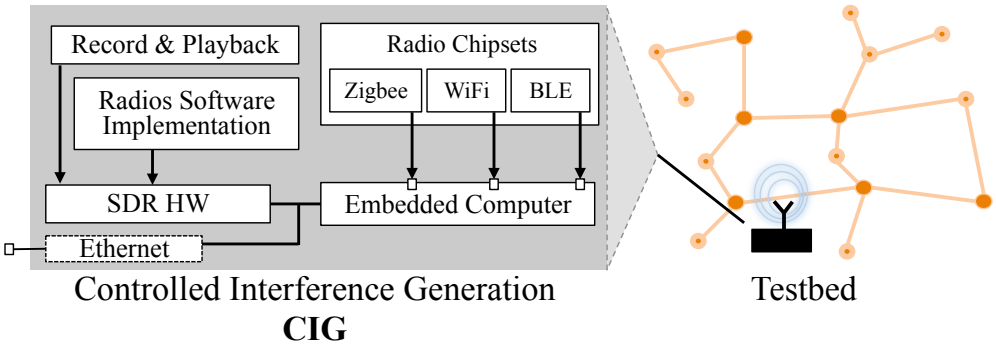

Figure 5.1: Schematic of our Controlled Interference Generation (CIG) framework, facilitating advanced wireless coexistence experimentation.

Patients wearing body sensors that monitor their vital functions cannot rely on instant alarms of their medical state if sensors are blocked by nearby radio communications.

Considering that radio interference has a non-negligible impact on the dependability of wireless networks. It is important that system designers attain a detailed understanding of how heterogeneous wireless systems and networks coexist and operate in the crowded unlicensed spectrum. The lack of proper testbed infrastructures that support generating repeatable and realistic interference patterns makes it challenging for researchers to test the robustness of their wireless protocols in the presence of interference. Hence, the goal of this chapter is to design and develop low-cost experimental facilities that allow researchers and system designers to replay realistic interference patterns in a simple and efficient way, consequently facilitate research in this direction.

Challenges. The impact of interference is highly stochastic in nature, and largely depends on the surrounding environment (e.g., on interferer location, traffic patterns, device manufacturer, and the number of active devices). Researchers working on wireless coexistence, typically use analytical models and simulation [84, 109], which are often abstract, less accurate, and fall short on capturing the complexity of the involved physical aspects, or alternatively use interference generated from actual wireless devices [54, 77, 71]. While the latter approach is more realistic, it is costly, labor intensive, and impractical as some of these devices can not be controlled in a systematic way (e.g., microwave oven, analog phone, etc.), especially when experiments are run in remote testbeds. Recently researchers have been working on augmenting testbeds with commodity hardware that is dedicated to generating controllable interference. For instance, Jamlab [20] is a recent approach, which makes use of

commodity hardware by utilizing a subset of the nodes in the testbed to generate controllable interference patterns. However, such systems have shortcomings in accuracy and the range of interference types they can support. Due to hardware limitations, such approaches are restricted to the fixed modulation schemes supported by the nodes used in the testbed and limited to the rate at which frequency hopping can be performed.

Contributions and Roadmap. In this chapter, we present Controlled Interference Generation (**CIG**), a Software Defined Radio (SDR) based solution for controlled interference generation, which can facilitate augmenting current testbeds with repeatable and realistic interference pattern generation (see Figure 5.1). CIG provides three modules for interference generation: *(i) Record and Playback*; this module features high precision record and playback. It can be used to record and playback various interferer patterns, but is particularly interesting for devices that are not feasible to be implemented in SDR, such as microwave ovens, and proprietary radios where we lack the know-how on their physical layer implementation. *(ii) Radio Software Implementation*; this module allows generation of interference from radios (i.e., the physical layer) implemented in software. For this, we implement or port the implementation of software radio of a set of prevalent interferers using GNU Radio and Universal Software Radio Peripheral (USRP). This set includes commercially available analog cordless phones, digital FHSS phones, security cameras, baby monitors, WiFi, and ZigBee devices. *(iii) Commercial Radio Chipsets*; this module allows the generation of interference patterns from a subset of commercial radio chipsets that are interfaced with an embedded computer within CIG. This further allows us to cover commercial software and hardware artifacts of different radio chipsets.

CIG is not bound to the set of interferer technologies presented in this chapter, and each of its modules can be extended to include new technologies. We provide a unified, simple to use interface for controlling CIG. We perform an initial validation of the generated interference patterns by correlating the generated and real interference in time and frequency domains. Furthermore, we analyze the impact of CIG generated interference on low-power networks to ensure accuracy and resemblance to the interference generated by actual RF interferers. Moreover, we provide insights on limitations and challenges of bringing some commercial radios to SDR.

In light of our contributions, Section 5.1 and 5.2 present the design and realization of CIG, Section 5.3 presents the system validation, and Section 4.8 provides brief concluding remarks. This chapter is based on the contributions made in [72].

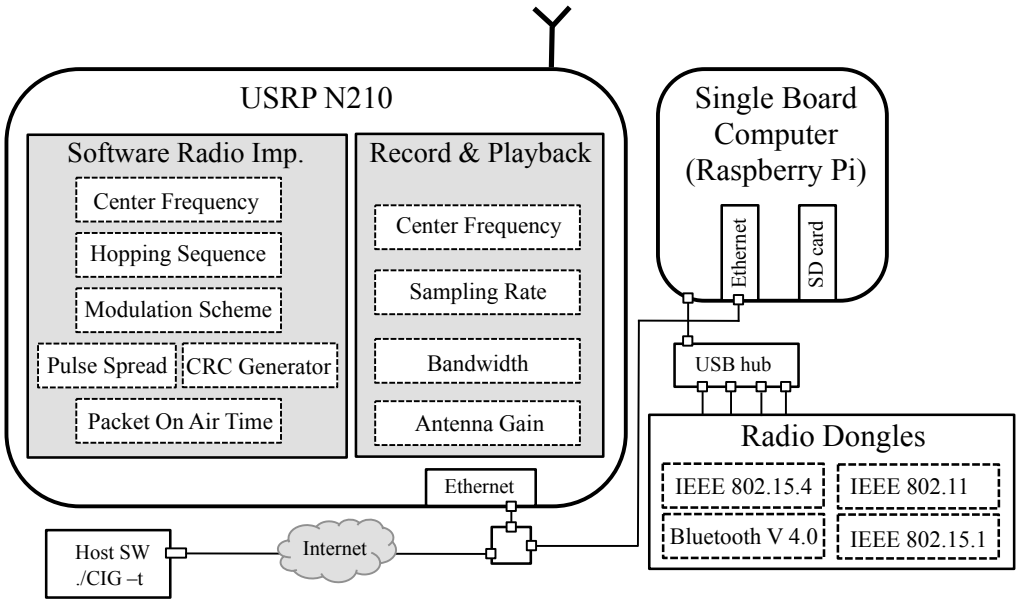

Figure 5.2: Architecture of CIG. Users can connect and control CIG from a remote host. The *Software Radio Implementations* and *Record and Playback* modules of CIG reside on USRP N210. The third module resides on a single board computer that enables the generation of interference from off-the-shelf radio dongles.

5.1 CIG Design Overview

We now present a high-level design overview of CIG, as illustrated in Figure 5.2. CIG consists of three modules for generating controllable interference. In the current prototype of CIG, we focus on incorporating a set of interferer technologies that are prevalent in the unlicensed bands. Our considered set of interferers covers low/high power, narrow/wide band, analog/digital, and channel hopping/fixed frequency interferers. This set represents common underlying properties adopted by most radio technologies.

Record and Playback. This module of CIG is realized on SDR and allows recording temporal and spectral patterns of a particular interference and playing back these patterns as energy pulses emitted in the spectrum. For a large body of interference mitigation research, particularly solutions residing in MAC and upper layers (e.g., clear channel assessment, interference avoidance, channel sampling for free channel discovery, and channel occupancy patterns for opportunistic MAC scheduling) it is sufficient to focus on temporal and spectral characteristics of interferers. The modulated signal type thereby is of less relevance. Moreover, interferers that are not inherently RF radios, such as microwave ovens or closed radios, which cannot be implemented on SDR, are appropriate candidates to be represented through the playback module of CIG.

Software Radio Implementations. This module allows interference generation of a set of prevalent interferers. We enable this by implementing physical layers of these interferers in software while aiming to achieve an authentic physical layer behavior. This module can be used while developing interference mitigation schemes where the type of modulated interfering signal is relevant. This is particularly relevant to physical layer solutions, such as interference source classification [78], interference suppression, and cancellation [54]. Moreover, it allows verifying whether emerging radios [100] and wireless systems can cause harm for competing technologies and quantify the impact.

Commercial Radio Chipsets. Reaching hardware-like efficiency and predictability with the software implementation of radios on SDRs is challenging and not always feasible. With this module, we have the possibility of generating interference from standard off-the-shelf radio chipsets. Thus, it allows covering the impact of commercial software and hardware artifacts of different radio chipsets and overcoming limitations of the SDR platform. The restrictions of the class of SDR platforms we employ in CIG are twofold: *(i)* Due to strict timing requirements, carrier sensing is hard to implement in software (e.g., 802.11 backoff). *(ii)* Due to strict frequency tuning capabilities, it is hard to achieve high frequency hopping rate in software (e.g., Bluetooth exhibits a hopping rate of 1600 hops/s). Note that more affordable and capable SDR platforms are populating the radio market every year. Therefore, we anticipate that future generations of affordable SDR platforms will overcome these limitations.

5.2 Realization

In this section, we elaborate on CIG's hardware and software architecture. We first give a brief overview of our platform and then discuss implementation aspects of the modules.

5.2.1 Platform

The hardware platform of CIG consists of two main components (see Figure 5.2). The software defined radio based component, is where the *Record and Playback* and *Software Radio Implementations* modules of CIG are realized. The second component is a low-power computer that controls the *Commercial Radio Chipsets*. We provide a unified interface in the form of extendable scripts that interact with the corresponding CIG component to generate interference. Users can use the interface to connect and control a remote CIG instance, located in a testbed.

Software Radio. For the SDR hardware, we rely on the Ettus USRP N210 [44], which is equipped with 100 M samples/s 14-bit ADCs and 400 M samples/s 16-bit DACs. It is connected to a host computer via a Gigabit Ethernet port and can stream up to 25 M samples/s to/from host applications. For the RF front-end, we use the SBX radio daughterboard [43]. The SBX board incorporates a wide band transceiver that operates from 400 MHz to 4400 MHz. It provides up to 40 MHz of instantaneous bandwidth and up to 100 mW of transmission power.

For development, we rely on GNU Radio [53], an open source software toolkit for building software radios. GNU Radio provides libraries for signal processing blocks to implement software-defined radios that can be coupled with generic radio platforms. In order to build a typical wireless radio stack, flow graphs, composed of a sequence of *Digital Signal Processing* (DSP) blocks, are created (see Figure 5.3). Moreover, a state machine selects the corresponding flow graph to process incoming samples. These DSP blocks are created in C++ and connected in a python wrapper to build the flow graphs. For example, the receiver of a DSSS analog phone has blocks for clock synchronization, channel equalization, Costas loop for phase and frequency correction, BPSK demodulator, symbol to constellation mapper, and direct-sequence despreader. Different blocks are integrated into separate flow graphs, each addressing different communication tasks, such as acknowledgment packets, and inbound and outbound communication. In the last step, the flow graphs are assembled into a DSSS cordless phone receiver state machine.

Embedded Computer Board. We use a Raspberry Pi as a single-board embedded computer which hosts a quad-core ARM Cortex-A7 controller [128]. It serves as a low cost and small form factor Linux platform to interface with off-the-shelf radio chipsets, as illustrated in Figure 5.2.

5.2.2 Record and Playback

Now we describe how to conduct RF record and playback using USRPs.

Interferers. The *Record and Playback* module is not bound to any specific interferer. This module can be used to record and playback RF radio technologies or playback (third party) recorded files or synthesized RF signals. We record RF signals of three interfering technologies operating in the unlicensed bands, namely: *(i)* Microwave oven (Clatronic MWG 758 oven), *(ii)* Analog DSSS cordless phone (Vtech GZ2456 cordless handset system [152]), and *(iii)* Wireless camera (Philips SCD 603 digital video baby monitor [117]). We refer to Table 5.1 for technical details

RF Technology	Vendor & Product Name	TX Power (dBm)	Channel Width (MHz)	Modulation Scheme	Spectrum Range (GHz)
Analog Phone	Vtech GZ2456	n/a	0.1 (Static)	DSSS and BPSK	2.41 - 2.42
Analog Phone	Uniden TRU 4465-2	n/a	0.08 (Static)	DSSS and GFSK	2.40 - 2.48
FHSS Cordless Phone	Uniden DCT6485-3HS	21	0.8 (FH)	GFSK and FHSS	2.41 - 2.47
Wireless Camera	Philips SCD 603	20	1.125 (FH)	BPSK	2.42 - 2.46
Wireless Camera	Genica C-501	20	0.1 (Static)	GFSK	2.41 - 2.47
IEEE 802.15.4	XBee XBP24-AWI-001	4	2 (Static)	DSSS and O-QPSK	2.40 - 2.48
Bluetooth (Class 2)	Bluetooth V2.0 EDR	4	1 (FH)	GFSK	2.40 - 2.48
BLE (Bluetooth V.4.0)	BLED112	4	2 (FH)	GFSK	2.40 - 2.48
Microwave Oven	Clatronic MWG 758	60	-	-	2.44 - 2.48
IEEE 802.11	RTL8192cu Chipset	17	20 (Static)	DSSS, DBPSK	2.40 - 2.48

Table 5.1: Characteristics of the considered RF technologies supported by CIG.

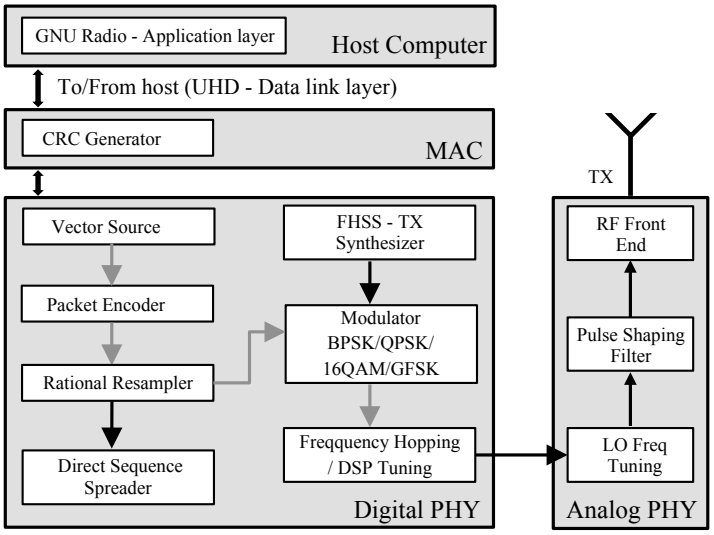

Figure 5.3: Simplified USRP block diagram to signal flow graph mapping. As an example, USRP implementation of the wireless camera is indicated by DSP blocks connected with gray arrows.

about the interferer devices used in this project. We select these particular technologies, representing three typical CTI behaviors, namely: frequency sweeping, frequency static, and high rate frequency hopping, respectively, to analyze the system's record and playback capabilities.

Record. We record 50 million samples by configuring the USRP to tune to the respective device's operational bandwidth and center frequency (f_c), as listed in Table 5.1. We perform the recording in an office environment. However, to maximize the correlation between the recorded and the actual signal, the recording can be performed in an anechoic chamber, which ousts the impact of nearby interfering signals on the recorded signal.

While the center frequency and the bandwidth need to be adjusted according to the wireless radio specifications of the interferer, the receive gain parameter needs to be adjusted according to the peak power of received signal and the SDR hardware specifications (i.e., the supported ADC range). The receive gain influences the accuracy of recorded signal, thus need to be adjusted to attain a unit amplitude of the recorded baseband signal, in order to use the full range of the 14-bits ADC without clipping. This does not necessarily correspond to the highest gain. For instance, recording a high-power microwave oven at 1 m distance, with the maximum gain of SBX (31.5 dB), results into signal clipping. Hence, it is necessary to select the receive gain in such a way that the clipping is avoided. For example for the microwave oven used in this project, the receive gain of 25 dB avoids clipping at 1 m distance.

Playback. The recorded signals are stored as 16 bit I/Q data samples. During playback, the recorded raw baseband data is sent to the USRP, which converts it to analog signal. The analog signal is then transmitted by the USRP by up-converting it to the RF signal. We configure the USRP's data rate (i.e., the rate of reading the recorded file) to match the recording sampling rate. The f_c is set according to the device specifications. The SBX daughter board has a nonlinear gain response when operating in a wide bandwidth [42]. Therefore, it is challenging to regenerate the wide-band recorded signal at the accurate power level, as the down-converted baseband signal does not match the actual transmit power specifications of the device. Hence, during playback, we set the transmit gain value to match the average power level and the peak power to the specified signal power (according to device specifications).

The accuracy of the playback signal is dependent upon hardware limitations of USRP, particularly the sampling rate, maximum transmit power, frequency tuning and settling time, and latency in the hopping rate imposed by the OS scheduling and Ethernet transmission time. We observe that the *Record and Playback* module is more suitable for narrow band interferers occupying static frequency channels, e.g., the DSSS cordless phone [14], provided that adequate device specifications are available to set the recording parameters. It is also suitable for frequency sweeping microwave ovens where the sweep to the next frequency channel typically occurs after 10-15 ms which provides sufficient time for retuning and settling to the next frequency. We observe that our platform can accurately capture the ON and OFF patterns of the microwave oven, over 40 MHz of bandwidth. However, for frequency hopping interferers, such as wireless cameras where the typical hopping rate is 400-600 hops/s, the frequency synthesizer is not able to capture all the packets, switch and settle to the next hop frequency in a bounded time to accurately represent the device specific frequency hopping nature.

5.2.3 Software Radio Implementations

We implement the physical layer (PHY) of five commercially available wireless interferers, operating in the unlicensed bands. We use the GNU Radio [53] framework to build the signal processing blocks and construct flow graphs of the considered radios. In the case of proprietary technologies, we implement the physical layer according to the description in the devices manuals with the support of our spectral analysis of the target device. Figure 5.3 shows the implemented flow graph of the wireless camera, as an example. Additionally, we implement a CRC generator in the sender software radio and a CRC checker in the receiver software radio; this enables the users to collect statistics about the

performance of active transmissions. Hence, researchers can additionally quantify the harm their wireless solutions introduce on other competing devices, such as wireless cameras where so far it is not trivial to quantify this impact. In the following, we elaborate on the implemented software radios we consider in this prototype of CIG:

Analog Cordless Phone. Our analog cordless handset [13] operates in a narrow frequency band [2410.2 - 2418.9] MHz. The user can configure the device to operate on one of the 30 supported channels, each 100 kHz wide. The phone uses DSSS to spread the BPSK modulated data. We use a vector source to generate bit streams followed by the spread spectrum block and connect the output to a BPSK modulator. We set the center frequency of the USRP sink block to match the f_c of the first supported channel (2.417 GHz). The user, however, can change channel configurations through the CIG's host software.

DSSS Cordless Phone. The phone base and the handset [14] communicate using digital spread spectrum and operate in the frequency band [2.407 - 2.478] GHz. The phone supports 28 possible channels, each 3 MHz wide, and shifts the operational channel automatically upon sensing interference. In our implementation, we provide the channel selection option to the user. The phone uses a data rate of 1.366 Mbit/s [54], employs digital spread spectrum, and transmits the data over GFSK modulation. We use the rational resampler block to achieve the specified data rate. The interpolation and decimation values can be derived from Equation 5.1 where the desired bit rate depends on the DAC sampling rate and the number of Samples per Symbol (SPS). We further connect this block to the DSSS block and GFSK modulator.

$$Bit\ Rate = DAC\ Rate/(Interpolation \times SPS) \tag{5.1}$$

Wireless Camera. We include two wireless cameras [117, 51] in our platform. The first is a wireless baby monitor [117]. It communicates with the video receiver using frequency hopping over 61 channels, each of which has a bandwidth of 1.125 MHz and uses BPSK modulation scheme. The second is a wireless monitoring camera [51]. It supports four different channels (2.414 GHz, 2.432 GHz, 2.450 GHz, and 2.468 GHz) and occupies a wide bandwidth of 16 MHz. We perform spectral analysis of these technologies to examine the on-air packet time, hopping sequence, and hopping rate. For the Philips baby monitor, we observe an average packet on-air time of 2.2 ms with a hopping rate of 450 hops/s.

We use respective blocks to generate packets and modulate them as specified in the device specifications. We connect the modulated output to the frequency hopping block.

The USRP N210 has two stages of frequency tuning: *(i)* RF front-end which translates between the RF and the intermediate frequency (IF), and tunes the frequency as close as possible to f_c. *(ii)* DSP, which translates from the IF to the baseband, accounts for the error in frequency tuning, and digitally sets the necessary offset to tune to the desired f_c. To achieve faster-hopping rates in the order of 2 ms tuning time, we fix the RF front-end frequency at the center of the band and hop via shifting in the FPGA only by using timed transactions and tune request objects [45]. We generate the signal at baseband and use the FPGA to convert the signal digitally to the correct frequency. We also schedule the frequency changes and streaming commands a priori to hop faster and deterministically, using timed transactions. We set the channel changes to cover all the channels specified within the operational bandwidth. The time is set to achieve the maximum number of hops possible through our implementation which is 280 hops/s.

FHSS Cordless Phone. The phone base and handset [150] communicate using FHSS, hopping over 90 channels in the range [2.4075 - 2.472] GHz, with a channel width of 800 kHz and GFSK modulation. The discussion we provided on the wireless camera implementation applies here, given that both technologies employ the same underlying signal spreading scheme, i.e., frequency hopping, only with slight changes in channel bandwidth and hopping rate.

5.2.4 Commercial Radio Chipsets

Additionally, we enable interference generation from off-the-shelf radio dongles for a set of prevalent wireless communication standards. We attach to our platform radio chipsets of various technologies, such as IEEE 802.11 (b/g/n) [158], Bluetooth class 2 [19], Bluetooth Low Energy [146], and ZigBee [160]. The transmission power, channel number, and data traffic parameters can be configured by the user via the host software to emulate various application traffic patterns.

5.3 Validation

To quantify to what extent one can rely on CIG for resembling interference in wireless experimentation, we conduct two types of analysis: *(i) Spectral analysis*; we perform a quantitative spectral analysis to validate CIG's accuracy in the time and frequency domains. *(ii) Network impact analysis*; we subject a small 802.15.4 network to interference generated by CIG and latter to interference generated by genuine interferer sources.

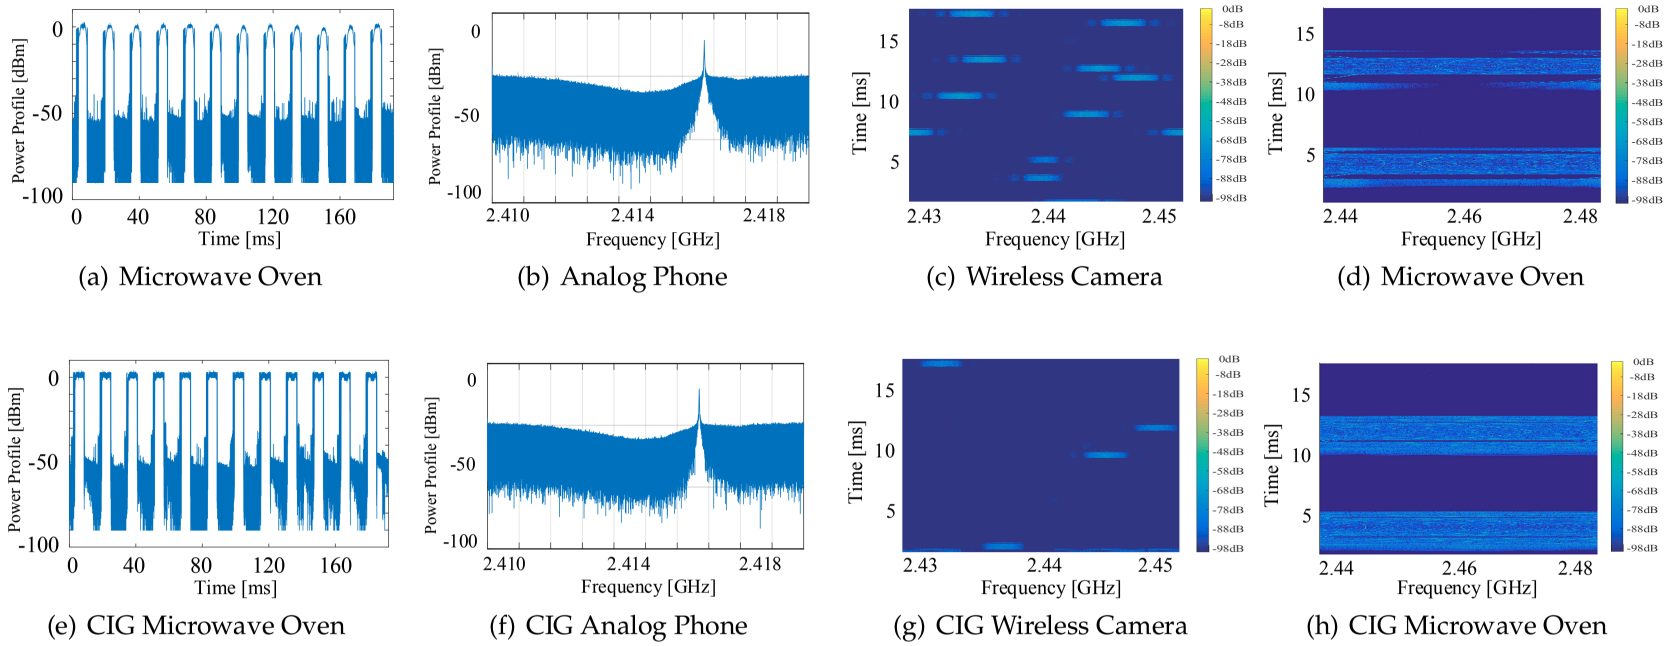

Figure 5.4: Comparison of interference patterns generated by actual interferers in the first row and CIG in the second row. (a) and (e) depict time profiles of the microwave oven. (b) and (f) depict periodograms of the analog phone. (c) and (f) depict spectrograms of the wireless camera. (d) and (h) depict spectrograms of the microwave oven.

We then compare their impact on the network performance. We focus our cross-validation on the SDR modules of CIG and assume that the implementations of the off-the-shelf module to be inherently correct. Our experimental setup consists of two instances of CIG, the discussed interferer technologies and a pair of low-power sensor nodes (TelosB motes).

5.3.1 Temporal Accuracy

We analyze the temporal characteristics of interference generated by CIG and compare it to that of the represented interference source. For this, we record the interference signal from CIG and the corresponding interferer device. Afterward, we compare the pulse duration and number of pulses in a given time period for each technology. For instance, as depicted in 5.4(a) and 5.4(e) for the microwave oven, we observe an equal number of pulses and similar timing behavior. In order to quantify CIG time accuracy, we compute the correlation coefficient between the regenerated signal and the originally recorded signal. For this we represent the original and regenerated signals as binary series $x(i)$ and $y(i)$ with i = $1, \ldots, N$ and N as the number of samples considered in the correlation ($N = [1, 10^6]$). The series take 0 value for clear channel and 1 for a busy channel, given a threshold of -45 dBm (typical clear channel assessment threshold for 802.15.4).

For the microwave oven, where the signal exhibits an ON and OFF pattern, the average cross-correlation coefficient over the length of the samples is 0.926 with a standard deviation of 0.0764. This high accuracy is due to the good performance of the SDR hardware in playing back the recorded samples without a noticeable jitter. In the case of analog DSSS phone, we observe a high cross-correlation value of 0.998. The wireless camera uses frequency hopping. Hence, to validate its temporal behavior, we compare the on-air packet time and the number of packets generated in a given time frame. Figures 5.4(c) and 5.4(g) visualize the general trend. We observe an average cross-correlation coefficient of 0.930 for each packet. However, we reach only 62.2% of the required hopping rate which is due to hardware limitations of the used SDR platform; we have elaborated more on this issue in Section 5.2.

5.3.2 Spectral Accuracy

In order to quantify the spectral accuracy of CIG, we consider aspects representing particular spectral patterns of the considered interferers. That is the static frequency behavior of analog phone where the signal peak lies at the center frequency of the selected channel,

the frequency sweeping behavior of microwave ovens where the sweeping occurs within the second half of the ISM band, and frequency hopping behavior of the wireless camera. We analyze the power spectral density and consider 95% occupied bandwidth for comparison. We compare the center frequency of the signal in case of the analog phone which lies at 2.417 GHz in both cases (see Figure 5.4(b) and 5.4(f)). The occupied bandwidth is 100 kHz for the actual phone and 107 kHz for CIG showing a reasonable accuracy for the analog phone. In the case of microwave oven, we validate the frequency sweeping behavior by comparing the spectrograms of the actual microwave oven and that generated by CIG for microwave oven, depicted in Figure 5.4(d) and 5.4(h). We observe a high energy present on the channel corresponding to microwave ON cycles for both of the cases. The average bandwidth occupied by the ON cycle amounts to approximately 284 kHz for the actual microwave and the generated signal by CIG. For the wireless camera, it is more challenging to compare and validate the channel switching pattern used in frequency hopping due to the absence of a particular sequence. Hence, we only compare the average bandwidth occupied by each packet which is 2.22 MHz for actual camera signal and 2.38 MHz for CIG generated signal.

5.3.3 Impact on the Communication Link

In the following, we study the impact of interference on the performance of an 802.15.4 link subjected to interference generated by original interferers and as compared to that generated by CIG. For the communication link, we use a pair of TelosB nodes. We evaluate various setups, but highlight here the following setup: The transmitter sends 1000 packets, each with a length of 50 bytes and CCA enabled at a transmit power of 0 dBm with an interval of 100 ms to a receiver placed 4 m away.

The transmitter logs CCA status before each transmission. The receiver logs statistics about received packets including RSSI, LQI reading, and the induced power level on the channel. We select the communication channel to overlap with the channel used by the interference source, or the one within the interferer used frequency hopping sequence.

In our experiments, interference generated by CIG exhibits in most cases similar impact on the communication link as interference generated by original interferers. The Packet Reception Rate (PRR) obtained for CIG's microwave oven, is 6.2% lower than the original oven. This is due to USRP's transmit power adjustment during signal playback which results in an increased noise level at the OFF periods of the microwave oven operational cycle. This consequently leads to slightly higher packet losses for receivers at distances affected by the residual noise.

Similarly, we observe a lower Link Quality Indicator (LQI) (indicating bad link quality) and higher noise readings, which only vary within 2 dBm.

In the case of the analog phone, the 802.15.4 transmitter kept backing off thus communication was not possible. This is because the phone continuously emits energy in the medium, thus monopolizing it completely. For both CIG phone and the phone device, we measure similar noise levels and LQI values. While disabling CCA (as we explored in Chapter 2 to allow communication during persistent interference), CIG results into a similar performance as the original device. Hereby, the PRR remains almost the same, showing a reasonable accuracy for analog phone interferer.

For the wireless camera, the PPR is 13.3% higher for CIG generated interference. This is due to the hopping rate limitations and consequently lower packet transmission rate. The average LQI and noise values for both interference sources are, however, in the same range. Moreover, we measure similar average RSSI values (variance of ±2 dBm) during packet reception, in both cases.

5.4 Summary

The number of wireless technologies operating in the unlicensed bands is significantly surging. This phenomenon brings unprecedented challenges for services and applications depending on the wireless medium for communication. A clear understanding of the specifics of radio spectrum sharing is important for the design and verification of wireless systems and protocols. To facilitate wireless systems and protocols testing and verification under heterogeneous interference patterns, we need to augment testbeds and experimental environments with tools that are capable of generating realistic and repeatable interference patterns, and yet easy to access and use.

In this chapter, we introduce CIG, a software-defined radio based controlled interference generator. CIG embodies a set of prevalent radio interferers in one device that can be installed in remote testbeds. CIG design incorporates playback capabilities to regenerate recorded interference patterns, as well as software radio implementation of a set of prevalent interferers operating in the unlicensed band. CIG is easy to use, install, and configure. We validate the spectral and temporal accuracy of the interference generated by CIG. In the design of CIG, we focus on emulating the physical properties of the interfering signals. Although CIG is not flawless, it is a competitive alternative solution in terms of fidelity, usability, extendability, and affordability.

In the design of CIG, we focus on reflecting on one dimension of interference generating, namely, reproducing the interfering physical signal (i.e., radio signal's physical properties) with high fidelity. For CIG to be comprehensive, we need to consider reflecting on other aspects of the interferer, such as location, orientation, and mobility.

For many interference mitigation solutions, reproducing the exact physical signal of the interfering source can be superfluous for testing. This particularly applies to solutions residing in MAC and upper layers (e.g., clear channel assessment, interference avoidance, channel sampling for free channel discovery, and channel occupancy patterns for opportunistic MAC scheduling). For testing this class of solutions, it is sufficient to focus on replicating the temporal and spectral characteristics of the interferers. In the design of CIG, we consider digitally implementing the exact coding and modulation schemes (PHY layer) of the interfering technologies. This allows researchers to test interference solutions that reside in the physical layer in addition to MAC and upper layers. CIG, hence, caters for an automated testing of a larger body of coexistence solutions.

This chapter is dedicated to designing a solution that can help researchers test the dependability of wireless networks deployed in indoor environments rich of radio interference. The numerous unanticipated hours we spent in conducting the experiments presented in this dissertation, led us to work on CIG, believing that researcher's time should be better invested fabricating novel solutions rather than racking with the ill-supported available experimental environments. We hope that the shared insights and design decisions presented in this chapter can help to design better experiment tools for radio coexistence research.

6

Conclusions and Outlook

To date, much of the devised radio frequency interference solutions have been focusing on resolving interference between devices of the same technology. As a consequence, current wireless systems are short of mechanisms to identify and adapt to dynamic sources of external interference. Utilizing non-overlapping segments of the spectrum has been the natural solution to avoid/tackle interference between different technologies. However, as the density of radio devices continue to increase, this solution will no longer suffice. Therefore, it is necessary to pursue alternative avenues to overcome the coexistence challenge in the scarce spectrum. This dissertation advocates for an alternative architecture that builds around mechanisms for wireless coexistence, and in its essence, focuses on designing radios that understand interference better and can reliably operate in occupied channels. The mechanisms and systems presented in this dissertation deliver essential building blocks for wireless networks to improve throughput and reliability in interference rich environments.

We now conclude this dissertation with a brief summary of our contributions and a discussion of directions for future work.

6.1 Contributions

In this dissertation, we argued that there are sufficient unutilized opportunities for low-power systems to coexist in shared channels. Hence, to alleviate the spectrum scarcity consequences, we need to revise current wireless designs to leverage these opportunities.

The systems presented in this dissertation adopt machine learning techniques and a cross-layer approach to increase radio's cognition of their environments. In particular, they exploit richer physical layer information and devise algorithms that automatically identify, calibrate, and correct for variations due to Cross-Technology Interference in the channel. To support our argument, this dissertation introduced three new systems: TIIM, CrossZig, and CIG that contribute to wireless systems coexistence.

TIIM. In this work, we have presented the problem of cross-technology interference, focusing on its implications on low-power wireless networks. We have shown that cross-technology interference has a non-negligible impact on the performance and availability of low-power wireless networks. Due to the heterogeneity and dynamicity of CTI, there is no one-size-fits-all solution to combat CTI. Correspondingly, we argued and showed that understanding the type and nature of interference is crucial for deciding how to mitigate it best. To seamlessly realize this, we developed a lightweight classifier that is trained to recognize channel fingerprints at which a particular coexistence solution can work best. In Chapter 3, we presented the design of TIIM, a lightweight *Technology-Independent Interference Mitigation* system that identifies, quantifies, and reacts to CTI in real-time. In the design of TIIM, we followed an unorthodox approach, where we employed machine learning to assist wireless nodes in recovering from interference. TIIM employs a lightweight machine learning classifier to *(i)* decide whether the communication is viable over the interfered link and *(ii)* characterize the ambient conditions and dynamically apply the best coexistence mitigation strategy. We developed a prototype of TIIM based on an off-the-shelf 802.15.4 radio platform. Our evaluation showed that TIIM, while exposed to extensive and heterogeneous interference, can achieve a total packet reception rate gain of 30% with an additional transmission overhead of 5.6%.

CrossZig. Current wireless designs still largely impose layer isolation, where the lower layers deliver fully correct packets for upper layers. Given this, conventional approaches to tackle wireless performance has focused on separately optimizing different layers of the networking stack. This rigid design fails to harness correctly received bits within corrupted packets and is oblivious of the rich ambient information embedded in the physical signals. Hence, reliability solutions developed for this design model are typically suboptimal. In recent years, cross-layer optimizations were profoundly advocated in the wireless community. This has been coupled with rapid developments in software defined radios that made it possible to demonstrate the potentials of cross-layer designs. In this dissertation, we showed how physical layer information and primitives

can be coupled with link layer to enhance low-power wireless systems coexistence and performance under interference. Notably, we showed that passing fine-grained physical layer information to upper layers enables the link layer to make more informed and intelligent decisions when reacting to interference. In Chapter 4, we presented CrossZig, a cross-layer wireless system design, that enables low-power wireless networks to exploit fine-grained information from the physical layer to make informed and intelligent decisions that can help them recover from varying sources of interference. CrossZig utilizes physical layer information to detect the presence of CTI in a corrupted packet and to apply an adaptive packet recovery which incorporates a novel cross-layer based packet merging scheme and an adaptive channel coding. We implemented a prototype of CrossZig for the low-power IEEE 802.15.4 in a software-defined radio platform. We showed the adaptability and the performance gain of CrossZig through experimental evaluation of the system performance under various interference patterns.

CIG. Wireless research testbed infrastructures often lack proper tools for enabling repeatable playback of realistic radio interference commonly found in real-world deployments. This can make it harder for researchers to benchmark their wireless coexistence solutions in remote testbeds. To tackle this challenge, we developed CIG, a tool that can extend current testbed infrastructures with capabilities to (re)run experiments under identical interference patterns. In Chapter 5, we presented CIG, a Controlled Interference Generator (CIG) framework that facilitates wireless coexistence research experimentation. In the design of CIG, we adopted a unified approach that incorporates a careful selection of interferer technologies (implemented in software), to expose networks to realistic interference patterns. We validated the resemblance of interference generated by CIG and interference from represented radio devices, by showing its accuracy in temporal and spectral domains.

This dissertation builds the above ideas into practical systems, integrates them within the IEEE 802.15.4 protocol stack, and provides prototype implementations of the proposed designs. Further, we evaluated them in wireless testbeds, demonstrating large gains in throughput and reliability in practice. This work demonstrate the benefits of alternative wireless network designs that can better coexist in shared channels.

6.2 Remaining Challenges & Future Directions

The systems in this dissertation addressed the challenges involved in low-power wireless systems coexistence in dense and diverse spectral

environments. The building blocks presented in this dissertation provide sound means necessary for better coexistence in the crowded spectrum. Nevertheless, the contributions described in this dissertation do not represent a universal solution to the uncoordinated wireless coexistence challenge. In this section, we first discuss possible directions for future research based on the work presented in this dissertation and then give a broader view on future work in low-power wireless systems coexistence.

Coexistence (TIIM and CrossZig). The work on coexistence presented in this dissertation can be extended to act on the following aspects.

• Low-power Wireless Technologies: For the contributions presented in this dissertation, we discussed, analyzed, and provided prototypes that comply with the IEEE 802.15.4 standard. However, most of the observations can be projected to analogous wireless technologies, such as 802.11 radios. We believe that TIIM's and CrossZig's building blocks can be beneficial to the coexistence of comparable radio technologies, provided that the underlying mechanisms leveraged by our systems are ported and adjusted in accordance with the physical layer details of the target wireless technology. We leave further investigations for future work.

• Highly Mobile and Multi-hop Wireless Networks: Some low-power wireless system's applications can be highly mobile in nature. The operation of the adaptive schemes presented in this dissertation is not optimized for networks that rapidly change as a result of mobility. Variations in such settings can occur at a pace higher than what our systems can cope with, hence, affect the system's stability inversely. Furthermore, TIIM and CrossZig are link-based solutions, ultimately directed at enhancing the link performance under interference. Therefore most of the experiments conducted in this dissertation target one-hop communication links. However, TIIM and CrossZig can be optimized for multi-hop networks. One possible direction is to allow propagation of the selected mitigation along consecutive hops and only trigger the search for a new mitigation scheme if the currently chosen scheme is not effective or begins to endure high cost. This can help minimizing the delay induced by TIIM and CrossZig across the network. We leave further investigations on this direction for future work.

• Other Uses of the Physical Information: Conventional Low Power Listening (LPL) mechanisms are susceptible to interference, this is mainly because they employ CCA to check the channel state. Current implementations of CCA rely on energy level to detect activities in the channel. Energy detection fails to differentiate between channel

activities arising from interference and those arising from target signals; this can consequently extrapolate the false wakeup problem. This has a considerable impact on Radio Duty Cycle's (RDC) performance and, hence, energy efficiency. Potential future work can focus on enhanced cross-layer MAC protocol designs, by optimizing the underlying CCA mechanism. Hereby, spectrum sensing techniques that are robust to CTI should be alternatively considered. For instance, combining the energy detection- with correlation-based detectors, which aims at detecting signals that have the modulation and spreading characteristics of the target signal. Moreover, soft values as we have examined in this dissertation can be a good indicator of the signal type, and can be an efficient alternative to energy detection. These approaches are however more power hungry than energy detection in general, but can deliver better overall energy efficiency for environments rich of interference. Moreover, scheduling algorithms used in multichannel MAC protocols such as TSCH can benefit further from a cross-layer design, where scheduling algorithms can adapt better to congestion. We leave further investigations of cross-layer optimizations for LPL and TSCH for future work.

• Resource Asymmetry in Low-power IoT Networks: In this work, we assume that resources are symmetric in the transmitter and receiver. Introduced solutions were designed to consider the constrained nature of the devices symmetrically. This is typically the case for networks with homogenous hardware resources (i.e., WSN). However, for IoT applications, hardware resources are commonly asymmetric, such that sensors are embedded in appliances, wearables, and within bodies and transmit their data to an unconstraint device referred to as the gateway. For combating CTI in such settings, one can take advantage of this inherent asymmetry. This asymmetry can be leveraged to enhance wireless coexistence, by porting the computationally expensive state of the art solutions to the unconstrained device to manage spectrum usage and assist constrained devices in deriving the adequate countermeasure.

• Spectrum Utilization: In this dissertation, we focused on enhancing spectrum utilization by employing mechanisms that harness opportunities that arise from variations in spectral occupancy across time. Similar opportunities can arise due to variations in spectral occupancy and channel conditions across *space and frequency*. For opportunities arising from space diversity, directionality offered by antenna beam steering can be harnessed to avoid interfering with a co-located interferer that simultaneously utilizes the same channel. Analogously, in multi-path networks, nodes can route packets to alternative paths that are less

affected by interference (i.e., exploit spatial diversity). Moreover, our systems have a narrow view of the radio spectrum that is limited to the 802.15.4 channel width. They focus on increasing the spectral efficiency over interfered channels, unacquainted of the state of the rest of the spectrum. Thus, our system lacks a comprehensive view of the radio spectrum to decide whether communication over an interfered channel is preferred over channel switching. To overcome the limitation of narrow spectrum perception in narrow-band channels, solutions such as cyclostationary analysis for bandwidth estimation of interfering signals, agile radios, and adaptive spectrum access can be employed. However, these approaches are not yet viable for constrained wireless systems. One possible solution is to outsource this task to a central entity that can be queried about the spectrum status, or if communication occurs between devices with asymmetric resources the more capable device takes this responsibility.

• Fairness and Coexistence: With the existing power asymmetry in the shared unlicensed bands, where low-power radios typically transmit at 0 dBm (several times lower than other radios), there is a large region where low-power transmitters can sense transmission from a high power interferer but not vice versa. Hence, low-power nodes unnecessarily abstain from transmitting and thus, suffer from starvation. To contain this effect, there is a need to rethink how the current CTI-oblivious CSMA protocol works (i.e., trigger devices to abstain transmissions without a good assessment of the potential harm a transmission can cause). This particularly applies for links subject to CTI, where using the default CCA can cause high false negatives. Access modalities that can allow opportunistic access to the interfered channel and, hence, allow utilization of frequency holes due to sparse frequency access and time variant characteristics of interferers can lead to a better utilization of the scarce spectrum. CTI-aware recovery mechanisms such as the one introduced in this dissertation can then help to alleviate the potential CTI damages.

Experimentation (CIG). In the design of CIG, much of our focus has been on emulating the physical properties of the interfering signals. We focus on reflecting on one dimension of interference generation, namely, reproducing the interfering physical signal (i.e., radio signal's physical properties) with high fidelity. For CIG to be comprehensive, we need however to consider reflecting on other aspects of the interferer source, such as realistically emulating location, orientation, and mobility patterns of the target interferers in future prototypes. Moreover, the current CIG design allows users to configure application traffic patterns and content manually. However, this requires the user to be knowledgeable about these aspects of the interference which can be cumbersome.

Therefore, providing support of realistic traffic pattern models is a feature we plan to address in future prototypes of CIG. Finally, integrating CIG in a public testbed is the natural next step for testing CIG (e.g., Flocklab [97], Indriya [37]).

Further Directions. With the number of wireless devices surging in the unlicensed bands, the coexistence challenge will become more pressing. Efforts have already started with exploring alternative means beyond utilizing the increasingly crowded unlicensed radio spectrum. For instance, resorting to visible light communication instead of radio waves, leveraging white spaces, or alternative licensed bands are some promising alternatives for low-power wireless systems. Pursuing the latter direction requires the development of efficient mechanisms for spectrum sensing, in order to increase the chances for low-power systems to coexist in these deserted bands. Furthermore, in this dissertation we advocate for bringing miniaturized intelligence into radios to address the complexity of the CTI problem. The concept of bringing intelligence to radios is not new. Communication concepts, such as cognitive radio, promise integration of intelligence into radios, such that they can sense, learn from, and adapt to their environment. To date, most of cognitive radio research focuses on licensed bands and has been restricted to policy-based radios that are hard-coded with rules on how to react in certain scenarios. Devising radios that utilize machine learning techniques, i.e., a learning-based cognitive radio, is a relatively unchartered research area. This dissertation shed light on the potential of using machine learning in one wireless communication application, namely automating countermeasure selection in the presence of interference. Interference in the unlicensed band is primarily from communication systems that follow systematic protocols which can be learned and exploited for better coexistence. However, the scope of machine learning applications in wireless communication is wide and needs further research exploration. The need of learning components in radios is more evident now, as the rise of active wireless devices implies that more RF optimizations and tuning is needed.

As the density and diversity of wireless devices populating the unlicensed spectrum continue to increase, there is an urgent need to rethink current wireless designs. Designing agile wireless protocols that can exploit opportunities that arise from variations in spectral occupancy across time, space, and frequency is inevitable for better utilization of the scarce spectrum. This dissertation demonstrates practical wireless systems that realize this by incorporating machine learning techniques and cross-layer algorithms to enhance low-power wireless systems coexistence in occupied channels.

Bibliography

[1] Extending LTE Advanced to Unlicensed Spectrum. White paper, Qualcomm Incorporated, 2013.

[2] LTE for Unlicensed Spectrum. White paper, Nokia Corporation, 2014.

[3] LTE in Unlicensed Spectrum: Harmonious Coexistence with Wi-Fi. White paper, Qualcomm Incorporated, 2015.

[4] Fadel Adib, Zachary Kabelac, and Dina Katabi. Multi-Person Localization via RF Body Reflections. In *Proceedings of the 12th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, pages 279–292, 2015.

[5] Fadel Adib, Hongzi Mao, Zachary Kabelac, Dina Katabi, and Robert C. Miller. Smart Homes That Monitor Breathing and Heart Rate. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI)*, pages 837–846, 2015.

[6] Cedric Adjih, Elmnanuel Baccelli, Eric Fleury, Gaetan Harter, Nathalie Mitton, Thomas Noel, Roger Pissard-Gibollet, Frederic Saint-Marcel, Guillaume Schreiner, Julien Vandaele, and Thomas Watteyne. FIT IoT-LAB: A Large Scale Open Experimental IoT Testbed. In *IEEE 2nd World Forum on Internet of Things (WF-IoT)*, pages 459–464, 2015.

[7] Adrian Wagstaff (Mass Consultants Limited). Estimating the Utilization of Key License-Exempt Spectrum Bands. Report, Online: `http://stakeholders.ofcom.org.uk/binaries/research/technology-research/wfiutilisation.pdf`, 2009.

[8] Ian F. Akyildiz, Won-Yeol Lee, Mehmet C. Vuran, and Shantidev Mohanty. NeXt Generation/Dynamic Spectrum Access/Cognitive Radio Wireless Networks: A Survey. *The International Journal of Computer and Telecommunications Networking*, 50:2127–2159, 2006.

[9] Beshr Al Nahas, Simon Duquennoy, Venkatraman Iyer, and Thiemo Voigt. Low-Power Listening Goes Multi-channel. In *Proceedings of the 10th IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS)*, pages 2–9, 2014.

[10] Urs Anliker, Jamie A. Ward, Paul Lukowicz, Gerhard Tröster, François Dolveck, Michel Baer, Fatou Keita, Eran B. Schenker, Fabrizio Catarsi,

Luca Coluccini, Andrea Belardinelli, Dror Shklarski, Menachem Alon, Etienne Hirt, Rolf Schmid, and Milica Vuskovic. AMON: A Wearable Multiparameter Medical Monitoring and Alert System. *IEEE Transactions on Information Technology in Biomedicine*, pages 415–427, 2004.

[11] Anthony Cuthbertson, ITproPortal. Report: UK faces Wi-Fi and mobile "spectrum crunch" by 2020, 2013.

[12] Ehsan Aryafar, Narendra Anand, Theodoros Salonidis, and Edward W. Knightly. Design and Experimental Evaluation of Multi-user Beamforming in Wireless LANs. In *Proceedings of the 16th ACM Conference on Mobile Computing and Networking (MobiCom)*, pages 197–208, 2010.

[13] Atmel. AT02845: Coexistence between ZigBee and Other 2.4GHz Products Description. Application Note, Online: http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.401.5086.

[14] AT&T. 2.4 GHz Cordless Telephone System E2725B. User Manual, Online: http://cdn-media-att.vtp-media.com/ecp/documents/product_Product/391/UserManual/2173/e2725b_manual_bkm.pdf.

[15] Bandspeed. Understanding the Effects of Radio Frequency (RF) Interference on WLAN performance and Security. White paper, 2010.

[16] Jan Beutel, Stephan Gruber, Andreas Hasler, Roman Lim, Andreas Meier, Christian Plessl, Igor Talzi, Lothar Thiele, Christian Tschudin, Matthias Woehrle, and Mustafa Yuecel. PermaDAQ: A Scientific Instrument for Precision Sensing and Data Recovery in Environmental Extremes. In *Proceedings of the 8th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 265–276, 2009.

[17] Ezio Biglieri, Andrea Goldsmith, Larry J. Greenstein, Narayan B. Mandayam, and H. Vincent Poor. *Principles of Cognitive Radio*. Cambridge University Press, 2012.

[18] Sanjit Biswas, John Bicket, Edmund Wong, Raluca Musaloiu-E, Apurv Bhartia, and Dan Aguayo. Large-scale Measurements of Wireless Network Behavior. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*, pages 153–165, 2015.

[19] Bluetooth Dongle. Online: https://www.hama.com/00092498/hama-nano-bluetooth-usb-adapter-version-40.

[20] Carlo Alberto Boano, Thiemo Voigt, Claro Noda, Kay Römer, and Marco Zúniga. JamLab: Augmenting Sensornet Testbeds with Realistic and Controlled Interference Generation. In *Proceedings of 10th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 175–186, 2011.

[21] Carlo Alberto Boano, Thiemo Voigt, Nicolas Tsiftes, Luca Mottola, Kay Römer, and Marco Antonio Zúñiga. Making Sensornet MAC Protocols Robust against Interference. In *Proceedings of the 7th European Conference on Wireless Sensor Network (EWSN)*, pages 272–288, 2010.

[22] Donald G. Brennan. Linear Diversity Combining Techniques. *Proceedings of the Institute of Radio Engineers*, pages 1075–1102, 1959.

[23] Vladimir Brik, Eric Rozner, Suman Banerjee, and Paramvir Bahl. DSAP: A Protocol for Coordinated Spectrum Access. In *Proceedings of the First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, pages 611–614, 2005.

[24] Milind M. Buddhikot, Scott Miller, Kevin Ryan, Paul Kolodzy, and Jason Evans. DIMSUMNet: New Directions in Wireless Networking Using Coordinated Dynamic Spectrum Access. In *Proceedings of the 14th IEEE Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, pages 78–85, 2005.

[25] Maxim Buevich, Dan Schnitzer, Tristan Escalada, Arthur Jacquiau-Chamski, and Anthony Rowe. Fine-grained Remote Monitoring, Control and Pre-paid Electrical Service in Rural Microgrids. In *Proceedings of the 13th ACM International Conference on Information Processing in Sensor Networks (IPSN)*, pages 1–11, 2014.

[26] Danijela Cabric, Shridhar Mubaraq Mishra, Daniel Willkomm, Robert Brodersen, and Adam Wolisz. A Cognitive radio approach for usage of virtual unlicensed spectrum. In *Proceedings of the 14th IST Mobile Wireless Communications Summit*, pages 1–4, 2005.

[27] Viveck R Cadambe and Syed Ali Jafar. Interference Alignment and Degrees of Freedom of the K-User Interference Channel. *IEEE Transactions on Information Theory*, pages 3425–3441, 2008.

[28] Matteo Ceriotti, Luca Mottola, Gian Pietro Picco, Amy L. Murphy, Ştefan Gună, Michele Corrà, Matteo Pozzi, Daniele Zonta, and Paolo Zanon. Monitoring heritage buildings with wireless sensor networks: The Torre Aquila deployment. In *Proceedings of the 8th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, 2009.

[29] Ranveer Chandra, Ratul Mahajan, Thomas Moscibroda, Ramya Raghavendra, and Paramvir Bahl. A Case for Adapting Channel Width in Wireless Networks. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*, pages 135–146, 2008.

[30] Octav Chipara, Chengjie Wu, Chenyang Lu, and William Griswold. Interference-aware real-time flow scheduling for wireless sensor networks. In *Euromicro Conference Real-Time Systems (ECRTS)*, pages 67–77, 2011.

[31] Kaushik R. Chowdhury and Ian F. Akyildiz. Interferer Classification, Channel Selection and Transmission Adaptation for Wireless Sensor Networks. In *Proceedings of the IEEE International Conference on Communications (ICC)*, pages 297–301, 2009.

[32] Brent Chun, David Culler, Timothy Roscoe, Andy Bavier, Larry Peterson, Mike Wawrzoniak, and Mic Bowman. PlanetLab: An Overlay Testbed for Broad-coverage Services. *ACM SIGCOMM Computer Communication Review*, pages 3–12, 2003.

[33] Cisco. 20 Myths of Wi-Fi Interference: Dispel Myths to Gain HighPerforming and Reliable Wireless. White paper, Online: http://www.cisco.com/c/en/us/products/collateral/wireless/spectrum-expert-wi-fi/prod_white_paper0900aecd807395a9.pdf, 2007.

[34] Cisco CleanAir. Online http://www.cisco.com/en/US/netsol/ns1070/.

[35] Antonio Criminisi and Jamie Shotton. *Decision Forests for Computer Vision and Medical Image Analysis*. Springer Publishing Company, 2013.

[36] Dariush Divsalar, Marvin K. Simon, and Dan Raphaeli. Improved parallel interference cancellation for CDMA. *IEEE Transactions on Communications*, pages 258–268, 1998.

[37] Manjunath Doddavenkatappa, Mun Choon Chan, and Akkihebbal L Ananda. Indriya: A Low-Cost, 3D Wireless Sensor Network Testbed. In *Proceedings of the Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TridentCom)*, pages 302–316, 2012.

[38] Richard O. Duda, Peter E Hart, and David G. Stork. *Pattern Classification*. John Wiley & Sons, Inc., New York, NY, 2012.

[39] Adam Dunkels, Bjorn Gronvall, and Thiemo Voigt. Contiki - A Lightweight and Flexible Operating System for Tiny Networked Sensors. In *Proceedings of the 29th Annual IEEE International Conference on Local Computer Networks (LCN)*, pages 455–462, 2004.

[40] Simon Duquennoy, Beshr Al Nahas, Olaf Landsiedel, and Thomas Watteyne. Orchestra: Robust Mesh Networks Through Autonomously Scheduled TSCH. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems (SenSys)*, pages 337–350, 2015.

[41] Joakim Eriksson, Fredrik Österlind, Niclas Finne, Nicolas Tsiftes, Adam Dunkels, Thiemo Voigt, Robert Sauter, and Pedro José Marrón. COOJA/MSPSim: Interoperability testing for wireless sensor networks. In *Proceedings of the 2nd ICST International Conference on Simulation Tools and Techniques (SIMUTools)*, 2009.

[42] Ettus Research, A National Instruments Company. SBX Nonlinearity. Online: http://files.ettus.com/performance_data/sbx/SBX-without-UHD-corrections.pdf.

[43] Ettus Research, A National Instruments Company. SBX USRP Daughterboard (400 MHz - 4.4 GHz). Online: https://www.ettus.com/product/details/SBX.

[44] Ettus Research, A National Instruments Company. Universal Software Radio Peripheral. Online: https://www.ettus.com/.

[45] Ettus Research, A National Instruments Company. USRP Tuning Process. Online: http://files.ettus.com/manual/page_general.html.

[46] Farpoint Group. Online: http://www.farpointgroup.com/.

[47] FCC Staff Technical Paper. Mobile Broadband: The Benefits of Additional Spectrum, 2010.

[48] George S. Fishman. *Principles of Discrete Event Simulation*. John Wiley & Sons, Inc., New York, NY, 1978.

[49] Hossein Fotouhi, Mário Alves, Marco Zuniga, Nouha Baccour, Claro Noda, Thiemo Voigt, Kay Romer, and Carlo Boano. *Radio Link Quality Estimation in Low-Power Wireless Networks*. Springer International Publishing, Heidelberg, July 2013.

[50] William A. Gardner. Exploitation of spectral redundancy in cyclostationary signals. *IEEE Signal Processing Magazine*, 8(2):14–36, 1991.

[51] Genica. 2.4 GHz 4-Channel Wireless Receiver and 4 Wireless Infrared Color Cameras. Online: http://www.genica.com.

[52] David Gesbert, Stephen Hanly, Howard Huang, Shlomo Shamai Shitz, Osvaldo Simeone, and Wei Yu. Multi-Cell MIMO Cooperative Networks: A New Look at Interference. *IEEE Journal on Selected Areas in Communications*, 28(9):1380–1408, 2010.

[53] GNU Radio Website. Online: http://www.gnuradio.org.

[54] Shyamnath Gollakota, Fadel Adib, Dina Katabi, and Srinivasan Seshan. Clearing the RF Smog: Making 802.11n Robust to Cross-technology Interference. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*, pages 170–181, 2011.

[55] Shyamnath Gollakota, Haitham Hassanieh, Benjamin Ransford, Dina Katabi, and Kevin Fu. They Can Hear Your Heartbeats: Non-invasive Security for Implantable Medical Devices. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*, pages 2–13, 2011.

[56] Shyamnath Gollakota and Dina Katabi. Zigzag Decoding: Combating Hidden Terminals in Wireless Networks. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*, pages 159–170, 2008.

[57] Krishna Gomadam, Viveck R Cadambe, and Syed A Jafar. A Distributed Numerical Approach to Interference Alignment and Applications to Wireless Interference Networks. *IEEE Transactions on Information Theory*, 57(6):3309–3322, 2011.

[58] Antonio Gonga, Olaf Landsiedel, Pablo Soldati, and Mikael Johansson. 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)*, pages 186–193, 2012.

[59] Google Spectrum Database. Online: `https://www.google.com/get/spectrumdatabase/`.

[60] Ramakrishna Gummadi, David Wetherall, Ben Greenstein, and Srinivasan Seshan. Understanding and Mitigating the Impact of RF Interference on 802.11 Networks. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*, pages 385–396, 2007.

[61] Peng Guo, Jiannong Cao, Kui Zhang, and Xuefeng Liu. Enhancing ZigBee throughput under WiFi interference using real-time adaptive coding. In *Proceedings of the 33rd IEEE International Conference on Computer Communications (INFOCOM)*, pages 2858–2866, 2014.

[62] Daniel Halperin, Thomas Anderson, and David Wetherall. Taking the Sting out of Carrier Sense: Interference Cancellation for Wireless LANs. In *Proceedings of the 14th ACM International Conference on Mobile Computing and Networking (MobiCom)*, pages 339–350, 2008.

[63] Bo Han, Aaron Schulman, Francesco Gringoli, Neil Spring, Bobby Bhattacharjee, Lorenzo Nava, Lusheng Ji, Seungjoon Lee, and Robert R Miller. Maranello: Practical Partial Packet Recovery for 802.11. In *Proceedings of the 7th USENIX Conference on Networked Systems Design and Implementation (NSDI)*, pages 205–218, 2010.

[64] Vlado Handziski, Andreas Köpke, Andreas Willig, and Adam Wolisz. TWIST: A Scalable and Reconfigurable Testbed for Wireless Indoor Experiments with Sensor Networks. In *Proceedings of the 2nd ACM International Workshop on Multi-hop Ad Hoc Networks: From Theory to Reality (REALMAN)*, pages 63–70, 2006.

[65] Rahul Hariharan. *Enhancing Spectrum Utilization through Cooperation and Cognition in Wireless Systems*. PhD thesis, Massachusetts Institute of Technology (MIT), 2013.

[66] Jan-Hinrich Hauer, Andreas Willig, and Adam Wolisz. Mitigating the Effects of RF Interference Through RSSI-Based Error Recovery. In *Proceedings of the 7th European Conference on Wireless Sensor Networks (EWSN)*, pages 224–239, 2010.

[67] Simon Haykin. Cognitive Radio: Brain-empowered Wireless Communications. *IEEE Journal on Selected Areas in Communications*, 23:201–220, 2006.

[68] Frederik Hermans, Olof Rensfelt, Thiemo Voigt, Edith Ngai, Lars-Åke Norden, and Per Gunningberg. 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)*, pages 55–66, 2013.

[69] Frederik Hermans, Hjalmar Wennerström, Liam Mcnamara, Christian Rohner, and Per Gunningberg. All Is Not Lost: Understanding and Exploiting Packet Corruption in Outdoor Sensor Networks. In *Proceedings of the 11th European Conference on Wireless Sensor Networks (EWSN)*, pages 116–132, 2014.

[70] Mike Hibler, Robert Ricci, Leigh Stoller, Jonathon Duerig, Shashi Guruprasad, Tim Stack, Kirk Webb, and Jay Lepreau. Large-scale Virtualization in the Emulab Network Testbed. In *USENIX Annual Technical Conference (ATC)*, pages 113–128, 2008.

[71] Anwar Hithnawi. Poster Abstract: Exploiting Physical Layer Information to Mitigate Cross-Technology Interference Effects on Low-Power Wireless Networks. In *ACM Conference on Embedded Networked Sensor Systems SenSys (SenSys)*, pages 23:1–23:2, 2013.

[72] Anwar Hithnawi, Vaibhav Kulkarni, Su Li, and Hossein Shafagh. Controlled Interference Generation for Wireless Coexistence Research. In *Proceedings of the ACM workshop on Software Radio Implementation Forum (SRIF)*, pages 19–24, 2015.

[73] Anwar Hithnawi, Su Li, Hossein Shafagh, Simon Duquennoy, and James Gross. Poster: Cross-Layer Optimization for Low-power Wireless Coexistence. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems (SenSys)*, pages 443–444, 2015.

[74] Anwar Hithnawi, Su Li, Hossein Shafagh, James Gross, and Simon Duquennoy. CrossZig: Combating Cross-Technology Interference in Low-Power Wireless Networks. In *Proceedings of the 15th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 1–12, 2016.

[75] Anwar Hithnawi, Hossein Shafagh, and Simon Duquennoy. Poster Abstract: Low-Power Wireless Channel Quality Estimation in the Presence

of RF Smog. In *Proceedings of the 10th IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS)*, pages 137–138, 2014.

[76] Anwar Hithnawi, Hossein Shafagh, and Simon Duquennoy. Understanding the Impact of Cross Technology Interference on IEEE 802.15.4. In *Proceedings of the 9th ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation and Characterization (WiNTECH)*, pages 49–56, 2014.

[77] Anwar Hithnawi, Hossein Shafagh, and Simon Duquennoy. TIIM: Technology-Independent Interference Mitigation for Low-power Wireless Networks. In *Proceedings of the 14th ACM International Conference on Information Processing in Sensor Networks (IPSN)*, pages 1–12, 2015.

[78] Steven Hong and Sachin Katti. DOF: A Local Wireless Information Plane. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*, pages 230–241, 2011.

[79] Steven Hong, Jeffrey Mehlman, and Sachin Katti. Picasso: Flexible RF and Spectrum Slicing. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*, pages 37–48, 2012.

[80] Jun Huang, Guoliang Xing, Gang Zhou, and Ruogu Zhou. Beyond Co-existence: Exploiting WiFi White Space for ZigBee Performance Assurance. In *Proceedings of the 18th IEEE International Conference on Network Protocols (ICNP)*, pages 135–146, 2010.

[81] IEEE Standards Association. IEEE 802.15.4 Standard. Online: `http://standards.ieee.org/findstds/standard/802.15.4-2011.html`, 2011.

[82] Iperf. A TCP, UDP, and SCTP Network Bandwidth Measurement Tool. Online: `https://github.com/esnet/iperf`.

[83] Venkat Iyer, Matthias Woehrle, and Koen Langendoen. Chrysso - 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)*, pages 449–457, 2011.

[84] Venkatraman Iyer, Matthias Woehrle, and Koen Langendoen. Chamaeleon: Exploiting Multiple Channels to Mitigate Interference. In *Proceedings of the 7th International Conference on Networked Sensing Systems (INSS)*, pages 65–68, 2010.

[85] Kyle Jamieson and Hari Balakrishnan. PPR: Partial Packet Recovery for Wireless Networks. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*, pages 409–420, 2007.

[86] Joseph M. Kahn, Randy Katz, and Kristofer S. J. Pister. Next Century Challenges: Mobile Networking for "Smart Dust". In *Proceedings of the*

5th ACM Conference on Mobile Computing and Networking (MobiCom), pages 271–278, 1999.

[87] Song Min Kim and Tian He. FreeBee: Cross-Technology Communication via Free Side-Channel. In *Proceedings of the 21st ACM International Conference on Mobile Computing and Networking (MobiCom)*, pages 317–330, 2015.

[88] Song Min Kim, Shuai Wang, and Tian He. cETX: Incorporating Spatiotemporal Correlation for Better Wireless Networking. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems (SenSys)*, pages 323–336, 2015.

[89] JeongGil Ko, Chenyang Lu, Mani B Srivastava, John A Stankovic, Andreas Terzis, and Matt Welsh. Wireless Sensor Networks for Healthcare. *Proceedings of the IEEE*, pages 1947–1960, 2010.

[90] Shankar Kumar and Chockalingam Ananthanarayanan. Parallel Interference Cancellation in Multi-carrier DS-CDMA Systems. In *Proceedings of the IEEE International Conference on Communications*, pages 2874–2878, 2004.

[91] Jamieson Kyle. *The SoftPHY Abstraction: from Packets to Symbols in Wireless Network Design*. PhD thesis, Massachusetts Institute of Technology (MIT), 2008.

[92] FCC Lab. Report On Trends in Wireless Devices. Online: `http://www.fcc.gov/oet/info/documents/reports/wirelessdevices.doc`, January 2011.

[93] Kaushik Lakshminarayanan, Samir Sapra, Srinivasan Seshan, and Peter Steenkiste. RFDump: An Architecture for Monitoring the Wireless Ether. In *Proceedings of the 5th International Conference on Emerging Networking Experiments and Technologies (CoNEXT)*, pages 253–264, 2009.

[94] Yoonmyung Lee, Suyoung Bang, Inhee Lee, Yejoong Kim, Gyouho Kim, Mohammad Ghaed, Pat Pannuto, Prabal Dutta, Dennis Sylvester, and David Blaauw. A Modular 1 mm^3 Die-Stacked Sensing Platform with Low Power I^2 C Inter-Die Communication and Multi-Modal Energy Harvesting. *IEEE Journal of Solid-State Circuits*, pages 229–243, 2013.

[95] Philip Levis, Nelson Lee, Matt Welsh, and David Culler. TOSSIM: Accurate and Scalable Simulation of Entire TinyOS Applications. In *Proceedings of the 1st ACM International Conference on Embedded Networked Sensor Systems (SenSys)*, pages 126–137, 2003.

[96] Chieh-Jan Mike Liang, Nissanka Bodhi Priyantha, Jie Liu, and Andreas Terzis. Surviving Wi-Fi interference in low power ZigBee networks. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems (SenSys)*, pages 309–322, 2010.

[97] Roman Lim, Federico Ferrari, Marco Zimmerling, Christoph Walser, Philipp Sommer, and Jan Beutel. FlockLab: A Testbed for Distributed, Synchronized Tracing and Profiling of Wireless Embedded Systems. In *Proceedings of 12th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 153–165, 2013.

[98] Roman Lim, Marco Zimmerling, and Lothar Thiele. Passive, Privacy-preserving Real-time Counting of Unmodified Smartphones via ZigBee Interference. In *Proceedings of the 11th IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS)*, pages 115–126, 2015.

[99] Kate Ching-Ju Lin, Nate Kushman, and Dina Katabi. ZipTx: Harnessing Partial Packets in 802.11 Networks. In *Proceedings of the 14th ACM International Conference on Mobile Computing and Networking (MobiCom)*, pages 351–362, 2008.

[100] Vincent Liu, Aaron Parks, Vamsi Talla, Shyamnath Gollakota, David Wetherall, and Joshua R Smith. Ambient Backscatter: Wireless Communication out of Thin Air. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*, pages 39–50, 2013.

[101] Heinrich Meyr, Marc Moeneclaey, and Stefan A. Fechtel. *Digital Communication Receivers: Synchronization, Channel Estimation, and Signal Processing*. John Wiley & Sons, Inc., New York, NY, 1997.

[102] Microsoft Spectrum Observatory. Online: https://observatory.microsoftspectrum.com/.

[103] Miercom. Cisco CleanAir Competitive Testing. Lab Test Report DR100409D, Online: https://www.cisco.com/c/dam/en/us/solutions/collateral/enterprise-networks/cleanair-technology/Miercom_Report_DR100409D_Cisco_CleanAir_Competitive_for_22Apr10.pdf, 2010.

[104] Joseph Mitola. *Cognitive Radio: An Integrated Agent Architecture for Software Defined Radio*. PhD thesis, Royal Institute of Technology (KTH), 2000.

[105] Mobashir Mohammad, XiangFa Guo, and Mun Choon Chan. Oppcast: Exploiting Spatial and Channel Diversity for Robust Data Collection in Urban Environments. In *Proceedings of the 15th ACM/IEEE Conference on Information Processing in Sensor Networks (IPSN)*, pages 1–12, 2016.

[106] Thomas Moscibroda, Ranveer Chandra, Yunnan Wu, Sudipta Sengupta, Paramvir Bahl, and Yuan Yuan. Load-aware Spectrum Distribution in Wireless LANs. In *Proceedings of the 16th annual IEEE International Conference on Network Protocols (ICNP)*, pages 137–146, 2008.

[107] Razvan Musaloiu-E and Andreas Terzis. Minimising the Effect of WiFi Interference in 802.15.4 Wireless Sensor Networks. *International Journal of Sensor Networks*, pages 43–54, 2008.

[108] Nest Thermostat Wireless Specs. Online: https://store.nest.com/product/thermostat?selectedVariantId=T3007ES.

[109] Roland Neumeier and Gerald Ostermayer. Analyzing Coexistence Issues in Wireless Radio Networks – The Simulation Environment. In *Proceedings of the European Modelling Symposium (EMS)*, pages 599–604, 2013.

[110] ns-3 Simulator. Online: https://www.nsnam.org/.

[111] George Nychis, Thibaud Hottelier, Zhuocheng Yang, Srinivasan Seshan, and Peter Steenkiste. Enabling MAC Protocol Implementations on Software-defined Radios. In *Proceedings of the 6th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, pages 91–105, 2009.

[112] OMNeT++ Simulator. Online: https://omnetpp.org/.

[113] OpenMote. OpenMote Platform. Online: http://www.openmote.com/.

[114] Okuary Osechas, Johannes Thiele, Jo Bitsch, and Klaus Wehrle. Ratpack: Wearable Sensor Networks for Animal Observation. In *Proceedings of the 30th IEEE Annual International Conference of Engineering in Medicine and Biology Society*, pages 538–541, 2008.

[115] Jiajue Ou, Yuanqing Zheng, and Mo Li. MISC: Merging Incorrect Symbols Using Constellation Diversity for 802.11 Retransmission. In *IEEE Conference on Computer Communications (INFOCOM)*, pages 2472–2480, 2014.

[116] Marina Petrova, Lili Wu, Petri Mahonen, and Janne Riihijarvi. Interference Measurements on Performance Degradation between Colocated IEEE 802.11g/n and IEEE 802.15.4 Networks. In *Proceedings of the Sixth International IEEE Conference on Networking (ICN)*, pages 93–93, 2007.

[117] Philips SCD 603 digital video baby monitor. Online: http://www.usa.philips.com/c-p/SCD603_10/avent-digital-video-baby-monitor.

[118] Gian Pietro Picco, Davide Molteni, Amy L. Murphy, Federico Ossi, Francesca Cagnacci, Michele Corrà, and Sandro Nicoloso. Geo-referenced Proximity Detection of Wildlife with WildScope: Design and Characterization. In *Proceedings of the 14th ACM/IEEE International Symposium on Information Processing in Sensor Networks (IPSN)*, pages 238–249, 2015.

[119] Joseph Polastre, Robert Szewczyk, and David Culler. Telos: Enabling Ultra-low Power Wireless Research. In *Proceedings of the 4th ACM/IEEE International Symposium on Information Processing in Sensor Networks (IPSN)*, pages 1–6, 2005.

[120] Sofie Pollin, Ian Tan, Bill Hodge, Carl Chun, and Ahmad Bahai. Harmful Coexistence Between 802.15.4 and 802.11: A Measurement-based Study. In *Proceedings of the 3rd International Conference on Cognitive Radio Oriented Wireless Networks and Communications (CrownCom)*, pages 1–6, 2008.

[121] CREW Project. Cognitive Radio Experimentation World. Online: `http://www.crew-project.eu/`.

[122] Qifan Pu, Sidhant Gupta, Shyamnath Gollakota, and Shwetak Patel. Whole-home Gesture Recognition Using Wireless Signals. In *Proceedings of the 19th ACM International Conference on Mobile Computing and Networking (MobiCom)*, pages 27–38, 2013.

[123] Božidar Radunović, Ranveer Chandra, and Dinan Gunawardena. Weeble: Enabling Low-power Nodes to Coexist with High-power Nodes in White Space Networks. In *Proceedings of the 8th international ACM Conference on Emerging Networking Experiments and Technologies (CoNEXT)*, pages 205–216, 2012.

[124] Hariharan Rahul, Farinaz Edalat, Dina Katabi, and Charles G. Sodini. Frequency-aware rate adaptation and mac protocols. In *Proceedings of the 15th ACM Conference on Mobile Computing and Networking (MobiCom)*, pages 193–204, 2009.

[125] Hariharan Rahul, Farinaz Edalat, Dina Katabi, and Charles G Sodini. Frequency-aware Rate Adaptation and MAC Protocols. In *Proceedings of the 15th ACM International Conference on Mobile Computing and Networking (MobiCom)*, pages 193–204, 2009.

[126] Hariharan Rahul, Nate Kushman, Dina Katabi, Charles Sodini, and Farinaz Edalat. Learning to Share: Narrowband-friendly Wideband Networks. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*, pages 147–158, 2008.

[127] Rakesh Rajbanshi, Qi Chen, Alexander M. Wyglinski, Joseph B. Evans, and Gary J. Minden. Comparative Study of Frequency Agile Data Transmission Schemes for Cognitive Radio Transceivers. In *Proceedings of the First ACM International Workshop on Technology and Policy for Accessing Spectrum (TAPAS)*, pages 1–6, 2006.

[128] Raspberry Pi Model B+. Online: `https://www.raspberrypi.org/`.

[129] Shravan Rayanchu, Arunesh Mishra, Dheeraj Agrawal, Sharad Saha, and Suman Banerjee. Diagnosing Wireless Packet Losses in 802.11: Separating Collision from Weak Signal. In *Proceedings of the 27th IEEE Conference on Computer Communications (INFOCOM)*, pages 735–743, 2008.

[130] Shravan Rayanchu, Ashish Patro, and Suman Banerjee. Airshark: Detecting non-WiFi RF Devices Using Commodity WiFi Hardware. In

Proceedings of the ACM SIGCOMM Conference on Internet Measurement Conference (IMC), pages 137–154, 2011.

[131] Dipankar Raychaudhuri, Ivan Seskar, Max Ott, Sachin Ganu, Kishore Ramachandran, Haris Kremo, Robert J. Siracusa, Hang Liu, and Manpreet Singh. Overview of the ORBIT Radio Grid Testbed for Evaluation of Next-generation Wireless Network Protocols. In *IEEE Wireless Communications and Networking Conference (WCNC)*, pages 1664–1669, 2005.

[132] Rulequest Research. Data Mining Tools See5 and C5.0. Online `http://www.rulequest.com/see5-info.html`, 2014.

[133] Mina Sartipi and Faramarz Fekri. Source and Channel Coding in Wireless Sensor Networks Using LDPC Codes. In *Proceedings of the 3rd IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, pages 309–316, 2004.

[134] Florian Schmidt, Matteo Ceriotti, and Klaus Wehrle. Bit Error Distribution and Mutation Patterns of Corrupted Packets in Low-power Wireless Networks. In *Proceedings of the 8th ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation and Characterization (WiNTECH)*, pages 49–56, 2013.

[135] Souvik Sen, Romit Roy Choudhury, and Srihari Nelakuditi. CSMA/CN: Carrier Sense Multiple Access with Collision Notification. In *Proceedings of the 16th ACM International Conference on Mobile Computing and Networking (MobiCom)*, pages 25–36, 2010.

[136] Souvik Sen, Naveen Santhapuri, Romit Roy Choudhury, and Srihari Nelakuditi. AccuRate: Constellation Based Rate Estimation in Wireless Networks. In *Proceedings of the 7th USENIX Conference on Networked Systems Design and Implementation (NSDI)*, pages 175–190, 2010.

[137] Hossein Shafagh and Anwar Hithnawi. Poster: Come Closer - Proximity-based Authentication for the Internet of Things. In *Proceedings of the 20th ACM International Conference on Mobile Computing and Networking (MobiCom)*, pages 421–424, 2014.

[138] Shared Spectrum Company (SSC). Spectrum Occupancy Measurements, Spectrum Reports. Spectrum Reports, Online: `http://www.sharedspectrum.com/papers/spectrum-reports/`, 2009.

[139] Gollakota Shyamnath. *Embracing Interference in Wireless Systems*. PhD thesis, Massachusetts Institute of Technology (MIT), 2013.

[140] Axel Sikora and Voicu F. Groza. 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)*, pages 1786–1791, 2005.

[141] Jeff Slipp, Changning Ma, Nagesh Polu, James Nicholson, Martin Murillo, and Sajid Hussain. WINTeR: Architecture and Applications of a Wireless Industrial Sensor Network Testbed for Radio-Harsh Environments. *Communication Networks and Services Research, Annual Conference on*, pages 422–431, 2008.

[142] Sen Souvik. *Restructuring Wireless Systems using PHY Layer Information*. PhD thesis, Duke University, 2012.

[143] Kannan Srinivasan, Prabal Dutta, Arsalan Tavakoli, and Philip Levis. An Empirical Study of Low-Power Wireless. *ACM Transactions on Sensor Networks (TOSN)*, pages 1–16, 2010.

[144] Kannan Srinivasan, Maria A. Kazandjieva, Saatvik Agarwal, and Philip Levis. The β-factor: Measuring Wireless Link Burstiness. In *Proceedings of the 6th ACM Conference on Embedded Networked Sensor Systems (SenSys)*, pages 29–42, 2008.

[145] Hong Steven. *Architectures and Mechanisms for Wireless Coexistence*. PhD thesis, Stanford University, 2012.

[146] Bluegiga Technologies. BLED112 Bluetooth Smart Dongle. Data Sheet , Online: https://www.bluetooth.org/tpg/RefNotes/BLE112_Datasheet1.pdf.

[147] Texas Instruments. CC2420 Datasheet. Online: http://focus.ti.com/lit/ds/symlink/cc2420.pdf.

[148] Thread Group. Online: http://threadgroup.org/.

[149] David Tse and Pramod Viswanath. *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.

[150] Uniden DCT6485-3HS Cordless Handset System. User Manual, Online: https://www.uniden.com/File%20Library/FooterNav/Product%20Information/Owners%20Manuals/dct6465om.pdf.

[151] Ambuj Varshney, Luca Mottola, Mats Carlsson, and Thiemo Voigt. Directional Transmissions and Receptions for High-throughput Bulk Forwarding in Wireless Sensor Networks. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems (SenSys)*, pages 351–364, 2015.

[152] Vtech GZ2456 User Manual. User Manual, Online: https://images-na.ssl-images-amazon.com/images/I/B1WWMUJzpkS.pdf.

[153] Mythili Vutukuru, Hari Balakrishnan, and Kyle Jamieson. Cross-Layer Wireless Bit Rate Adaptation. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*, pages 3–14, 2009.

[154] Yan Wang, Jian Liu, Yingying Chen, Marco Gruteser, Jie Yang, and Hongbo Liu. E-eyes: Device-free Location-oriented Activity Identification Using Fine-grained WiFi Signatures. In *Proceedings of the 20th ACM International Conference on Mobile Computing and Networking (MobiCom)*, pages 617–628, 2014.

[155] Bo Wei, Wen Hu, Mingrui Yang, and Chun Tung Chou. Radio-based Device-free Activity Recognition With Radio Frequency Interference. In *Proceedings of the 14th International ACM/IEEE Conference on Information Processing in Sensor Networks (IPSN)*, pages 154–165, 2015.

[156] Elias Weingärtner, Hendrik Vom Lehn, and Klaus Wehrle. A Performance Comparison of Recent Network Simulators. In *Proceedings of the IEEE International Conference on Communications (ICC)*, pages 1287–1291, 2009.

[157] Geoffrey Werner-Allen, Patrick Swieskowski, and Matt Welsh. MoteLab: A Wireless Sensor Network Testbed. In *Proceedings of the 4th International ACM/IEEE Conference on Information Processing in Sensor Networks (IPSN)*, pages 1–6, 2005.

[158] WiFi Adapter for Raspberry Pi and Beagle Bone's USB. Online: `http://www.adafruit.com/product/1030`.

[159] Grace R. Woo, Pouya Kheradpour, Dawei Shen, and Dina Katabi. Beyond the Bits: Cooperative Packet Recovery Using Physical Layer Information. In *Proceedings of the 13th ACM International Conference on Mobile Computing and Networking (MobiCom)*, pages 147–158, 2007.

[160] XBee Pro Module. Online: `http://www.adafruit.com/products/964`.

[161] Wenyuan Xu, Wade Trappe, and Yanyong Zhang. Channel surfing: defending wireless sensor networks from interference. In *Proceedings of 6th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 499–508, 2007.

[162] Lei Yang, Wei Hou, Lili Cao, Ben Y Zhao, and Haitao Zheng. Supporting Demanding Wireless Applications with Frequency-agile Radios. In *Proceedings of the 7th USENIX Conference on Networked Systems Design and Implementation (NSDI)*, pages 65–80, 2010.

[163] Lei Yang, Wei Hou, Lili Cao, Ben Y. Zhao, and Haitao Zheng. Supporting demanding wireless applications with frequency-agile radios. In *Proceedings of the 7th USENIX Conference on Networked Systems Design and Implementation (NSDI)*, pages 5–5, 2010.

[164] Shengrong Yin, Qiang Li, and Omprakash Gnawali. Interconnecting WiFi Devices with IEEE 802.15.4 Devices without Using a Gateway. In *Proceedings of the International IEEE Conference on Distributed Computing in Sensor Systems (DCOSS)*, pages 127–136, 2015.

[165] Moustafa Youssef, Matthew Mah, and Ashok Agrawala. Challenges: Device-free Passive Localization for Wireless Environments. In *Proceedings of the 13th ACM International Conference on Mobile Computing and Networking (MobiCom)*, pages 222–229, 2007.

[166] Yan Yubo, Yang Panlong, Li Xiangyang, Tao Yue, Zhang Lan, and You Lizhao. ZIMO: Building Cross-technology MIMO to Harmonize Zigbee Smog with WiFi Flash Without Intervention. In *Proceedings of the 19th ACM International Conference on Mobile Computing and Networking (MobiCom)*, pages 465–476, 2013.

[167] Jun Zhao, Haitao Zheng, and Guang-Hua Yang. Distributed Coordination in Dynamic Spectrum Allocation Networks. In *Proceedings of the First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, pages 259–268, 2005.

[168] Qing Zhao, Lang Tong, Ananthram Swami, and Yunxia Chen. Decentralized Cognitive MAC for Opportunistic Spectrum Access in Ad Hoc Networks: A POMDP Framework. *IEEE Journal on Selected Areas in Communications*, pages 589–600, 2007.

List of Publications

The following list includes publications that form the basis of this thesis. The corresponding chapters are indicated in parentheses.

Anwar Hithnawi, Hossein Shafagh, Simon Duquennoy. **Understanding the Impact of Cross Technology Interference on IEEE 802.15.4.** In *Proceedings of the 9th ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation and Characterization (WiNTECH@MobiCom).* Maui, HI, USA. pp. 49–56, September 2014. (Chapter 2)

Anwar Hithnawi, Hossein Shafagh, Simon Duquennoy. **TIIM: Technology-Independent Interference Mitigation for Low-power Wireless Networks.** In *Proceedings of the 14th ACM International Conference on Information Processing in Sensor Networks (IPSN).* Seattle, WA, USA. pp. 1–12, April 2015. (Chapter 3)

Anwar Hithnawi, Su Li, Hossein Shafagh, James Gross, Simon Duquennoy. **CrossZig: Combating Cross-Technology Interference in Low-power Wireless Networks.** In *Proceedings of the 15th ACM International Conference on Information Processing in Sensor Networks (IPSN).* Vienna, Austria. pp. 1–12, April 2016. (Chapter 4)

Anwar Hithnawi, Su Li, Hossein Shafagh, Simon Duquennoy, James Gross. **Poster Abstract: Cross-Layer Optimization for Low-power Wireless Coexistence.** In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems (SenSys).* Seoul, South Korea. pp. 443–444, November 2015. (Chapter 4)

Anwar Hithnawi, Vaibhav Kulkarni, Su Li, Hossein Shafagh. **Controlled Interference Generation for Wireless Coexistence Research.** In *Proceedings of the Software Radio Implementation Forum (SRIF@MobiCom).* Paris, France. pp. 19–24, September 2015. (Chapter 5)

The following list includes publications that I have co-authored and are not part of this dissertation.

Hossein Shafagh, Lukas Burkhalter, Anwar Hithnawi. **Demo Abstract: Talos a Platform for Processing Encrypted IoT Data.** In *Proceedings of the 14th ACM Conference on Embedded Networked Sensor Systems (SenSys).* Stanford, CA, USA. pp. 1–2, November 2016.

Hossein Shafagh, Anwar Hithnawi, Andreas Dröscher, Simon Duquennoy, Wen Hu. **Talos: Encrypted Query Processing for the Internet of Things.** In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems (SenSys).* Seoul, South Korea. pp. 197–210, November 2015.

Hossein Shafagh, Anwar Hithnawi, Andreas Dröscher, Simon Duquennoy, Wen Hu. **Poster: Towards Encrypted Query Processing for the Internet of Things.** In *Proceedings of the 21th ACM International Conference on Mobile Computing and Networking (MobiCom).* Paris, France. pp. 251–253, September 2015.

Hossein Shafagh, Anwar Hithnawi. **Poster: Come Closer - Proximity-based Authentication for the Internet of Things.** In *Proceedings of the 20th ACM International Conference on Mobile Computing and Networking (MobiCom).* Maui, HI, USA. pp. 421–424, September 2014.

Anwar Hithnawi, Hossein Shafagh, Simon Duquennoy. **Poster Abstract: Low-Power Wireless Channel Quality Estimation in the Presence of RF Smog.** In *Proceedings of the 10th IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS).* Marina Del Rey, California, USA. pp. 137–138, May 2014.

Hossein Shafagh, Anwar Hithnawi. **Poster Abstract: Security Comes First, A Public-key Cryptography Framework for the Internet of Things.** In *Proceedings of the 10th IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS).* Marina Del Rey, California, USA. pp. 135–136, , May 2014.

Anwar Hithnawi. **Poster Abstract: Exploiting Physical Layer Information to Mitigate Cross-Technology Interference Effects on Low-Power Wireless Networks.** In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems (SenSys).* Rome, Italy. pp. 1–2, November 2013.

Florian Schmidt, Anwar Hithnawi, Oscar Punal, James Gross, Klaus Wehrle. **A Receiver-Based 802.11 Rate Adaptation Scheme with On-Demand Feedback.** In *Proceedings of the 23rd IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC).* Sydney, Australia. pp. 1–7, Septemper 2012.