Towards Reliable Communication in Low-Power Wireless Body Area Networks

vorgelegt von
Diplom-Informatiker
Jan-Hinrich Hauer
aus Berlin

von der Fakultät IV
Elektrotechnik und Informatik
der Technischen Universität Berlin
zur Erlangung des akademischen Grades

Doktor der Ingenieurwissenschaften
– Dr.-Ing. –

genehmigte Dissertation

Promotionsausschuss:

Vorsitzender: Prof. Dr.-Ing. Hans-Joachim Grallert
Berichter: Prof. Dr.-Ing. Adam Wolisz
Berichter: Prof. Dr. Roberto Verdone
Berichter: Prof. Dr.-Ing. Rolf Kraemer

Tag der wissenschaftlichen Aussprache: 30. Januar 2014

Berlin 2014
D83

Zusammenfassung

Es wird zunehmend die Ansicht vertreten, dass tragbare Computer und Sensoren neue Anwendungen in den Bereichen Gesundheitswesen, personalisierte Fitness oder erweiterte Realität ermöglichen werden. Die am Körper getragenen Geräte sind dabei mithilfe eines Wireless Body Area Network (WBAN) verbunden, d.h. es wird drahtlose Kommunikation statt eines drahtgebundenen Kanals eingesetzt. Der drahtlose Kanal ist jedoch typischerweise ein eher instabiles Kommunikationsmedium und die Einsatzbedingungen von WBANs sind besonders schwierig: Einerseits wird die Kanalqualität stark von den physischen Bewegungen der Person beeinflusst, andererseits werden WBANs häufig in lizenzfreien Funkbändern eingesetzt und sind daher Störungen von anderen drahtlosen Geräten ausgesetzt. Oft benötigen WBAN Anwendungen aber eine zuverlässige Datenübertragung.

Das erste Ziel dieser Arbeit ist es, ein besseres Verständnis dafür zu schaffen, wie sich die spezifischen Einsatzbedingungen von WBANs auf die intra-WBAN Kommunikation auswirken. So wird zum Beispiel analysiert, welchen Einfluss die Platzierung der Geräte auf der Oberfläche des menschlichen Körpers und die Mobilität des Benutzers haben. Es wird nachgewiesen, dass während regelmäßiger Aktivitäten wie Laufen die empfangene Signalstärke stark schwankt, gleichzeitig aber Signalstärke-Spitzen oft einem *regulären* Muster folgen. Außerdem wird gezeigt, dass in urbanen Umgebungen die Effekte von 2.4 GHz Radio Frequency (RF) Interferenz im Vergleich zu den Auswirkungen von *fading* (Schwankungen der empfangenen Signalstärke) eher gering sind. Allerdings führt RF Interferenz dazu, dass häufiger Bündelfehler auftreten, d.h. Fehler zeitlich korrelieren. Dies kann insbesondere in Anwendungen, die eine geringe Übertragungslatenz benötigen, problematisch sein.

Der zweite Teil dieser Arbeit beschäftigt sich mit der Analyse von Verfahren, die potentiell die Zuverlässigkeit der Kommunikation in WBANs erhöhen, ohne dass wesentlich mehr Energie verbraucht wird. Zunächst wird der Trade-off zwischen Übertragungslatenz und der Zuverlässigkeit der Kommunikation analysiert. Diese Analyse basiert auf einem neuen *Paket-Scheduling Algorithmus*, der einen Beschleunigungssensor nutzt, um die WBAN Kommunikation auf die physischen Bewegungen der Person abzustimmen. Die Analyse zeigt, dass unzuverlässige Kommunikationsverbindungen oft zuverlässig werden, wenn Pakete während vorhergesagter Signalstärke-Spitzen gesendet werden. Ferner wird analysiert, inwiefern die Robustheit gegen 2.4 GHz RF Interferenz verbessert werden kann. Dazu werden zwei Verfahren betrachtet: Ein bereits existierendes Verfahren, das periodisch einen Wechsel der Übertragungsfrequenz durchführt (*channel hopping*) und ein neues Verfahren, das durch RF Interferenz entstandene Bitfehler reparieren kann, indem der Inhalt mehrerer fehlerhafter Pakete kombiniert

wird (*packet combining*). Eine Schlussfolgerung ist, dass Frequenzdiversität zwar das Auftreten von Bündelfehlern reduzieren kann, dass jedoch die statische Auswahl eines Kanals am oberen Ende des 2.4 GHz Bandes häufig schon eine akzeptable Abhilfe gegen RF Interferenz darstellt.

Abstract

There is a growing belief that wearable computers and sensors will enable new applications in areas such as healthcare, personal fitness or augmented reality. The devices are attached to a person and connected through a Wireless Body Area Network (WBAN), which replaces the wires of traditional monitoring systems by wireless communication. This comes, however, at the cost of turning a reliable communication channel into an unreliable one. The wireless channel is typically a rather unstable medium for communication and the conditions under which WBANs have to operate are particularly harsh: not only is the channel strongly influenced by the movements of the person, but WBANs also often operate in unlicensed frequency bands and may therefore be exposed to a significant amount of interference from other wireless devices. Yet, many envisioned WBAN applications require reliable data transmission.

The goals of this thesis are twofold: first, we aim at establishing a better understanding of how the specific WBAN operating conditions, such as node placement on the human body surface and user mobility, impact intra-WBAN communication. We show that during periodic activities like walking the received signal strength on an on-body communication link fluctuates strongly, but signal strength peaks often follow a *regular* pattern. Furthermore, we find that in comparison to the effects of fading 2.4 GHz Radio Frequency (RF) interference causes relatively little packet loss — however, urban 2.4 GHz RF noise is *bursty* (correlated in time), which may be problematic for applications with low latency bounds.

The second goal of this thesis is to analyze how communication reliability in WBANs can be improved without sacrificing a significant amount of additional energy. To this end, we first explore the trade-off between communication latency and communication reliability. This analysis is based on a novel *packet scheduling algorithm*, which makes use of an accelerometer to couple WBAN communication with the movement patterns of the user. The analysis shows that unreliable links can often be made reliable if packets are transmitted at predicted signal strength peaks. In addition, we analyze to what extent two mechanisms can improve robustness against 2.4 GHz RF interference when adopted in a WBAN context: we analyze the benefits of *channel hopping*, and we examine how the packet retransmission process can be made more efficient by using a novel *packet combining algorithm* that allows to repair packets corrupted by RF interference. One of the conclusions is that while frequency agility may decrease "burstiness" of errors the static selection of a channel at the upper end of the 2.4 GHz band often already represents a good remedy against RF interference.

Acknowledgments

First of all, I would like to express my sincere gratitude to my advisor, *Prof. Adam Wolisz*, for the scientific mentorship and (financial) support throughout all those years, for the freedom to explore new ideas and for the very open and productive discussions we have had.

The work in this dissertation has been supported in part through the European Commission projects ANGEL, CONET and CREW. In particular, one of the presented ideas was born within the CONET project through the interaction with *Shashi Prabh* from Polytechnic Institute of Porto, Portugal. The CONET project also allowed me to collaborate with *Carlo Boano*, *Marco Zúñiga*, *Thiemo Voigt* and *Mário Alves*.

At the Telecommunication Networks Group *Dr. Vlado Handziski* was a great influence to me. Not only had he been an excellent supervisor of my Diploma thesis and sparked my activities in the TinyOS 2.x Core Working Group, but he also helped me grow as a professional with inexhaustible advice during my PhD. During the ANGEL-project I also received valuable guidance and support from *Dr. Andreas Willig* and I'm also grateful for the help from the students *Moksha Birk* and *Jasper Büsch*.

The pleasant and cooperative atmosphere at the TKN group was created for me by both, the "old generation": *Niels Karowski*, *Dr. Mathias Bohge*, *Dr. Daniel Willkomm*, *Sven Wiethölter* and *Filip Idzikowski*, as well as the "newer guys": *Manoj Rege*, *Hieu Le* and *Mikolaj Chwalisz*. Of course, life at TKN would not have been the same for me without *Onur Ergin*, our many cake-and-coffee breaks and the (pseudo?)-philosophical insights we shared. I'm also grateful to *Sven Spuida* for supporting me in creating various measurement setups, and last but certainly not least for the tireless support of *Petra Hutt*.

Finally, I wish to thank my girlfriend, my parents and my brother for their support and huge patience especially during the last months of writing.

Contents

List of Figures

xiii

xiv

List of Tables

Acronyms

6LoWPAN IPv6 over Low power WPAN

ACF Autocorrelation Function

ACK Acknowledgement Message

ADC Analog-to-Digital Converter

ADF Average Duration of Fades

AFH Adaptive Frequency Hopping

AMDF Average Magnitude Difference Function

AP Access Point

ARQ Automatic Repeat reQuest

ASN Absolute Slot Number

ATM Automated Teller Machine

BER Bit Error Rate

BLE Bluetooth Low Energy

CAP Contention Access Phase

CCA Clear Channel Assessment

CDF Cumulative Distribution Function

CDMA Code Division Multiple Access

CFP Contention-Free Period

CoAP Constrained Application Protocol

COPS Centralized Opportunistic Packet Scheduler

COTS Commercial-Off-The-Shelf

CPRR Corrupt Packet Reception Rate

CPU Central Processing Unit

CRC Cyclic Redundancy Check

CR Cognitive Radio

CSMA Carrier Sense Multiple Access

CSMA/CA CSMA with Collision Avoidance

CSS Chirp Spread Spectrum

CTS Clear To Send

DCO Digitally Controlled Oscillator

DSSS Direct-Sequence Spread Spectrum

DTW Dynamic Time Warping

EAP Exclusive Access Phase

ECG Electrocardiography

EDR Enhanced Data Rate

EEG Electroencephalography

EIRP Equivalent Isotropically Radiated Power

EMG Electromyography

ETSI European Telecommunications Standards Institute

ETX Expected Transmission Count

FCC U.S. Federal Communications Commission

FDMA Frequency Division Multiple Access

FDTD Finite Difference Time Domain

FEC Forward Error Correction

FHSS Frequency-Hopping Spread Spectrum

GFSK Gaussian Frequency Shift Keying

GPS Global Positioning System

GSM Global System for Mobile Communications

GSR Galvanic Skin Response

GTS Guaranteed Time Slot

HAA Hardware Abstraction Architecture

HBC Human Body Communications

HTTP Hyper Text Transfer Protocol

IC Integrated Circuit

IEEE Institute of Electrical and Electronics Engineers

IETF Internet Engineering Task Force

IP Internet Protocol

IPv6 Internet Protocol version 6

IQR Interquartile Range

ISM Industrial Scientific and Medical

LCR Level-Crossing Rate

LOS Line-Of-Sight

LP Linear Programming

LPL Low-Power Listening

LPP Low-Power Probing

LQI Link Quality Indicator

M2M Machine to Machine

MAC Medium Access Control layer

MAP Managed Access Phase

MBAN Medical Body Area Network

MCPS MAC Common Part Sublayer

MCU Micro-Controller Unit

MICS Medical Implant Communication Service

MIT Massachusetts Institute of Technology

MLE Maximum Likelihood Estimation

MLME MAC Sublayer Management Entity

MPDU MAC Protocol Data Unit

MSC Message Sequence Chart

MSDU MAC Service Data Unit

NACK Negative ACK

NLOS Non-Line-Of-Sight

OFDM Orthogonal Frequency Division Multiplexing

O-QPSK Offset Quadrature Phase-Shift Keying

OSI Open Systems Interconnection

OS Operating System

OTW Opportune Transmission Window

PAN Personal Area Network

PC Personal Computer

PHY Physical Layer

PIB PAN Information Base

PLR Packet Loss Rate

PPDU PHY Protocol Data Unit

PRR Packet Reception Rate

PSD Power Spectral Density

RAM Random-Access Memory

RAP Random-Access Phase

REPE RSSI-based Bit Error Position Estimation

RFID Radio-Frequency Identification

RF Radio Frequency

RISC Reduced Instruction Set Computing

RSSI Received Signal Strength Indication

RTS Request To Send

RWMM Random Walk Mobility Model

SCIP Solving Constraint Integer Programs

SD Secure Digital

SFD Start Frame Delimiter

SIG Special Interest Group

SINR Signal to Interference plus Noise Ratio

SMD Surface-Mount Device

TCP Transmission Control Protocol

TDMA Time Division Multiple Access

TEP TinyOS Enhancement Proposal

TSCH Timeslotted Channel Hopping

TSMP Time Synchronized Mesh Protocol

TWIST TKN Wireless Indoor Sensor Network Testbed

UDP User Datagram Protocol

USART Universal Asynchronous Receiver/Transmitter

USB Universal Serial Bus

UWB Ultra-Wideband

WBAN Wireless Body Area Network

WirelessHART Wireless Highway Addressable Remote Transducer Protocol

WLAN Wireless Local Area Network

WMTS Wireless Medical Telemetry Service

WPAN Wireless Personal Area Networks

WSN Wireless Sensor Network

Chapter 1

Introduction

The past decade has seen the rise of a technology that promises to join the physical with the virtual world. Tiny computers with sensing capabilities are now able to collect information about real objects with an unprecedented detail and pervasiveness. As the miniaturization of electronic components continues and computing capacity becomes evermore cheaper we may soon see tens of billions of these devices connected to the Internet. By interconnecting real-world sensing with the massive computational and storage capabilities in the cloud many new possibilities will emerge (as well as great challenges to protect our privacy): we may, for example, be able to better understand and control the traffic on the roads, the status of equipment and industrial machinery or the micro-climate on farms.

Connecting sensors to the Internet often involves wireless communication, because it facilitates rapid deployments even at hard-to-reach locations and enables mobile sensors. The possibilities of mobile sensing have lately brought people and their behavior into focus. Human-mounted sensors may, for examples, assist people in their interaction with objects, services or persons in their environment. They also enable mobile health monitoring, which is of key importance given that population ageing, an increasing number of chronic diseases and a prevalence of unhealthy lifestyles pose significant challenges for current and future societies [9]. The technology that is expected to pave the path towards these new possibilities are Wireless Body Area Networks (WBANs).

1.1 Wireless Body Area Networks

A WBAN is composed of several small untethered devices, which are attached to, or even implanted in, the human body. Each device includes some sensor(s) that may capture physiological parameters of the human body such as heart rate, blood pressure or Electrocardiography (ECG), and/or characteristics of the environment such as the ambient air quality or the proximity of certain objects. The sensor data are collected and possibly (pre)processed by a Micro-Controller Unit (MCU). With the help of short-range wireless transmission they are then often transferred from the sensor devices to a more capable device, a central controller such as the user's cell phone. There the data may be analyzed and presented to the user or transmitted via the Internet to a remote server for storage, processing and analysis.

WBANs are often regarded a special type of Wireless Sensor Network (WSN). WSNs also consist of sensors that enable wireless monitoring of environmental conditions. But while "classical" WSNs are static, deployed over some geographic area, WBANs are attached to a mobile human body. WBANs thus have the potential to facilitate the treatment of chronic diseases, reduce visits to doctors and hospitals and to help in post-operative monitoring [1]. They may also support elderly people in daily routines and medical conditions [10]. Other WBAN usage scenarios go beyond personalized health-care: WBANs can monitor and support firefighters or policemen in hazardous situations [11] as well as soldiers in the battlefield [12]; they may be used for sports and personal fitness monitoring [13], in multimedia applications or even interact with virtual objects in augmented reality scenarios [14].

Traditional health monitoring systems are bulky and use wired connections. Consequently the person's comfort and mobility are often severely constrained. WBANs enable less obtrusive monitoring by replacing the wires by wireless communication. This, however, comes at the cost of turning a reliable communication channel into an unreliable one: the wireless channel is an unstable medium for communication and it becomes even more dynamic when the nodes are mobile.

Yet, many envisioned WBAN applications require reliable data transmission. Especially in medical applications data loss is problematic, because it may lead to wrong or late diagnosis [15]. The reliable delivery of data is therefore considered one major requirement for WBANs [16, Chapter 1]. While non-functional properties such as privacy/security, standard-compliance or scalability, may be equally important in certain application scenarios, an overarching requirement in WBANs is energy-efficiency: one consequence of replacing the wires by wireless communication is that a WBAN device must have its own energy source. WBANs are therefore typically battery-powered. In many scenarios the batteries have to last for several weeks or months, and consequently, in addition to reliable communication, energy-efficient operation is one major design goal.

WBAN devices therefore consist of extremely low-power hardware components. For example, the MCU may contain only a 8- or 16-bit processor and a few kilobytes of memory. The largest fraction of the overall energy budget is, however, usually spent by the wireless transceiver as it typically consumes orders of magnitude more energy than the MCU. The transceiver must therefore not only consist of extremely low-power hardware components, but communication protocols must also use it efficiently, e.g. by shutting off the radio whenever possible. While this is also true in other low-power wireless networks, such as "classical" WSNs, communication in WBAN is faced with an additional characteristic: the mobility of the user, which as stated by Dutta and Culler "invalidates many assumptions implicit in today's low-power, static sensor network designs" [17].

In WBANs wireless communication occurs along and through the human body, communication quality is therefore strongly influenced by the posture (changes) of the person. Furthermore, since WBANs are mobile and often operate in unlicensed frequency bands they may temporarily be exposed to a significant amount of interference from other co-located wireless communication technologies. This often results in rather harsh and dynamic conditions for the WBAN communication processes. To compensate for the harsh conditions it may be tempting to simply increase the radio's transmission power. Unfortunately, this approach is often ill-suited for WBANs, be-

cause higher transmission power not only results in higher energy consumption, but WBAN devices should also "transmit lower power when possible in order to reduce interference to other devices and systems" (IEEE 802.15.4 Standard [18, Sect. 6.9.5]). High transmission power is also problematic, because it leads to greater absorption of electromagnetic energy by the human body, which is particularly crucial for implanted sensors. Reconciling the requirement of reliable data transmission and low-power wireless communication under the harsh channel conditions experienced in WBANs results in a set of challenges that are quite unique in comparison to those experienced in other wireless networks.

1.2 Goals of the Thesis

WBANs are a relatively new field of research and the first commercial solutions have only recently become available. In comparison to the amount of research that has been conducted in the WSN domain, there are relatively few experimental studies investigating the characteristics of wireless communication in a WBAN. Yet it is clear that the WBAN channel behaves quite differently from a WSN channel: in a WBAN the nodes are attached to a person who may change posture and move around in different indoor/outdoor (interference) environments, whereas WSNs typically consist of nodes that are spread over some geographic area and remain static once the network has been deployed. Since the conditions under which WSN and WBAN communication protocols operate are quite different many approaches that have been developed for WSNs may not be directly applicable in the WBAN domain. Instead, approaches that take the specific WBAN conditions into account are expected to be more suitable.

The goals of this thesis are twofold: first, we aim at establishing a better understanding of how the specific WBAN operating conditions, such as node placement on the human body surface and user mobility, impact intra-WBAN communication. This analysis relates to the intended communication among WBAN devices as well as the impact of interfering signals from other co-located devices. The investigation is important, because the harsh and dynamic conditions under which WBAN communication must operate are not well-understood.

The second goal of this thesis is to analyze how communication reliability in WBANs can be improved without sacrificing a significant amount of additional energy. To this end, we first explore the trade-off between communication latency and communication reliability. This analysis covers a novel *packet scheduling algorithm*, which couples the WBAN communication with the movement patterns of the user. Afterwards, we analyze to what extent two existing mechanisms can improve robustness against Radio Frequency (RF) interference when adopted in a WBAN context: we analyze the trade-offs of *channel hopping*, and we examine how the packet retransmission process can be made more efficient by using a novel *packet combining algorithm* that allows to repair packets corrupted by RF interference.

1.3 Contributions

Conceptually, the challenges addressed in this thesis can be mapped to the components of a well-known communication model that describes the quality of a wireless communication link with

its Signal to Interference plus Noise Ratio (SINR). The SINR model states that a sufficiently high — the actual threshold depends on hardware characteristics — SINR value leads to a very high probability of correctly detecting an incoming bit on the receiver side. Thus the SINR directly determines the quality of a wireless communication link. The SINR consists of three components:

$$SINR = \frac{P_{Signal}}{P_{Interference} + P_{Noise}} \tag{1.1}$$

P_{Signal} is the power of the (intended) received signal, $P_{Interference}$ represents the sum of the power of any interfering (unintended) signals and P_{Noise} is the noise power. In wireless systems P_{Signal} and $P_{Interference}$ may fluctuate significantly, sometimes over very short time intervals. In WBANs, however, these two parameters are generally more dynamic than in other wireless networks due to the user's mobility. Essentially, this thesis aims at (a) establishing a better understanding of the characteristics of P_{Signal} and $P_{Interference}$ in realistic WBAN scenarios and (b) providing an analysis of protocol mechanisms that can "improve" these parameters in order to increase communication reliability without consuming a significant amount of additional energy. To this end the first part of this thesis concentrates on P_{Signal} in the numerator in Eq. 1.1, and the second part deals with the $P_{Interference}$ parameter in the denominator (the P_{Noise} component in Eq. 1.1 is not addressed in thesis, because it depends on the electronic components of the receiver radio circuitry and is mainly influenced by thermal noise).

Signal Propagation ($\mathbf{P_{Signal}}$): We have conducted a large number of experiments with Commercial-Off-The-Shelf (COTS) WBAN hardware attached to various body positions of different subjects moving in realistic urban indoor and outdoor environments. Through careful experiment design we were able to focus only on the (intended) received power of an intra-WBAN communication link, i.e. only on the P_{Signal} parameter. In line with previous work we found that although the distances between WBAN nodes are small the wireless signal experiences severe attenuation through the human body. We also found that changes in the posture can result in significant fluctuations in received signal power. The main focus of our experiments, however, is on the impact of human mobility on received signal power. To this end we equipped different subjects with eight WBAN devices and monitored the behavior of P_{Signal} while a subject was walking. Our results show that the magnitude of P_{Signal} fluctuation is rather strong and that the received signal power exhibits correlation with the movements (step frequency) of the subject.

A major contribution of this thesis is to analyze to what extent the P_{Signal} dynamics may be considered in balancing communication reliability and latency. We show that by deferring communication until P_{Signal} is "stronger than average" a WBAN can indeed significantly improve the chances of successful communication and thus increase communication reliability. However, this approach comes at the cost of an increase in communication delay. To better understand this trade-off we analyzed a novel *packet scheduler* algorithm, that adapts the WBAN communication to the movement pattern of the user. It utilizes information from motion sensors to detect quasi-periodic user movements and applies this knowledge in the process of scheduling communication. An extensive experimental evaluation involving seven different subjects shows that

by introducing an average delay of about $0.7\,\mathrm{s}$ this mechanism can indeed significantly increase intra-WBAN communication reliability while spending only little additional energy.

RF Interference ($\mathbf{P_{Interference}}$): Due to cost reasons, WBANs often operate in unlicensed communication bands. The majority of commercial WBAN systems use the popular $2.4\,\mathrm{GHz}$ band, which is also shared with, for example, Wireless Local Area Network (WLAN). One consequence of using unlicensed bands is that WBANs are exposed to RF interference from other co-located wireless technologies. Interfering signals may disturb WBAN communication and thus degrade the reliability of the data transmission. In the WSN domain there have been several proposals to make communication robust against RF interference [19–22]. In WBANs, however, user mobility results in a quite different perception of RF interference, since the user often passes by many interfering sources within short time. Unfortunately, the impact of RF interference on WBAN communication in *realistic* environments is not well-understood. Therefore we have conducted one of the first experimental studies on how urban RF interference affects WBAN communication performance. Our results indicate that the overall impact of RF interference on WBAN communication reliability is small, but urban RF noise is generally "bursty" (correlated in time). In particular in applications with low latency bounds RF interference mitigation techniques may therefore be advisable.

To better understand how the effects of RF interference can be mitigated we have analyzed the applicability of two well-known techniques: *packet combining* and *channel hopping*. One widely-used approach to increase communication reliability is to simply send an additional copy of a packet in case the previous transmission was unsuccessful. Such retransmissions come at the price of increased resource usage, in particular energy. In this thesis we investigate how the retransmission process in WBANs can be improved in the presence of RF interference. To this end we introduce and analyze a packet combining algorithm that is able to reduce the energy consumption of the retransmission process, which in effect allows to increase reliability given a fixed energy budget. Finally, we analyze the applicability of channel hopping in WBANs. Channel hopping mitigates the effects of RF interference by continuously changing the communication frequency. An interfering source that affects only some frequencies will thus only temporary degrade communication performance. Channel hopping has by now been introduced in many low-power wireless networks, although its benefits are are debatable ([23] vs. [24]). The topic has, however, not received much attention in the context of WBANs. The final contribution of this thesis is therefore to better understand to what extent channel hopping can improve communication reliability in WBANs.

To analyze the above-mentioned mechanisms one option is to use analytical models or network simulators. A generic analytical model for WBAN communication, however, does not exist, because WBAN operating conditions are too diverse and the model would need to take a large number of parameters into account. Instead, sometimes empirical data is used to construct distributions from which a simulator may perform probabilistic sampling. Unfortunately, this implies that the simulation is not able to accurately model temporal dependencies. Our measurements, however, show that when the user is mobile (walking) there is a significant amount of temporal correlation in the P_{Signal} and $P_{Interference}$ signal.

For many years the WSN community has been acknowledging that experimental research is of key importance when dealing with wireless communication. The construction of the many experimental test facilities [25–28] gives evidence of this trend. These facilities have enabled many empirical WSN studies on wireless communication [29–31] and the impact of interference [21, 23, 24]. They have also fueled a growing consensus that wireless systems require experimental evaluation, which has lately manifested in the trend to experimentally evaluate WSN protocols in multiple testbeds [32–34]

In WBANs the channel conditions are often harsher than in many WSN scenarios. Therefore, rather than relying on a particular model, the main evaluation method in this thesis is based on experiments with real WBAN hardware involving several subjects in various realistic indoor/outdoor environments. Furthermore, we adopt the method of "replaying" the results that we have collected in our measurements in a simulator. The advantage is that we can repeat an analysis under identical, realistic conditions. The disadvantage is that there are only a limited number of simulation runs possible, because the number of experimentation scenarios is limited (which is also the main drawback of performing experiments in the first place).

1.4 Outline

The rest of this thesis is structured as follows: first some background information on WBAN applications and related network technologies is provided; then the scope and limitations of the thesis are discussed; and in the following four chapters the above research questions are answered, each chapter focusing on one question. The thesis completes with a conclusion chapter that summarizes the main results and provides directions for future research. This is detailed in the following overview:

- **Chapter 2 – Background:** An introduction to WBAN applications, sensors, platforms and protocol architectures is provided. The characteristic properties of wireless on-body communication are discussed. Recent standardization efforts such as IEEE 802.15.6, IEEE 802.15.4e and Bluetooth low energy are introduced. Finally, our IEEE 802.15.4 Medium Access Control layer (MAC) implementation "TKN15.4" is introduced. The latter has been published in [35–37].

- **Chapter 3 – Scope of the Thesis:** A motivation for the research questions addressed in this thesis is given. The goals of this thesis are detailed and the methods of achieving them are explained.

- **Chapter 4 – The Wireless On-Body Channel During Walking:** This chapter studies the impact of human mobility on on-body communication. Results from an experimental study involving eight different node positions and five different subjects are compared. Parts of this work have been published in [38].

- **Chapter 5 – Opportunistic Packet Scheduling:** Based on the finding in the previous chapter a packet scheduling algorithm is proposed that exploits periodic mobility for on-body communication. The algorithm utilizes the output from motion sensors to schedule

communication events. It also optimizes channel access by computing non-overlapping schedules using a linear programming technique. The chapter includes an extensive experimental performance evaluation of the proposed solution. The algorithms presented in this chapter have previously been published in [38] and the initial idea was introduced in [39].

- **Chapter 6 – The Impact of External RF Interference on On-Body Communication:** In this chapter we report on an extensive measurement study that captured urban 2.4 GHz RF noise in several different realistic environments. We also analyze the impact of controlled as well as uncontrolled urban RF interference on an WBAN communication link. Parts of this work have been published in [40] and [41].

- **Chapter 7 – RF Interference Mitigation:** Based on the findings in the previous chapter two interference mitigation strategies are investigated in this chapter. First, the potential benefits of channel hopping are evaluated via a trace-driven simulation, which involves a large dataset of urban RF noise measurements. Second, we investigate how a receiver can detect RF interference during packet reception and by closely monitoring received signal strength estimate the bit error positions in corrupt packets. We propose a packet combining extension to the IEEE 802.15.4 MAC that retransmits only corrupted bits to save energy and communication bandwidth. Parts of this work have been published in [42].

- **Chapter 8 – Conclusions:** The thesis is concluded by a short review of the proposed solutions and a summary of the major contributions. In addition, directions for future research are provided.

Chapter 2

Background

This chapter provides background information that facilitates the understanding of the following chapters. We start with an overview of WBAN applications and sensors and report on current WBAN hardware platforms and protocol architectures. Afterwards we review the characteristics of wireless on-body communication and discuss current WBAN channel models. The last part of this chapter is dedicated to WBAN communication protocols with a particular focus on recent standardization efforts such as IEEE 802.15.6, IEEE 802.15.4e and Bluetooth Low Energy (BLE).

2.1 Wireless Body Area Networks

WBAN applications allow to measure physiological parameters of the human body. This makes it possible to assess the general health situation or monitor special medical conditions; to understand body reactions during sports and personal fitness activities; or to monitor the body status in hazardous situations. WBAN applications may also involve monitoring the behavior or location of a person, for example, to assist the person in his/her activities and daily routines; and WBANs can monitor the person's immediate environment to identify certain interesting events (or threats) or to assist the person in his/her interaction with the environment. In order to understand what amount of data WBAN applications produce and, correspondingly, what traffic rates the communication processes must support the following sections review representative WBAN applications and give an overview of the sensors that these applications use. By analyzing the typical sampling rates and precision of sensor data types we derive the data rates that these applications produce. Afterwards we report on current WBAN platforms and protocol architectures to contrast the data rates with the capabilities of current systems.

2.1.1 Applications

WBAN application scenarios can be classified into medical, assisted living, personal fitness, disaster relief and augmented reality scenarios, although a single application may belong in multiple categories. While some (medical or sports) applications are already available in commercial products today, many of the applications described below are in research/prototype stadium.

WBAN *medical* applications, are often categorized into clinical and at home (or outdoor) scenarios. Applications belonging to the first category assist in the monitoring of patients in hospitals in order to detect deterioration of their health status. Such checks are currently often performed manually by medical personnel and WBAN are envisioned to greatly facilitate tasks such as postoperative nursing care. For example, critical events in hospitals that require to displace a patient after operation to intensive care are rare, but when they occur they are typically preceded by changes in vital signs that can be observed several hours earlier [43]. A WBAN can be used to monitor body parameters such as heart rate, ECG or blood oxygen saturation levels to detect deteriorations in advance. Several WBAN prototypes have been deployed in hospital scenarios in the past [44–46] and these pilot studies clearly indicated the benefits of using WBANs for real-time monitoring in clinical settings.Medical applications are also used in the context of home care or outdoor mobile monitoring. One application is continuous cardiac monitoring for high risk patients. For example, Zhou et al. [47] have developed a WBAN system that allows real-time continuous ECG monitoring. Their system uses a 4-lead ECG and is capable of cardiac arrhythmia detection. The data are transmitted wirelessly to a local server (laptop or mobile phone) and/or forwarded to a remote server for inspection by medical personnel. Other medical applications are a seizure monitor for Epilepsy patients [48], estimation of the motor complications in patients with Parkinson's disease [49] or the detection of walking gait impairment with an accelerometer to identify injuries after surgery [50].

There are also several applications envisioned to support (elderly) people in their daily routines. These *assisted living* applications aim at improving personal safety and social contact or increasing the independence of elderly people. For example, mobility may be used as an indicator of the activity and health status of an elderly person. Woods et al. have proposed a system that uses motion sensors to monitor the behavior of a person [51]. Their system exploits the fact that most people have behavioral trends that follow a 24 h cycle. Whenever the person's behavior deviates from this rhythmic behavioral pattern, the system detects these changes, and may generate an appropriate alarm to caregivers. A similar system is proposed by Fernandez et al. [52], which additionally includes environmental sensors, such as magnetometers that detect opening of doors. Motion sensors are also used for wearable *sports and personal fitness* applications, which have lately received increasing attention from industry. These applications allow to monitor the status of the body during workout with the help of, for example, heart-rate belts or pedometers.

WBANs are also envisioned to help in *hazardous situations* or mass casualty events. For example, Wilson et al. have developed a system to support firefighters involved in industrial firefighting and emergency response scenarios [11]. With their system temperature, smoke, and carbon monoxide may be monitored to identify dangerous situations. Sensor nodes are also used to localize people and find safe routes. To this end they are either pre-installed inside the building, or firefighters drop sensor nodes "as they go, creating a smart breadcrumb trail that can be followed to escape danger zones" [11]. George at al. proposed a similar disaster response system that includes audio sensors to reconstruct audio signals, such as human voice, using a distributed low-frequency sampling approach [53]. Yet another set of WBAN applications is related to military scenarios: by providing remote real-time access to physiological data and enabling faster assessment of medical conditions WBANs can help to support soldiers in the

Sensor	Sampling rate (Hz)	Precision (bits)	Data rate (bps)
ECG (6 leads)	1000	12	72000
EEG (12 leads)	350	12	50000
EMG	20000	16	320000
Heart rate	4	12	48
Body temperature	15	8	120
Blood oxygen saturation	1	12	12
Galvanic skin response	32	12	384
Blood pressure	100	12	1200
Blood glucose	0.003	12	0.04
Accelerometer	20	12	240
Gyroscope	20	12	240
Magnetometer	5	12	60
Air quality (ozone)	0.2	12	2.4

Table 2.1: Characteristics of common WBAN sensors, sources: [1–7].

battlefield [12] or categorize injured soldiers by urgency [54].

Finally, Barrie et al. describe how they use WBANs as an interface for *interaction with the environment* [14]. They use body motion sensors to identify user gestures and map them to certain pre-defined commands, such as regulating the volume of an audio player. They also envision scenarios that realize mixed reality situations for gaming or human-to-human interaction. Closely connected to this topic is the area of *location-based services*, which have lately gained increasing popularity: these services identify the geographic location of a person and help to discover persons or objects in the environment. For example, theses services may help to find a nearby Automated Teller Machine (ATM) or restaurants but also request assistance in case of an medical condition. Currently these services use mobile phones equipped with Global Positioning System (GPS), however, WBANs may also support similar applications in an Ad-Hoc fashion.

2.1.2 Sensors

WBAN sensors measure a physical quantity and convert it into a signal that can be processed by the WBAN node, typically involving an Analog-to-Digital Converter (ADC). In this section we give an overview of sensors that are commonly used in WBAN applications. Since the focus of this thesis is on WBAN communication we concentrate on the involved sampling and data rates, because they define the communication traffic. This information is also summarized in Table 2.1.

ECG sensors measure the electrical activity of the of the heart. This allows to determine the rate and regularity of heartbeats and thus diagnose heart abnormalities. The amount of ECG electrodes involved may vary, for example, there are 3-lead, 5-lead or 12-lead ECGs. ECG data rates are comparably high, because the signal bandwidth is rather large and data are captured on

multiple electrodes in parallel. *Electroencephalography (EEG)* sensors measure the electrical activity of the brain. Typically electrodes are placed on a person's scalp at distinct locations. EEG can, for example, be used to determine the epileptic activity of a person. Similar to ECG sensors, EEG sensors result in a high data rate, because multiple electrodes are involved, but signal bandwidth is typically lower than for ECG. *Electromyography (EMG)* sensors measure the electrical activity of muscles. EMG is used to analyze movements, muscle activation level and medical abnormalities. EMG helps to detect neuromuscular diseases and disorders of motor control. The data rates are high, because the signal bandwidth can be up to $10\,\mathrm{kHz}$ [1].

Heart rate may be monitored to determine medical conditions, i.e. abnormalities such as unnaturally fast or slow heart rate or rhythmic deviation. It is typically measured in beats per minute. Heart rate monitoring does not involve high data rates, because sampling once every few seconds is sufficient. *Body temperature* can be analyzed to detect a possible medical condition. Lately thermometers that measure the temperature near the eardrum have become popular, because this point reflects the body core temperature very well. Measuring body temperature involves rather low data rates and WBAN temperature sensors may be small and inexpensive. *Blood pressure* is a principal vital sign that describes the systemic arterial pressure and is typically measured at a person's upper arm. Abnormal blood pressure can indicate the risk of cardiovascular disease, often the heart's output or the blood vessels' resistance. The involved data rates may be several hundred bits per second.

Blood oxygen saturation (SpO_2) is related to hemoglobin, which is the protein in the blood that can carry oxygen. SpO_2 describes the ratio of hemoglobin that is saturated with oxygen to unsaturated hemoglobin. Pulse oximetry is the method that allows determining SpO_2, for example by measuring the amount of light that is absorbed at translucent part of the body during the cardiac cycle. Pulse oximetry is part of standard monitoring during intensive care, but nowadays also used during mountain climbing or aviation to detect altitude sickness. The data rate is very low, because signal bandwidth is rather small. *Galvanic Skin Response (GSR)* sensors measure the electrical conductance of the skin, which is influenced by moisture level (sweat). Because sweat glands are controlled by the sympathetic nervous system, GSR sensors may be used determine psychological or physiological arousal. Significant changes of GSR may be expected on the order of hundreds of milliseconds, therefore data rates are low. *Blood glucose* monitoring is particularly important for people affected by diabetes mellitus. Typically access to blood is required, but new technologies based on ultrasound or dielectric spectroscopy could enable non-invasive long-term monitoring. Taking a sample every few minutes is more than sufficient, thus the data rates for measuring blood glucose are low.

Motion sensors are common in many WBAN applications. An accelerometer can be used to measure acceleration of body parts, which may be used to identify a movement pattern or body posture. Accelerometers are low-cost and have a small form-factor, therefore tri-axial accelerometers are often integrated in many WBAN platforms. Gyroscopes measure orientation and rotation and allow a more accurate recognition of movement within a 3D space. Finally, magnetometer may also be used to measure orientation relative to the earth surface. When (only) human mobility is involved the signal of motion sensors typically has a bandwidth of not much more than $20\,\mathrm{Hz}$. *Air quality sensing* with WBANs has lately attracted increased attention. Air pollution is a serious health problem in many urban areas, especially in developing

Figure 2.1: Functional decomposition of a WBAN node architecture.

countries. Traditionally, air quality is measured with stationary units, but recently integration of small air quality sensors into cell phones or WBAN nodes has been suggested by the research community [7, 55, 56]. Air quality sensors may, for example, measure the exposure to carbon monoxide (CO), nitrogen dioxide (NO_2), or ozone (O_3). A sampling interval of 5 seconds is sufficient [7], therefore the resulting data rates are rather low.

The typical sampling and data rates are summarized in Table 2.1. Note, however, that these values may differ depending on how that data are actually used by the application. For example, an ECG sensor must be sampled with 1000 Hz to measure heart rate variability only in certain risk situations (e.g. after heart transplant), but for healthy persons a sampling rate of 125 Hz is usually sufficient [57]. Similarly, if only the envelope of an EMG signal is required, then a sampling rate of 100 Hz is sufficient [58].

With respect to latency medical monitoring applications have rather diverse requirements: the acceptable time between collection and processing/visualization of data may span from 50 ms for capsule endoscopes [59] over 500 ms for ECG sampling [60] to several minutes for blood glucose monitoring [61]. For vital signs (temperature, blood pressure, pulse or respiratory rate) several tens of seconds up to minutes may suffice [46]. However, in general the quality of the medical care delivery is expected to improve as vital signs can be monitored in real-time [62, 63].

Table 2.1 allows to draw an interesting conclusion: there is a small number of sensors/applications that require very high sampling and data rates (ECG, EEG and EMG) while the other sensors produce data on the order of 100 byte or less per second. The latter is true, for example, for the collection of vital signs which is in general satisfied with low data rates [46]. The corresponding traffic rates are therefore low enough to be supported by current WBAN platforms as explained in the following section.

2.1.3 Hardware Platforms

WBAN platforms may be composed of COTS components or based on a custom design. The first approach has the advantage that development of the platforms is relatively fast and cheap. Custom designs take longer to develop but can be tailored to the application thus often have a more convenient form factor and lower power consumption. The majority of WBAN devices are composed of the five functional parts depicted in Fig. 2.1. The core component is a processing

element, which is powered by an energy source — typically a battery[1] — and interfaces to the sensors, for example, via an ADC. The processing element may be a low-power power MCU, such as the 16-bit TI MSP430 [65], that is capable of performing basic data processing at a few MHz of processing frequency. The processing element is also connected to a wireless communication device, the wireless transceiver. The transceiver (or "radio") is used to transmit or receive data in the WBAN. In certain WBAN applications the radio may not be fast enough to transmit all data while they are being collected. In this case some data may need to be stored on persistent storage, for example on the Flash memory of a micro Secure Digital (SD) card. Even if all data are transmitted over the radio, applications may keep a backup on the local storage element of a node.

When developing a WBAN hardware platform one design goal is to minimize its power consumption in order to maximize the lifetime of the application. The hardware modules on a WBAN platform consume different amounts of power. For example, Table 2.2 provides some information for the modules on the Shimmer2r [8] platform, which is a representative WBAN platform described in more detail below. It can be seen that the MCU typically consumes considerably less power than the radio chip and therefore efficient wireless communication is an important requirement in WBAN. Storing data on an external SD card may also consume a significant amount of power, possibly more than transmitting the data via the radio. The table also shows that the power consumption of WBAN sensors may differ by several orders of magnitude. However, the information provided in the table has to be interpreted in light of the fact that in practice most components will be duty-cycled. Therefore a component's startup time as well as the conversion time for sensor readings must be considered as well, when assessing the power consumption of a WBAN platform.

Various WBAN hardware platforms have been developed by the research community and recently the first WBAN devices have also become available commercially. Furthermore, in many academic studies multi-purpose sensor nodes, such as the popular Telos [66], have been used as WBAN platform. Except for the medical sensors a "multi-purpose" platform like Telos often consist of similar hardware components as COTS WBAN nodes and it often also allows to attach additional hardware. For example, the Telos node has a 10-pin expansion connector which has been used to connect an ECG board and other medical sensor boards [67]. In the following we will first describe the platforms that have been used to evaluate some of the aspects described in this thesis; afterwards we will provide an overview of some alternative WBAN platforms.

The majority of experiments performed within the scope of this thesis involved a popular and representative WBAN platform: Shimmer2 [8]. A small number of experiments were also conducted with the more generic WSN platform Telos [66]. The two platforms are in fact quite similar as they use the same MCU and radio. The MCU is a Texas Instruments MSP430F1611 [68], which belongs to the MSP430 family of MCUs. MSP430 MCUs are specifically designed for ultra-low-power applications and incorporate a 16-bit Reduced Instruction Set Computing (RISC) Central Processing Unit (CPU), peripherals and a clock system. The MSP430F1611 has 48 KB Flash, 10 KB of Random-Access Memory (RAM) and a flexible clock system sourced by an internal Digitally Controlled Oscillator (DCO) and/or two external

---

[1]Although lately the topic of energy harvesting has received an increasing interest in the research community [64], the vast majority of current WBAN platforms relies on batteries as an energy source.

Hardware component	Operating current (mA)	Sleep current (mA)	MCU interface
MCU (MSP430F1611) @ 8 MHz	5	0.02	-
IEEE 802.15.4 Radio (CC2420)	20	< 0.03	SPI
Bluetooth Radio (RN-42)	45	−	SPI
SD card	47	0.15	SPI
Magnetometer	0.8	< 0.05	ADC
Accelerometer	0.5	0.003	ADC
ECG (3-lead)	0.18	−	ADC
EMG (2-lead)	0.18	−	ADC
GPS	< 48	0.015	I2C
GSR	0.06	−	ADC
Barometric pressure	0.7	0.0001	I2C

Table 2.2: Current consumption of Shimmer2r hardware components at regulated voltage of 3 V, source: Shimmer [8] datasheets.

oscillators. It also contains a 12-bit ADC, two independent timers and two Universal Asynchronous Receiver/Transmitters (USARTs). Both, Telos and Shimmer2, are equipped with the IEEE 802.15.4-compliant Texas Instruments CC2420 wireless transceiver [69]. The CC2420 operates in the 2.4 GHz Industrial Scientific and Medical (ISM) band, it uses Offset Quadrature Phase-Shift Keying (O-QPSK) modulation and has a data rate of 250 kbps. A detailed overview of the IEEE 802.15.4 standard [18] is provided below in Section 2.3.2. The Telos platform is available in different variants. The variant used for this thesis (*Tmote Sky*) allows to run the CPU clock at 8 MHz. It includes a temperature, light and humidity sensor and has an additional external Flash memory of 1 MB capacity. Like all Telos platforms the Tmote Sky has a MSP430 MCU, CC2420 radio and is powered by a pair of AA batteries.

Like the Telos platform, *Shimmer2* integrates the Texas Instruments MSP430 MCU and the CC2420 transceiver [69]. However, its form factor, enclosure and the availability of various medical sensorboards make the Shimmer platform a more suitable WBAN platform. The Shimmer2 platform also incorporates a Bluetooth radio, which can be used as an alternative to the IEEE 802.15.4 radio. In addition, each Shimmer2 node has an integrated 3-axis accelerometer (Freescale MMA7260Q [70]), which can, for example, be used to detect posture changes of the person monitored by the WBAN. Shimmer also provides a micro SD card slot and in our experiments each Shimmer2 nodes is always equipped with a 2 GB micro SD card, which is sufficient to store any traces that we collect during any of our measurements. The additional sensorboards that can be plugged into the Shimmer2 platform are ECG, EMG, GSR, magnetometer or GPS boards. The Shimmer2r is powered by a 450 mAh rechargeable Li-ion battery.

The research community has developed several other WBAN hardware platforms. Like the Telos platform the so-called BSN (Body Sensor Network) platform, designed by Lo and Yang [71], uses an MSP430 MCU and a CC2420 Radio. It can be extended by ECG and SpO_2 sensorboards and has an on-board Flash memory with a capacity of 512 KB. The Eco node

developed by Park et al. is a wearable ECG platform [72], which is composed of sensors that do not require direct contact to the skin. The platform is based on a Nordic VLSI nRF24E1 chip which integrates a 2.4 GHz RF transceiver core with a data rate of 250 kbps and a 8-bit 8051-compatible CPU. The MITes (MIT Environmental Sensors) platform developed a few years earlier at Massachusetts Institute of Technology (MIT) uses the same nRF24E1 chip [73]. MITes includes an accelerometer and is designed to monitor people's interaction with objects in the environment. The Berkeley Tricorder [74] is a WBAN platform that is intended for ambulatory health monitoring. It allows to measure ECG, EMG, SpO_2, Respiration, and motion. Like the Shimmer2 platform the Berkeley Tricorder is equipped with an MSP430 MCU and can store data on an micro SD card. The Berkeley Tricorder transmits data via a Bluetooth radio. The Epic family designed by Dutta et al. allows to decompose sensor network platforms into "modules", which consist of one or more packaged Integrated Circuits (ICs), and "carriers" that connect modules with application-specific sensors [75]. Epic has been used to create badges that allow to people to collect and analyze environmental data with the help of air quality sensors; Epic has also been used to create a platform that can be used for motion detection. The Epic core module is built on a MSP430 MCU and CC2420 radio. The examples given so far indicates that WBAN research platforms typically include either a IEEE 802.15.4 or Bluetooth radio and a low-power MCU, and that the combination MSP430 MCU and CC2420 radio seems to be wide-spread. An more thorough overview of platforms is provided in [16, Chapter 6] and [15, Appendix A].

Some of the described platforms (for example, Shimmer2, Telos or Epic) are commercially available, but primarily target the research community. There exist a few commercial wireless systems that are intended for operation in hospital: the Welch Allyn Micropaq device [76] integrates wireless ECG or SpO_2 sensors that communicate via WLAN with a central monitoring station. The system logs data and provides an audible alarm when it detects a deterioration in patient or equipment conditions (such as low battery status). The company also provides an ECG monitor that uses Bluetooth communication. Corventis is a company that has developed a wireless device to monitor heart rate (variability), respiration, posture and ECG [77]. The device is small and can be attached to the body like a plaster. Data is transmitted to a central station using a proprietary communication protocol. In the domain of sports and personal fitness lately body-worn pulse and heart rate wireless monitoring devices have become available commercially. These devices may communicate via with a user's mobile phone or wrist-worn watch via Bluetooth [78] or the ANT+ protocol designed for fitness applications [79]. While "classical" Bluetooth may often consume a significant amount of energy (c.f. Table 2.2), BLE is a promising new radio technology whose power consumption is comparable to IEEE 802.15.4 (see Sect. 2.3.2). Unfortunately, so far there exist no WBAN platforms that use BLE communication.

A comparison of WBAN sensor sampling rates and the capabilities of current WBAN platforms allows to draw the conclusion that the data rates of current low-power radios (around 250 kbps with IEEE 802.15.4 and Bluetooth low energy) are suitable to support the majority of sensors (at least 10 out of 13 sensors shown in Table 2.1). In these applications a node is expected to transmit on average no more 150 byte of data per second (c.f. blood pressure monitoring in Table 2.1). Applications that involve high-frequency sampling (ECG, EEG and EMG) require faster radios — in this case "classical" Bluetooth may be an option as it provides a data rate of 1 Mbps (which can in Bluetooth 3.0 + HS (High Speed) even be extended up to 24 Mbps). In

(a) Star topology.   (b) Two-hop extended star topology.   (c) Multihop topology.
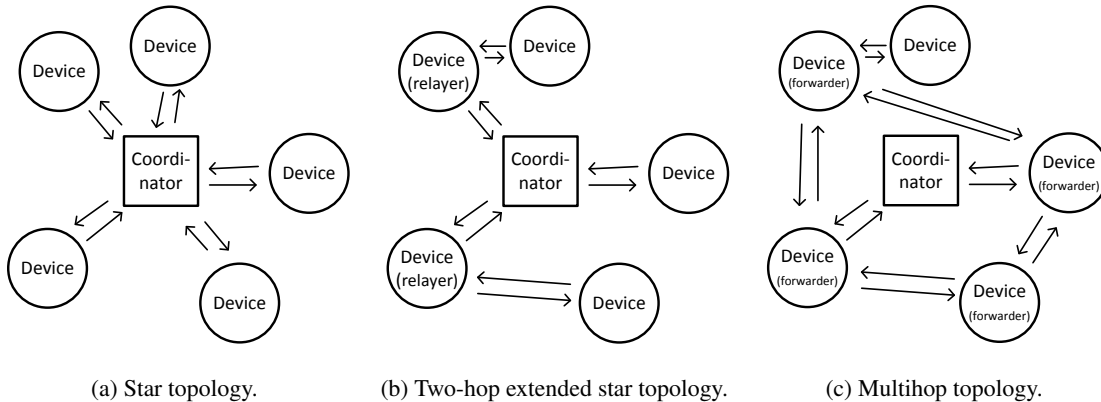
Figure 2.2: WBAN network topologies.

this thesis, however, we concentrate on applications with low data rate of no more than 150 byte per second. For these applications the data rates provided by IEEE 802.15.4 or Bluetooth low energy are sufficient.

2.1.4 Topologies and Protocol Architectures

WBAN communication can be classified into on-body and off-body communication. On-body communication refers to communication among nodes attached to the (same) human body and is also called intra-WBAN communication. When a body-worn node exchanges data with a communication peer in the environment, for example, a stationary gateway, this is referred to as off-body (or extra-WBAN) communication. Communication among two WBAN nodes attached to different persons is a special case of off-body communication sometimes referred to as body-to-body communication. When a WBAN enables both, off-body and on-body communication, different communication technologies may be involved, for example, intra-WBAN communication may be based on IEEE 802.15.4 while off-body communication may involve cellular networks [80]. In the following we concentrate on on-body communication, because it is the type of communication that we focus on the remaining chapters of this thesis.

WBAN on-body communication is often realized as a logical star: among the N nodes in the network there exists one special node referred to as "coordinator" (or "hub" or "controller"), with whom the remaining N-1 devices may communicate. Direct communication among the remaining N-1 devices does not occur, even though it may physically be possible. Sometimes a star topology is augmented with relayer nodes, which are capable of forwarding a packet to the coordinator on behalf of the relayed node. When there is a maximum of one relayer involved between a device and a coordinator then the topology is sometimes called a "two-hop extended star topology" (IEEE 802.15.6 standard [81]). When two or more intermediate nodes are chained such that a packet may travel more than two hops to reach the destination the network forms a multihop topology.[2] A multihop network may enable peer-to-peer communication among the

---

[2]The distinction between a star topology extended by relaying capabilities and a multihop network is fuzzy — conceptually the first is realized in the link layer and the latter in the network layer.

Architecture	On-body communication	Off-body communication	Body-to-body communication?
MobiHealth [80]	IEEE 802.15.4 or BT	3G	no
CodeBlue [85]	IEEE 802.15.4	IEEE 802.15.4 + MultiHop	yes
WHMS [83]	IEEE 802.15.4 or BT	3G or WLAN	no
SMART [44]	-	IEEE 802.11	no
ALARM-NET [86]	IEEE 802.15.4	IEEE 802.15.4 + MultiHop	no
MEDiSN [45]	-	IEEE 802.15.4 + CTP	no

Table 2.3: WBAN architectures (research prototypes).

devices. The described topologies are visualized in Figure 2.2.

The selection of the network topology has implications for the protocol architecture. For example, in a network that realizes a star topology the physical and data link layer are required, but a routing protocol does not have to be present. This is also true for a two-hop extended star topology using relayers, because the relaying functionality is part of the link layer [81]. While there seems to be an agreement that (extended) star topology is a "common choice" [82] in WBANs (e.g. [83, 84]) there currently exists no well-established protocol architecture for WBANs. The lower two layers of the Open Systems Interconnection (OSI) reference model (Data Link and Physical layer) are typically realized by one of three IEEE 802.15 standards. IEEE 802 describes a set of communication standards for personal, local and metropolitan area networks, which all cover the Physical Layer (PHY) and MAC layer. The IEEE 802.15 standards target Wireless Personal Area Networks (WPAN), and the most relevant standards for WBAN communication are IEEE 802.15.1 (Bluetooth), IEEE 802.15.4 and IEEE 802.15.6. Section 2.3.2 provides an overview of these standards.

Above the link layer a common protocol architecture for WBANs is not yet well-established. There have been proposals for dedicated WBAN network layer protocols [87] or the adoption of existing WSN protocols such as ZigBee [88]. The protocol architecture may also involve middleware that provides additional services to the application, e.g. definition and serialization of data structures as well as providing services that allow the application to communicate using the published/subscribe interaction pattern [85]. Several research projects such as MobiHealth [80], CodeBlue [85], Wearable Health Monitoring Systems (WHMS) [83], ALARM-NET [86] or MEDiSN [45] have developed their own WBAN protocol architectures, which are typically adapted to a particular usage scenarios (a comparison is shown in Table 2.3).

In the WSN domain during the past years there has been a growing trend towards enabling Internet Protocol (IP)-connectivity [89], which allows to establish interoperability with a large number of other devices and to leverage many existing tools and protocols. The author of this thesis expects that the same trend will catch on in the WBAN domain. To this end WBANs may implement the IPv6 over Low power WPAN (6LoWPAN) protocol [90], which is an adaptation layer that can be used to transmit Internet Protocol version 6 (IPv6) packets over 802.15.4 links. One possible alternative for the upper layers is the Internet Engineering Task Force (IETF) Constrained Application Protocol (CoAP) [91] on top of the of User Datagram Protocol (UDP).

Application	*Application Layer*
CoAP	
UDP	*Transport Layer*
IPv6 (6LoWPAN)	*Network Layer*
IEEE 802.15.4 MAC	*Data Link Layer*
IEEE 802.15.4 PHY	*Physical Layer*

Figure 2.3: An IP-based WBAN protocol architecture.

CoAP implements an interaction model similar to the client/server model of Hyper Text Transfer Protocol (HTTP), but it interchanges messages asynchronously over unreliable UDP rather than Transmission Control Protocol (TCP). Its low overhead makes it very useful for Machine to Machine (M2M) communication. For example, Jacobsen et al. have proposed such a CoAP- and IP-based protocol architecture for the WBAN domain [92]. Figure 2.3 visualizes such a generic protocol architecture.

2.2 Wireless On-Body Communication

The previous Section 2.1.4 introduced the concepts of on-body, off-body and body-to-body communication. In the following we provide a more detailed description of the wireless on-body channel. We first discuss the frequency bands available for WBAN communication, then we provide a brief revision of the wireless propagation mechanisms and finally we focus on the particular characteristics of the wireless on-body channel and corresponding channel models.

There exist several books on wireless communication, the following overview in Section 2.2.2 is based on the book by Rappaport [93], whenever no other reference is given. Therefore, for detailed information please refer to the original book. Note also that in the following we restrict ourselves to an explanation of the propagation channel, which means that impact of antennas or RF circuits are ignored.

2.2.1 Frequency Bands

The electromagnetic spectrum is a limited resource, at least the portion that is practically envisioned for WBAN communication (up to 60 GHz [94, 95]). The national regulation agencies (in Germany: "Bundesnetzagentur", in US: U.S. Federal Communications Commission (FCC).) manage the frequency allocation and spectrum usage. The vast majority of the radio spectrum is reserved, e.g. for military or public safety, or sold/leased to a specific user. Only a small portion is unlicensed, also called license-free, spectrum. For the unlicensed spectrum, users do not have to apply individually, instead anyone can use it under certain conditions related to, for exam-

ple, maximum output power, use of spread spectrum techniques or duty cycles.[3] Because they are often shared by many users and technologies the main disadvantage of license-free bands is that there may be RF interference deteriorating the communication quality. An exclusive user of a licensed bands does not have (or can control) this problem with proper frequency planning. However, the exclusive acquisition of licensed spectrum can be quite costly and time-consuming.

On the other hand, there exist licensed bands that are available for certain users or applications, sometimes for only a small registration fee. For example, the FCC has recently opened the $2360 - 2400$ MHz band for medical WBANs, which are defined as "medical body-worn devices all of which transmit or receive non-voice data or related device control commands for the purpose of measuring and recording physiological parameters and other patient information or performing diagnostic or therapeutic functions via radiated bi- or unidirectional electromagnetic signals" [96]. However, as this spectrum was previously assigned to the aeronautics industry, medical WBANs may only use it on a secondary basis. This means they not only have to accept RF interference from the primary users (commercial test pilots), but are also responsible to not cause any interference to the primary. As this example shows, some licensed bands are therefore not exclusive; but since the pool of users is typically smaller than in unlicensed bands, these bands are still quite attractive. One may therefore classify RF bands in:

1. **licensed**: (a) the band can be used exclusively by a certain user who has purchased/leased it from the national agency; (b) the band may be used by a certain group of users under certain conditions after authorization (e.g. registration for a nominal fee) at the national agency.

2. **unlicensed**: the band can be used by anyone with certified equipment and under certain conditions (but without prior registration)

A frequency band that is licensed in one country may be unlicensed in another and therefore practically only a small number of unlicensed bands are available world-wide. For mobile networks, such as WBAN, that may be "traveling" through different regulatory domains this may pose a problem.

Licensed Bands

Table 2.4 lists some licensed frequency bands that may be used by medical WBAN applications. The Medical Implant Communication Service (MICS) allows communication with electronic implants, for example, pacemakers. The amount of radiated energy is very low (Equivalent Isotropically Radiated Power (EIRP) of 25 μW) and the bandwidth of a channel is rather small (the maximum usable bandwidth is 300 kHz). The majority of WBAN applications described in Section 2.1.1 use non-invasive sensors, and will thus not be able to communicate in MICS bands. The MICS band are the core of the Medical Device Radiocommunications Service (MedRadio) bands, which specifies further "wing" bands as additional adjacent spectrum (401-402 MHz and 405-406 MHz).

---

[3]Unlicensed bands are not unregulated: the transceivers must have been certified by the appropriate regulatory authorities.

Band	Frequency range (MHz)[4]	Availability	Application example
MICS	402-405	world-wide	medical implants
WMTS	608-614, 1395-1400, 1427-1432	US	respiration monitoring
MBAN	2360 - 2400	US	ECG monitoring

Table 2.4: Licensed bands for medical applications.

The Wireless Medical Telemetry Service (WMTS) bands are reserved for medical telemetry devices. The bands were allocated in the year 2000 by the FCC in response to a more intensive use of certain television channels, which had until then been used by medical wireless telemetry on a secondary basis. The WMTS bands are divided in three separate blocks and available only in the US. Some of these frequencies are shared with other users, and on some frequencies the WBAN has to operate on secondary basis (which may even change from city to city). The allocation of WMTS bands has been criticized, because the total bandwidth is rather small (13 MHz) and not contiguous, and bands partially suffer from (legal) adjacent channel interference. In practice, rather than using WMTS devices, several hospitals have been using IEEE 802.11 networks for healthcare applications, some people claiming these networks already meet the requirements for supporting life-critical applications [97].

In late 2012 the FCC allocated a new band (Medical Body Area Network (MBAN)) that allows medical WBANs (as defined above) to operate in the 2360-2400 MHz band. Most of the spectrum, however, overlaps with spectrum allocated for the aeronautics industry and MBANs have to operate on secondary basis. Healthcare providers will have to register their equipment at FCC (this process is envisioned to start around mid-2013). Currently the MBAN band is only available in the US, but the European Telecommunications Standards Institute (ETSI)) is considering to suggest adopting it in Europe [98]. However, there exist certain disadvantages of using MBAN bands for WBAN application, which we discuss below, after introducing the unlicensed bands.

Unlicensed Bands

Several unlicensed bands may be used for WBAN communication. An overview of the most common bands is provided in Table 2.5. Although they are available world-wide the 13,56 MHz, 27 MHz and 40.6 MHz ISM bands will likely not be attractive for most WBAN applications due to their limited available bandwidth. The 433 MHz ISM is available world-wide, has sufficient bandwidth to achieve data rates of up to 200 kbps [99] and good signal propagation characteristics due the comparably low frequency. For example, free-space path loss at 433 MHz is about 15 dB smaller than at 2.4 GHz (Eq.2.1). On the other hand, the 433 MHz band may require large antennas (a 1/4 wavelength monopole antenna would be 16 cm long); furthermore, the band may already be quite crowded in some areas due to recent standardization efforts such as DASH7 [99]. A successful pilot study with 433 MHz transceivers attached to human body have

---

[4]Note that for bands that are available "world-wide", the frequency range may differ by country and the table shows the German frequency range.

Frequency range (MHz)[4]	Availability	Technology example
13.553−13.567	world-wide	Radio-Frequency Identification (RFID)
26.957−27.283	world-wide	babyphones
40.66−40.70	world-wide	wireless remote controls
433.05−434.79	world-wide	DASH7 [99] devices
868−870[5]	EU	wireless thermometer
902−928	US	remote metering
2400−2483.5	world-wide	WLAN, IEEE 802.15.4
5725−5875	world-wide	WLAN
6000−8500[6]	world-wide	UWB

Table 2.5: Unlicensed communication bands.

been reported by Kim et al. [100].

The 915 MHz and 868 MHz ISM band offer several megahertz of bandwidth and the antenna sizes are acceptable for many WBAN applications. Low-power transceiver modules are available commercially and signal propagation characteristics are quite good. A drawback of the 868 MHz band is its fragmentation into several small sub-bands, which have different restrictions regarding transmit duty cycle or operating modes. The main disadvantage, however, is that the bands are not available world-wide: the 868 MHz band is not unlicensed in the US, while 915 MHz may not be used without license in the Europe. For WBAN applications that target a specific location (e.g. hospital), this is unproblematic; however, applications with unrestricted user mobility would be unfeasible in practice.

The 2.4 GHz ISM band is available world-wide and offers a contiguous bandwidth of more than 80 MHz. Antennas can be very small and the regulatory restrictions with respect to output power are more relaxed than in the sub-GHz bands. However, body shadowing may become a problem, because in the 2.4 GHz band path-loss is more pronounced than in the 868/915 MHz bands (c.f. also Section 2.2.3). Furthermore, the 2.4 GHz band is shared with several other wireless technologies, notably WLAN (IEEE 802.11b/g/n) and consequently RF interference can become a severe problem. Yet, the majority of current WBAN systems use this band and the emergence of several low-power 2.4 GHz RF communication standards (Section 2.3.2) indicates that this band will continue to remain one of the prevalent frequency bands for WBAN communication.

In comparison to the 2.4 GHz band the 5 GHz band is (currently) less likely congested. However, path loss is significantly higher and due to smaller wavelengths signals do not penetrate well through solid objects, including the human body. Therefore the 5 Ghz is mainly used by WLAN (IEEE 802.11a/n), but has hardly been considered for WBAN communication: the low transmission power of WBAN transceivers can often not compensate for the high path loss and at higher frequency the creeping wave effect (Section 2.2.3) is less pronounced. Finally,

---

[5]The 868 MHz band is not contiguous, but segmented into several small chunks.

[6]In Germany the maximum power can be emitted by Ultra-Wideband (UWB) devices between 6 GHz and 8.5 GHz, but more (in fact the entire) spectrum can be used if power is sufficiently low [101].

UWB has been considered a promising technology for WBANs. UWB transceivers spread the signal over a very large bandwidth, $> 500\,\mathrm{MHz}$ or 20% of the arithmetic center frequency (FCC definition). This results in a tiny rise in the noise floor of narrowband transceivers, i.e. the interference for the latter should practically be negligible. UWB systems may, for example, use short burst of energy shifted in time to convey information. UWB is capable of achieving high data rates (several hundred Mbps) and is robust against multipath fading, because it covers a large bandwidth. In 2002 the FCC opened the 3.1-10.6 GHz spectrum for UWB technologies under certain conditions such as maximum power spectral density of -41.3 dBm/MHz. Europe adopted similar rules a few years later. In practice, however, radio chips have data rates far below the theoretical estimations, which caused many UWB vendors to stop operation. Furthermore, progress in the standardization has been slow, e.g. the IEEE 802.15.3a task group addressing UWB communication in WBANs was dissolved.[7] However, if low-power UWB chips are becoming available commercially, UWB may become an interesting technology for WBAN applications. A thorough overview of UWB for WBANs is provided in [16, Chapter 16].

Discussion

Licensed frequency bands are favorable, because the amount of RF interference is typically greatly reduced. Of the three licensed bands introduced above, the MICS and WMTS bands will likely only satisfy a small number of WBAN applications described in Section 2.1.1. However, the MBAN band seems attractive, because it provides substantial contiguous bandwidth. Also standards from the 2.4 GHz ISM band, such as IEEE 802.15.4, can be adapted easily by re-tuning radios, which is expected to be rather cost effective [98]. On the other hand, there exist several reasons, why the MBAN band may not be suitable for several WBAN applications: first, not all WBAN applications are used for "diagnostic or therapeutic functions", but several WBAN applications are envisioned for non-medical purposes (c.f. Section 2.1.1); these WBAN applications cannot use the MBAN bands. Second, the operation in the 2360-2390 MHz range are restricted to indoor operation in health care facilities only. This will limit the number of WBAN applications considerably. Third, operation on secondary basis is mandatory, which means communication will always include overhead for detecting primaries; furthermore, when a primary is present communication may have to be disabled completely. Finally, the network has to be implemented as a star topology and off-body (including body-to-body) communication is not allowed. Instead, a separate telecommunications system must be used to transmit data from the body to an external entity. This requirement means that most of the envisioned network architectures introduced in Section 2.1.4 are unsuitable.

It is therefore rather likely that in the foreseeable future a significant number of WBAN applications will continue to use unlicensed bands. UWB may become an important technology for WBAN applications; at the moment, however, this remains unclear, because interest from industry has decreased significantly over the last few years. Therefore, the 2.4 GHz ISM band

---

[7]The IEEE 802.15.4a standard also covers an impulse UWB PHY, but it is intended mainly for ranging application and thus unsuitable for WBAN applications. Furthermore, Wireless Universal Serial Bus (USB) (Universal Serial Bus) is based on UWB, but the power consumption would be to high for most WBAN applications.

will likely continue to play an important role, due to its world-wide availability, large bandwidth and acceptable on-body propagation characteristics. On the other hand, as explained in Chap. 6 RF interference is a serious problem for 2.4 GHz WBAN devices. An important question that is addressed by this thesis in Section 7.2 and Section 7.3 is therefore on mechanisms that mitigate the effects of 2.4 GHz RF interference in WBANs.

2.2.2 Wireless Communication Basics

Typically, the propagation of the electromagnetic wave propagation is modelled on two levels of abstraction: large-scale fading (path loss) models describe how *average* received signal strength decays with distance. And small-scale fading models describe the fast signal strength variations over small distances or time intervals. Note that in this thesis we use the term *fading* to describe the aggregate effect of both, large- and small-scale fading. The effect that signal strength decays logarithmically with distance is called large-scale fading or *path loss*. When there exists a line-of-sight path in free space with no obstacles in the environment that could affect signal propagation, then the *free space propagation model* may be used to predict signal strength between a transmitter and receiver separated by *d* meters:

$$P_r(d) = P_t \left(\frac{\lambda}{4\pi d} \right)^2 \tag{2.1}$$

where $P_r(d)$ is the received signal power at distance d, P_t is the transmitted power and λ is the wavelength in meters. The underlying assumption is that energy is emitted isotropically; thus, at a given distance d the energy will be uniformly distributed over the surface of a sphere with radius d (whose area is $4\pi d^2$). In practice, however, path loss will be influenced by objects in the propagation environment, which may attenuate signal strength. In this case the *average* received signal strength is often modeled as

$$P_r(d)[dBm] = P_t[dBm] - PL(d_0)[dB] - 10\,\alpha\,log\left(\frac{d}{d_0} \right) \tag{2.2}$$

where P_t is the transmitted power, $PL(d_0)$ is the average path loss, measured at a close reference distance d_0, α is the path loss exponent, which defines how fast path loss increases and d is the transmitter-receiver separation. Typically, α is derived by curve fitting a set of field measurement results; alternatively, an α value may be adopted if it was measured in a similar environment. For example, α values between 3 and 5 are common for rural to urban environments (with higher values are for urban environments) [102, Chapter 3]. Even in the same propagation environment path loss for a given transmitter-receiver distance may vary as it depends on the particular type, amount and configuration of intermediate/surrounding objects. Eq. 2.2 does not take this into consideration, instead it always predicts the same path loss for a fixed transmitter-receiver separation. Therefore, Eq. 2.2 is often extended by a normally distributed (normal in dB) random variable, that represents random *shadowing* effects:

$$P_r(d)[dBm] = P_t[dBm] - PL(d_0)[dB] - 10\,\alpha\,log\left(\frac{d}{d_0} \right) + X_\sigma \tag{2.3}$$

where X_σ is log-normally distributed random variable with standard deviation σ. For example, in mobile communications σ typically takes values between $5\,\text{dB}$ to $12\,\text{dB}$ [103, Chapter 1]. Eq. 2.3 has been considered as the basis for modelling the large-scale fading effects in on-body communication, which we will discuss in Section 2.2.4. In addition, there exist various more specific large-scale propagation models for outdoor/indoor and urban/rural propagation environments based on a certain terrain or building structure; details can be found in [93, Chapter 3].

In addition to path loss, which is responsible for average signal strength variations over larger distances, signal strength may undergo fast variations over small distances or time intervals. These characteristics are described by fading models. The basis for fading models are three physical phenomena: electromagnetic waves travelling from a transmitter to a receiver may be reflected from various objects in the environment. Furthermore, when a wave encounters an obstacle it may bend around the object creating secondary waves which may, for example, reach a receiver on a Non-Line-Of-Sight (NLOS) path. Finally, the wave may hit a rough surface or object which is small compared to the wavelength and this interaction may scatter the wave and diffuse its energy in various directions.

These three propagation mechanisms - reflection, diffraction and scattering - may cause multiple versions of the signal to arrive at the receiver at slightly different time offsets as they may take different paths. One consequence is that the signal copies may interfere with each other and as a result the magnitude of the received signal can vary rather drastically. This phenomenon is called multipath fading, small-scale fading or simply *fading*. In a multipath fading channel, many signal copies may reach the receiver and each copy of the signal taking a path i with distance d_i will yield a certain delay $\frac{d_i}{c}$ (with c being the speed of light), which corresponds to a specific phase shift and attenuation. The fluctuation of the delays of all signal copies is described by the channel-specific *delay spread*, which may range from a few tens of nanoseconds (indoor) to tens of microseconds (outdoor). At the receiver the set of multipath components combine through superposition (linearly) into the received signal. This process may, for example, involve destructive interference: when two frequency components have a phase difference of π they interfere destructively and may even cancel each other out in case they have similar magnitude. Therefore, the specific physical arrangement and environment has large impact on the received signal strength, because it determines the delay values, phase offsets and thus amount of interference among the multipath components. If the transmitter, receiver and all objects in the environment are static the channel is *time-invariant*, implying that received signal does not change over time. Only as the transmitter, receiver or objects in the environment start to move, the channel becomes *time-variant*: then the phase and magnitude of the multipath components may change over short time scales and the received signal strength will show fast variation over time. In addition, if the distance between transmitter and receiver changes - for example, when a mobile receiver passes by a fixed transmitter - the velocity contributes to a change in the carrier frequency (*Doppler shift*). Therefore the received spectrum will have components in a larger spectrum than the transmitted signal (the amount of spectral broadening depends on the wavelength and speed is described by the *Doppler spread*). Then, with a rate of change proportional to the velocity of transmitter or receiver, multiple signal copies arrive each having an individual delay, magnitude, phase and Doppler shift.

First Order Statistics

If one had detailed knowledge of the path lengths of the signal copies one might attempt to compute the impact of attenuation and phase shift for the individual multipath components; the superposition (i.e. sum) of these components could then be used to determine the received signal.[8] Practically, this is often not possible. Instead, the behavior of the received signal is often characterized through the statistical description of a fading model. For example, the Rayleigh fading model is based on the assumption that there exist multiple paths and none of them is dominant (no Line-Of-Sight (LOS) path). Then one may assume that the phase offsets of the signal copies are independent and uniformly distributed (between 0 and 2π). When a large number of signal copies arrive over different paths the central limit theorem can be applied and it can be shown that the received signal strength follows a Rayleigh distribution. When there is a dominant path (LOS), then the Ricean fading model is more applicable. As described in Section 2.2.4 both, Rayleigh and Ricean, have been used to model fading in WBANs. According to these models fading can result in rather drastic deep fades. For example, a practical Rayleigh fading margin of 20-30 dB should be assumed [104]. This is particular problematic if no movement is present, because then a longer-lasting deep-fade can occur.

Second Order Statistics

The amount of time over which the channel is assumed to be stable is called *coherence time*. Different definitions for this metric exist, for example coherence time may represent the time over which the amplitudes of two signals have a correlation above 0.5 [93, Chapter 4]. Coherence time is inversely proportional to Doppler spread (both are are based on the speed of the receiver / transmitter). Therefore practical approximations, such as $T_C \approx \frac{9}{16\pi f_m}$, have been derived to determine the coherence time T_C given the maximum amount of Doppler shift f_m [93, Chapter 4]. The maximum Doppler shift can be derived from speed v and wavelength λ via $f = \frac{v}{\lambda}$. Thus, for example, at 2.4 GHz transmission frequency and human walking speed of 1 $\frac{m}{s}$ the mentioned approximation results in a coherence time of 22 ms. As reported by several empirical WBAN studies a value of 20 ms is indeed a realistic bound for the coherence time in WBANs scenarios when the subject is in motion (c.f. [105, 106]). To better characterize how the signal behaves in time second order statistics may be used. Two popular techniques, also applied in the context of WBANs, are Level-Crossing Rate (LCR) and Average Duration of Fades (ADF). LCR is the average rate at which the signal crosses a certain signal level threshold. And ADF represents the average duration that the signal level lies below a certain threshold. Analytical expressions have been derived that describe LCR and ADF for some of the well-known fading distributions, such as Rayleigh fading [93, Chapter 4].

Noise and Interference

Path-loss, shadowing and fading are multiplicative effects: the transmitted signal is multiplied with the superposition of the effects described above. At this level of abstraction the propagation

---

[8]Ray tracing techniques built on this approach, however, such techniques require detailed knowledge of the environment and are very computation-intensive.

Figure 2.4: Signal strength measurements for a subject performing a sit-stand-sit movement.

channel is typically symmetric, i.e. the signal will undergo the same effects in both directions. In addition to these multiplicative effects, the signal may be affected by additive effects, mainly noise and interference. These effects make the channel asymmetric, as they may have different characteristics at transmitter and receiver. Thermal motion of particles is present in all matter and inside the components of the receiver circuitry it creates *noise* which distorts the signal. Such noise is often modelled as spectrally flat, which is then referred to as *white noise*. In addition to noise induced by the receiver components the signal may be affected by *interference*. Interference describes unwanted signals in the same portion of the spectrum as the carrier of the intended signal. Interference may be caused by electrical phenomena in the atmosphere (e.g. lightning) or be of cosmic origin (e.g. the sun). However, the occurrence of these natural interference sources is normally much less troublesome than artificial or so-called *man-made interference*. The latter stems from electrical equipment in the proximity of the receiver. The source of man-made interference may be electrical motors or digital electronic equipment that is badly shielded. But the most frequent interference source are typically the signals of other transmitters.

All effects discussed so far – path-loss, shadowing and fading as well as noise and interference – impact the quality of a communication link. The SINR model introduced in the previous Chap. 1 combines these effects relating received signal power to the noise and interference effects: the nominator in Eq. 1.1 represents the signal after it has been affected by the multiplicative effects and the denominator is the sum of additive components, i.e. noise and interference. The SINR is an important metric, because it determines the probability of correctly decoding a symbol during the process of converting the analog signal to a digital sequence of symbols (demodulation). Analytical expressions exist that describe for several well-known modulation schemes the relationship between SINR and symbol error probability, typically under the assumption of white noise with a constant spectral density. For details refer to Proakis et al. [107, Chapter 4].

27

2.2.3 Channel Characteristics

The principles of wireless communication described in the previous Section 2.2.2 are naturally also valid in the context of WBANs, after all the underlying physical principles do not change. Yet the channel characteristics of WBANs are quite unique, due to the following reasons:

- In WBANs the inter-node distances are small, but the relative change in distance may be large. Furthermore, in WBANs movement of the human body may result in frequent alternation between LOS and NLOS communication, and thus introduce rather dynamic shadowing. This can translate into large path-loss variation as can be seen in Fig. 2.4: this figure shows a 15-second signal strength trace measured by WBAN nodes in an office environment. The nodes were attached to chest, left hand and right foot (receivers) while the sender was located in the right trouser pocket (Shimmer2r platform using 0 dBm output power[9]). During the 15-second interval the subject was first sitting still on a chair, then standing up and then sitting down again. It can be seen that the path loss can vary by up to 20 dB among the different postures and that during the transition phases (standing up / sitting down) the variation is very agile. This is in line with the theoretical explanation given above in Section 2.2.2: for example, consider two WBAN nodes with an inter-node distance of 20 cm and assume the log-distance path loss model Eq. 2.2 with a path loss exponent of $\alpha = 3$. If, through a change in body posture the distance extends to 50 cm, the increase in path loss will be about 12 dB.[10] To make an analogy: in the world of cellular networks such an increase in path loss is comparable to a mobile located at 20 m from a basestation and moving to a distance of 50 m (under the same model assumptions).

- The distance between the antenna and the body surface has a rather large impact on the path loss of an on-body communication link. For example, Roelens et al. have shown for half-wavelength dipole antennas that an antenna height between 0.5 to 5.0 cm may result in path-loss variation up to 20 dB [110]. When the body is in motion (and especially when the nodes are not tightly fixed to the body surface) this effect can result in a rather dynamic path loss. Furthermore, in practice antenna radiation patterns are not perfectly omnidirectional, thus a change in the angle between sender and receiver antenna (introduced by turning a limb) may cause fluctuation in signal strength.

- Signal reflections and scattering from body parts and nearby objects in the environment contribute to multipath fading. Due to human mobility the (propagation) environment may be quite dynamic. Reflections from the ground often have a significant contribution to fading [111]. The amount of energy being reflected is determined by the material specific reflection coefficient [112, Chapter 8]. As the person moves the material of the ground (as well as other objects in the environment) changes and thus the amount of signal reflection can vary.

---

[9]The center frequency was set at a non-standard value of 2483 MHz, which is also sufficiently spaced apart from WLAN/Bluetooth and thus rather interference-free

[10]A path loss exponent of $\alpha = 3.3$ between torso and arm has been reported by Reusens et al. [108]. However, this value may be optimistic: the IEEE 802.15.6 channel modeling subcommittee has adopted a path loss model based on $\alpha = 6.6$ [109]. With $\alpha = 6.6$ the example would result in an increase in path loss of 26 dB.

- In WBANs the diffraction components is often significant: finite-difference time-domain simulation has shown that waves that are diffracted around the human body can represent a rather strong signal component [113]. Such a diffracted signal component is commonly also referred to as *creeping wave* or *surface wave*. Creeping waves do not imply that more energy is present at the boundary; rather diffraction on the curved boundary guides the wave such that it can extend beyond edges or curves [114]. The magnitude of the diffraction component depends on the wavelength and characteristics of the diffracting object. In WBANs creeping waves may self-interfere, for example when waves travel clockwise and counter-clockwise around the torso [113].

- Finally, one consequence of human mobility is that WBANs may be exposed to a range of different RF interference environments. Especially in the license-free ISM bands this is a serious problem. For example, as shown in this thesis, co-located WLANs can lead to drastic deterioration of 2.4 GHz WBAN communication performance. Since WBAN are mobile but many interferers (such as WLAN Access Points (APs)) are often static the quality of RF interference in WBANs is often different from traditional networks: a mobile WBAN may experience gradual increase in RF interference as the person approaches the interference source with a peak at the closest distance and afterwards a decrease as the person withdraws from the interferer.

These effects result in significant spatio-temporal signal strength variations, which are quite unique. In the next section we will give an overview on how these effects have influenced the development of WBAN channel models.

2.2.4 Channel Models

Since the WBAN channel has some properties which are quite different from other wireless channels, the research community has started to investigate how to better model these channels. As the operating conditions of WBAN are diverse it is difficult to establish generic analytical models. Instead, models are often developed for a specific environment (e.g. indoor vs. outdoor) or human activity. In addition, WBAN channel models are typically developed for a certain frequency range, because the wavelength has large impact on the signal propagation. For example, scattering is especially pronounced in the 60 GHz band, where wavelengths are short and thus surfaces appear rougher [115]. While there also exist models for off-body [116] and body-to-body [117] communication, in this thesis we focus on narrowband *on-body* signals at a frequency around 2.4 GHz. Therefore, in the following we will describe only models that match these criteria, if not stated differently.

Fading Distributions

A common approach is to model the WBAN channel with the help of statistical distributions that describe the signal amplitude. As such distributions abstract from any time dependency or temporal correlation in the signal, this approach assumes that the channel is wide-sense stationary process. Various authors have derived WBAN fading distributions, typically by determining
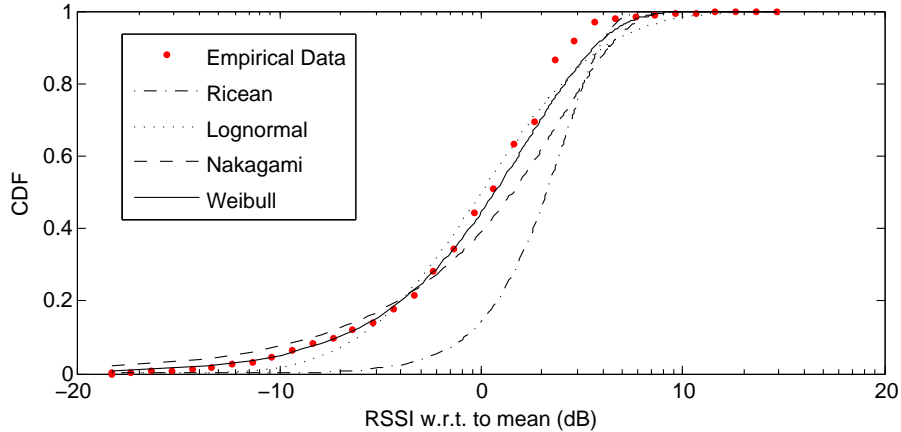
Figure 2.5: Empirical CDF and fading distributions.

the best fit out of well-known distributions like Ricean or Lognormal with their empirical data. For example, after having conducted on-body propagation measurements in four different environments Nechayev et al. [118] concluded that Ricean fading represents the best fit for their data. Others have reported Rayleigh [119], Nakagami [120], and yet others Weibull [106] or Lognormal [121] distributions. An overview is provided by Cotton et al. [122]. The latter also conclude their survey by explaining that there is currently no consensus in the literature on one model for signal strength distributions; rather, certain models seem to be the best choice for certain transmitter-receiver arrangements and usage scenarios.

For example, Fig. 2.5 shows the distribution of signal strength measured by us over several minutes on a WBAN node attached to the chest (receiver) with a sender located in the right trouser pocket of a person performing office work in front of a Personal Computer (PC). We used Maximum Likelihood Estimation (MLE) to determine the parameter estimates for different fading distributions that maximize the probability of the empirical data. In this particular case, the Weibull and Lognormal distributions are good fits.

Fading distributions are sometimes complemented by path-loss models. For example, one indoor path-loss model that the IEEE 802.15.6 channel modeling subcommittee has adopted is the standard log-distance path loss model with log-normal shadowing shown in Eq. 2.3; they suggest a rather high path loss exponent of $\alpha = 6.6$ for an indoor hospital scenario [109]. On the other hand, Fort et al. have compared different path-loss models and concluded that, based on their empirical data, this model is a rather bad choice [119]. Instead, the authors suggest the following path-loss model, which has also been adopted by the IEEE 802.15.6 channel modeling subcommittee [123]:

$$PL(d)[dB] = -10 \, log \left(P_0 \, e^{-m_0 \cdot d} + P_1 \right) + X_\sigma \tag{2.4}$$

where P_0 represents average losses close to the transmitter antenna, P_1 is the "average attenuation of components in an indoor environment radiated away from the body and reflected back towards the receiving antenna", and m_0 is the average decay rate in dB/cm of the creeping wave that diffracts around the body [119].

One less common approach to model the WBAN channel is based on the Finite Difference Time Domain (FDTD) concept, which has, for example, been used by Ryckaert et al. [113]. FDTD is based on numerical analysis to find solutions for Maxwell's equations by calculating the E and H fields in space. This requires the material of objects in the environment to be precisely specified in terms of their permittivity and conductivity. For example, the body shape may be represented by a set of simple geometric objects, such as elliptical cylinders or by detailed human phantom. The permittivity and conductivity parameters for body skin, fat or muscles may be used to describe the body structure. Once an anatomic model of the human body has been created FDTD simulation can be used to determine path-loss or more detailed information, such as delay spread, between arbitrary positions on the human body. It has also been suggested to combine FDTD with animation software to better understand the impact of human movements [124]. However, due to its computational complexity FDTD simulation is quite time-consuming and rarely used.

Modelling Temporal Variations

The fading and path-loss models described in the previous paragraph have an obvious shortcoming, namely they do not consider the temporal correlation of the channel. In mobile scenarios this is not realistic as explained in Section 2.2.3). In fact, even when a person is standing still wide-sense stationarity may not be given, because the breathing often creates noticeable periodic signal strength patterns [125]. The following gives an overview of techniques that have been used to model temporal variation (second order statistics) in WBANs.

One metric to describe the temporal variability of the wireless channel is coherence time, which was introduced in Section 2.2.2. Several authors have experimentally investigated WBAN coherence time. For example, Smith et. al have conducted an empirical study with subjects walking and running [106]. Based on their measurement results they concluded that the channel coherence time (defined as a correlation coefficient above 0.7) is "36 to 73 ms for walking, with a mean of 48 ms, and from 23 to 66 ms, with a mean of 31 ms for running" [106]. Fu et al. performed similar experiments and measured a coherence time (defined as a correlation coefficient above 0.5) of 125 ms for walking and 27 ms for jogging scenarios (and 381 ms when standing still) [105].

Two commonly used second-order characteristics are LCR and ADF (Section 2.2.2). LCR describes the expected rate at which the signal crosses a certain threshold, and for many fading distributions analytical closed-form expressions for LCR have been derived. Similarly, closed-form expressions for ADF, which is the expected time the signal stay below a certain signal level, have been derived. An overview of the analytical models is given in [122]. An empirical study that has measured LCR and ADF in several WBAN scenarios is presented by Hall et al. [126]. The authors selected a threshold of 10 dB below the link's mean signal strength and reported an LCR between 0.2 Hz and 0.8 Hz and an ADF of 60 ms to 250 ms when the subject was walking in different indoor and outdoor environments.

LCR and ADF are not able to describe channel periodicities. Yet, even by simple visual inspection of the WBAN signal in the time-domain it is sometimes obvious that the wireless channel exhibits periodic fluctuations when the subject performs repetitive patterns, such as walking (c.f. [127, Fig.17] or [122, Fig.11.2]). Others have analyzed the Autocorrelation Func-

31

tion (ACF) and revealed periodicities that match a subject's cyclical movement pattern [111]. Lately, there has been an increasing interest to better model such channel periodicities. Cai et al. have derived a channel model that captures signal strength variation during walking [128]. They suggest to model the WBAN channel — between hip and wrists — by two states, which represent the closest and farthest transmitter-receiver distance (and thus smallest and largest path loss). Each of the states is represented by a log-normal distribution of signal amplitudes; furthermore, the time period that describes a full cycle through both states is also log-normally distributed. More recently, Roberts et al. have studied temporal correlation of signal strength when the subject is walking [129]. They suggest to complement the traditional path loss model shown in Eq. 2.3 or Eq. 2.4 by an additional term that captures channel periodicity as follows:

$$PL(t)[dB] = PL(d) + 2\sigma_{PL} \cdot sin(2\pi \cdot f_{PL} \cdot t) \tag{2.5}$$

where $PL(d)$ is the path loss calculated by Eq. 2.3 or Eq. 2.4, σ_{PL} is the standard deviation of the signal fluctuation and f_{PL} is the frequency of the movement pattern (e.g. inverse of the subject's step or stride period). Note that the independent variable is now in the time domain, i.e. parameter d is now time-dependent. The authors suggest to derive f_{PL} by representing the "amount of user movement" with a value between 0 and 1, for which they provide basic examples; and σ_{PL} based on the relative movement among two sensors again with a value between 0 and 1 in conjunction with fitting empirical data.

Finally, Tselishchev et al. [130] have proposed to apply the popular Gilbert-Elliot channel model [131] to capture the channel dynamics in WBANs. This model represents the channel by a Markov chain with two states: a "good" state which has a Bit Error Rate (BER) of (close to) 1 and "bad" state, with BER of 0.5. In [130] the authors derived transition probabilities from a large pool of experiment traces and used the model to compare different slot assignment techniques standards, which is explained in more detail in Section 2.3.1

WBAN Network Simulators

Although several WBAN channel models have been proposed in the past, network simulators that support on-body communication are rare. Extending existing simulators with WBAN fading distributions (c.f. Section 2.2.4 and Section 2.2.4) is not sufficient, because these distributions do not consider temporal correlation of the channel. Since human movements largely determine the channel dynamics, accurate modelling of the human mobility may in fact be considered a prerequisite to modelling the on-body communication channel. The effect of mobility on (relative) arrangement of the nodes on the body could then, for example, be used to apply a path-loss model with (varying) inter-node distance or shadowing.

Many network simulators have been extended by mobility models that assume uncorrelated node mobility. Random Walk Mobility Model (RWMM) [132] is one popular example. When modelling on-body communication such models are inappropriate, because WBAN nodes are attached to the same body and thus their movements are correlated. Group mobility models have been proposed, but they are often based on the assumption that there is a one-to-one relationship between nodes and persons; group mobility models then try to model the movement of these individuals based on, for example, a certain social behavior that governs the movement of the

people. Such models are also not directly applicable for WBAN. Only recently, Nabi et al. have developed a mobility model for WBAN that supports mobility of multiple nodes attached to the same body [133]. It is designed as an add-on to the MiXiM framework [134] of the OMNeT++ network simulator [135]. Their model allows to specify a sequence of postures, where each posture is characterized by a set of parameters such as relative node position and the velocity of movement of a node within a given sphere.

Another approach was recently proposed by Boulis et al., who suggest to model temporal changes with the help of conditional probabilities derived from WBAN experiment traces [121]. They assume that the channel does not change within 10 ms, it is highly correlated up to 50 ms, it is partially correlated between 50 ms and about 1-2 s and it is completely independent if the time interval is larger than 5 s. After quantizing the empirical data, they compute a set of conditional distributions which can be used to determine current signal strength given a certain value was observed at a given time in the past (if the time interval is larger than 5 s they simply sample from the ensemble fading distribution). This model has recently been integrated in the WSN simulator Castalia [136], which is a WSN and WBAN simulator based on the OMNeT++ network simulator. The WBAN network simulator extensions described above have, however, been available only recently it remains to be seen how the community adopts these tools.

2.3 Physical and Data Link Layer Protocols

As explained in Section 2.1.4 intra-WBAN communication is often based on a logical star topology. Therefore the lower two protocol layers (Physical and Data Link) of OSI the model are required, but network and transport layers are often omitted. This section provides an overview of WBAN communication protocols that realize the lower two layers of the OSI reference model. We first discuss MAC protocols that have been proposed by academia and then provide an overview of PHY/MAC communication standards applicable in the WBAN domain.

2.3.1 MAC Protocols

Since wireless communication occurs in a broadcast medium a method is needed to share the medium among the different transmitters in order to avoid collisions. The task of a MAC protocol is to regulate access to the channel. In energy-constrained networks, such as WSNs, the MAC should also minimize additional "sources of inefficiency" [137]: apart from avoiding collisions a MAC protocol should reduce the amount of idle listening (the receiver's radio is enabled, but no transmission is ongoing) and overhearing (a packet for a different destination is received). In addition, control overhead, such as the number of control bits added to a packet in form of a MAC protocol header, should be minimized. These four goals — minimizing collisions, idle listening, overhearing and control overhead — are also valid for WBAN MAC protocols and have to be addressed in light of the general application requirements of energy efficiency, transmission reliability and latency. Depending on the particular WBAN application, additional requirements such as low jitter and non-functional properties like security may also have to be considered by the MAC protocol.

Over the past ten years the academic community has developed a myriad of MAC protocols

for WSN (for an overview refer to recent survey papers [138, 139]). Some of these protocols have also been considered for use in WBANs. However, the particular application requirements and operating conditions of WBANs imply that some design decisions may be unsuitable. For example, WBANs often realize a star topology, while WSN MAC protocols are typically designed to support multi-hop networks. WSNs often consist of identical sensors, whereas WBAN nodes may incorporate different sensors with individual traffic characteristics; furthermore, the core of a WBAN is typically a more capable (and less energy-restricted) node, possibly integrated in the user's cell phone. Also, the number of nodes in a WSN may greatly exceed the number of WBAN nodes and a WBAN may have tighter constraints with respect to reliability or delay when medical data are involved. Furthermore, in contrast to some WSNs that generate rather sporadic traffic (e.g. a few packets per hour in an environmental monitoring application), WBAN applications may often generate substantial (continuous) traffic (c.f. Table 2.1). Last but not least, due to mobility the wireless channel in a WBAN may have rather different characteristics than a WSN channel as elaborated in Section 2.2.3. Therefore, the research community has lately proposed several MAC protocols specifically targeting the WBAN domain. In the following we provide a brief overview of these protocols.

MAC protocols are often classified according to the following basic channel access methods: Time Division Multiple Access (TDMA) splits time into (periodic) slots and assigns a subset of the slots to a certain user for exclusive access; Carrier Sense Multiple Access (CSMA) requires a transmitter to check whether another transmission is in progress before trying to send; Frequency Division Multiple Access (FDMA) allocates one or several frequency bands (channels) to a certain user so that users do not interfere when they transmit concurrently; and Code Division Multiple Access (CDMA) requires a transmitter to apply a code to (spread) the information, which allows multiple transmissions to be multiplexed over the same channel. In practice, however, a communication protocol (stack) is often a hybrid mix like the latest revision of the IEEE 802.15.4 MAC standard (c.f. Section 2.3.2): with Direct-Sequence Spread Spectrum (DSSS) the physical layer involves concepts related to CDMA, channel access is slotted into distinct time intervals, some of which are exclusive (TDMA) while others require contention-based access (CSMA); and transmissions may be scheduled on different non-overlapping frequencies (FDMA) in order to avoid collisions.

The IEEE 802.15.4 MAC has also been the source of inspiration for several WBAN MAC protocols. For example, Yun et al. propose OD-MAC (on-demand MAC) [140], which allows to dynamically modify the superframe structure of the IEEE 802.15.4 MAC. The beacon interval is adapted based on the currently required bandwidth, packet delivery ratio and energy efficiency. Taking the requirements of all nodes into account a new superframe structure is defined that maps real-time traffic to the TDMA portion of the 802.15.4 MAC and the remaining traffic to the CSMA interval. The authors show simulation results that indicate that in comparison to the standard 802.15.4 MAC their approach can cut the energy consumption in half. Similarly to OD-MAC, BodyMAC [141] is a TDMA-MAC protocol that dynamically adapts bandwidth to support different traffic (e.g. burst or periodic) patterns and utilizes the gateway compute a TDMA schedule. BodyMAC has been compared to the IEEE 802.15.4 MAC in a simulation study and the results show that it can outperform the standard in terms of delay and energy efficiency. Zhang et al. have also proposed a MAC protocol that modifies the IEEE 802.15.4 su-

perframe structure to support different traffic classes [142], which are mapped to different time intervals to support control and bursty or periodic data traffic. These proposals are representative for a class of WBAN MAC protocols that try to support different types of WBAN traffic or packet priorities by (dynamically) structuring time into different contention-free or contention-based intervals and mapping certain traffic classes onto them. One main shortcoming of these proposals is, however, that they do not take the specific WBAN operating conditions — in particular the unique channel characteristics and RF interference environments — into account. Since the protocols are often evaluated with standard (WSN) simulators that do not model these effects it is not entirely clear how they would perform in real-life scenarios.

Some MAC protocols are more closely tailored to the WBAN domain: H-MAC [143] is a TDMA protocol that uses the heartbeat pattern of the subject wearing the WBAN nodes to perform the necessary time synchronization. The authors report that with their protocol periodic control packets (beacons) are unnecessary, because sensors (ECG, blood pressure or SpO_2) can extract timing information from a common time source: the heart. In a simulation study based on real-world ECG and blood pressure data, they show that their approach can extend network lifetime due to the lack of explicit control traffic. Su et al. have also proposed a TDMA MAC for WBANs that regulates channel access by taking into consideration the properties of the battery, channel conditions, and packet queues [144]. The authors propose a duty-cycling approach in order leverage the battery's recovery process and thus improve lifetime of a node. They also suggest to use an adaptive modulation scheme to take channel conditions (based on a Nakagami-m model parameterized for WBANs) into account. Time is structured in "periodical frame periods" bounded by beacons, which contain an active (TDMA) and an inactive (radio is off) phase, similarly to the IEEE 802.15.4 standard. With a customized event-driven simulator that integrates a battery discharge model they show that their protocol can outperform other protocols such as the IEEE 802.15.4 MAC in terms of energy efficiency.

Like the proposals described in the paragraph before, the last two protocols have not been evaluated experimentally. What has become accepted practice in the WSN domain, namely the experimental evaluation of protocols in realistic environments (testbeds), is often often not the state-of-the-art for WBAN protocol evaluation. Unfortunately the WBAN domain shows a **shortage of experimental protocol evaluation**. This is problematic, because network simulators are often not capable of capturing the channel and interference characteristics unique to WBANs. Boulis et al., who have performed a large number of WBAN channel measurements at NICTA, argue that an efficient WBAN MAC protocol must take the unique channel characteristics into account when assigning dynamic time slots or selecting a retransmission policy [145], which is non-trivial because "the movement of the human body has a dramatic effect on the strength of the received signal" [145]. In another study the authors have evaluated several variable packet scheduling strategies based on the popular Gilbert-Elliot channel model, which represents the channel by a Markov chain with two states ("good" and "bad"). The authors derived transition probabilities from a large pool of experiment traces and used the model to compare different slot assignment techniques. Depending on the amount of knowledge that the scheduler (coordinator in a star topology) has, a simple greedy schedule or a "flipping" strategy (inverting the order of the previous schedule) is performing best.

2.3.2 Communication Standards

The analysis in Sect. 2.1.3 has shown that many envisioned WBAN applications generate low data rates of no more than 150 byte per second (c.f. Table 2.1). In this section we review three low-power communication standards that match this requirement. There also exist several proprietary WBAN communication solutions [76, 78, 79], which are typically not interoperable. Communication standards, however, enable interoperability and often represent a major step towards greater adoption of a technology.

IEEE 802 describes a set of communication standards for personal, local and metropolitan area networks, which all cover the lower two layers (Data Link and Physical) of the OSI reference model. The IEEE 802.15 standards most relevant for WBAN communication are IEEE 802.15.1 (Bluetooth), IEEE 802.15.4 and IEEE 802.15.6. In this section we provide an overview of these standards. Bluetooth low energy and IEEE 802.15.6, however, are comparably new standards and so far no commercially available WBAN platforms uses these technologies.

Bluetooth Low Energy

Bluetooth is a standard for wireless short-range (up to $10\,\mathrm{m}$) communication. Bluetooth specifications are managed by the Bluetooth Special Interest Group (SIG), which has several thousand member companies. In 2002 the first release of the IEEE 802.15.1 was published standardizing Bluetooth v1.1 and in 2005 the IEEE 802.15.1 group released the IEEE 802.15.1-2005 [146], which is a standard adaptation of the v1.2 specification. Since then the Bluetooth SIG has released further Bluetooth specifications up to the current v4.0.

Bluetooth supports a range of applications and has traditionally been used to connect keyboards, head-sets or printers to computers or cell phones.[11] The Bluetooth SIG has specified several variants of Bluetooth, all of them having the following properties in common: a Bluetooth network, also called piconet, realizes a star topology and consists of one master and one or more slave nodes;[12] the channel access method is TDMA; communication occurs in the $2.4\,\mathrm{GHz}$ ISM band based on Frequency-Hopping Spread Spectrum (FHSS) over an $80\,\mathrm{Mhz}$ bandwidth; and Bluetooth can achieve theoretical data rates of $1\,\mathrm{Mbps}$ or more.

The Bluetooth variant most suitable for WBAN applications is likely BLE. BLE has evolved from Nokia's WiBree, and is a part of the Bluetooth v4.0 specification, which was published in 2010. In the following we provide a technical overview of BLE, and explain how it differs from the traditional, Enhanced Data Rate (EDR) Bluetooth.[13] Bluetooth includes upper layer protocols and application profiles, but we focus on the lower protocol layers (Data Link and Physical), because only they are relevant to this thesis.

BLE splits the $2.4\,\mathrm{GHz}$ ISM band into 40 non-overlapping channels, each with a channel width of $2\,\mathrm{MHz}$ (previous Bluetooth variants: 79 channels with a width of $1\,\mathrm{MHz}$). Out of the 40 channels 3 are used for device discovery and connection establishment and 37 for data

---

[11]Bluetooth was originally created by Ericsson with the intention to wirelessly connect their cellular phones with other devices

[12]A Bluetooth star topology can be extended to a multihop network, but this is beyond the Bluetooth specification.

[13]BLE and classic Bluetooth are incompatible, however, dual-mode devices are available that implement both variants in one device.

Figure 2.6: Two consecutive Bluetooth low energy connection events.

transmission. This considerable increase connection times to a few ms, which is much faster than "classical" Bluetooth that used 16 or 32 advertising channels. The physical layer uses Gaussian Frequency Shift Keying (GFSK) and provides a data rate of 1 Mbps. BLE uses frequency-hopping, but rather than a constant hopping frequency (800 or 1600 Hz in previous versions of Bluetooth), BLE stays on a single channel for an entire *connection event*. As indicated in Fig. 2.6 a connection event is a periodic time interval, during which a sequence of packets can be exchanged between a master and slave. Since TDMA is used for channel access, these packets must fit into the slot assigned by the master to the slave device. Within a connection event packets are exchanged until neither master nor slave have any data left (indicated by a flag the packet header), or the reception of two consecutive packets has failed. In either case, further data transmission will happen in the subsequent connection event on a different channel. An Adaptive Frequency Hopping (AFH) algorithm is used to select the next channel; the algorithm can decide to temporarily skip certain channels in case they are interfered, which it can measure periodically.

A slave has to be active only within its own connection event. The time between two consecutive connection event (*connInterval*, c.f. Fig. 2.6) can be configured between 7.5 ms and 4 s; and an additional parameter defines the number of connection events that a slave can skip (keep the radio off) after a previously successful connection event. Due to these mechanism BLE can achieve significantly lower power consumption than classic Bluetooth. BLE devices have only recently become commercially available, but studies already indicate that with proper parametrization BLE can achieve network lifetime of several years [147]; other measurements indicate that power consumption of BLE devices is comparable to that of IEEE 802.15.4 devices [148]. This is mainly because (a) both protocols allow very low duty cycles, (b) due to its TDMA approach BLE does not require a radio to perform a Clear Channel Assessment (CCA) before tranmission, which saves energy, and (c) BLE radios theoretically provide 4 times higher data rate than IEEE 802.15.4 transceivers but consume comparable amount of current [148]. On the other hand, application throughput of BLE has been reported to be often much lower than theoretical limits (well below 100 kbps [147, 149]). Nonetheless, BLE is a promising candidate technology for WBAN applications, which has also been demonstrated through pilot projects performing glucose monitoring [150] or ECG monitoring [151] with BLE devices.

IEEE 802.15.6

In 2012 the Institute of Electrical and Electronics Engineers (IEEE) 802.15 Task group 6 released the IEEE 802.15.6 standard [81] specifying "short-range, wireless communications in the vicinity of, or inside, a human body (but not limited to humans)" with data rates up to 10 Mbps.

	contention-based	contention-based	contention-free / improvised and unscheduled	contention-based	contention-based	contention-free / improvised and unscheduled		contention-based
B1	EAP1	RAP1	MAP1	EAP2	RAP2	MAP2	B2	CAP

←————————————IEEE 802.15.6-2012 superframe————————————→

Figure 2.7: The IEEE 802.15.6 superframe structure.

The standard covers three different physical layers: (1) a narrowband PHY that supports different licensed bands such as the MICS (420 MHz to 450 MHz) or MBAN bands (c.f. Table 2.4) but also the unlicensed 868 MHz, 915 MHz, or 2.4 GHz ISM band. (2) An UWB PHY using either impulse radio UWB and wideband frequency modulation, with a channel bandwidth of about 500 MHz in one of 11 channels either low-band (3-5 GHz) or high-band (6-10 GHz). And (3) a Human Body Communications (HBC) PHY based on capacitive coupling (electric field communication) at a center frequency of 21 MHz.

An IEEE 802.15.6 WBAN forms a star topology with one central controller (called hub) and zero or more nodes. In addition, a two-hop extended star topology is supported, which allows the hub and another WBAN node to exchange packets via a relayer WBAN node. The 802.15.6 MAC supports several different modes of operation. There exists a beacon mode and a non-beacon mode. In beacon mode the hub transmits periodic beacons which mark time intervals called superframe (beacon period). A superframe is subdivided into several different (partially optional) intervals which allow certain types of channel access: either random channel access based on CSMA with Collision Avoidance (CSMA/CA) or slotted ALOHA, or scheduled (contention-free) access or so-called "improvised and unscheduled" access based on polling (hub uses data request messages to obtain messages from devices) or posting (hub sends a notification messages to devices). Furthermore, the MAC supports 4 different types of acknowledgement policies including the use of cumulative acknowledgements. In non-beacon mode either a superframe structure is used or not. In the first case the hub transmits Timed (T-Poll) frames to inform devices about superframe timing.

Fig. 2.7 visualizes the superframe structure in the IEEE 802.15.6 beacon mode. Each superframe begins with the transmission of a beacon frame by the hub, which is followed by an Exclusive Access Phase (EAP) reserved for high-priority traffic such as "emergency or medical implant event report" 2.7. The EAP is followed by a Random-Access Phase (RAP) which can be used for the remaining traffic. Channel access in both, EAP and RAP, is based on CSMA/CA or slotted ALOHA (defined by a MAC sublayer parameter). The RAP is followed by a Managed Access Phase (MAP), which is used for scheduled, improvised or unscheduled channel access. In the MAP the hub controls the devices by either reserving slots for devices or by sending poll or post packets to request device data or distribute information to devices, respectively. As can be seen in Fig. 2.7 the EAP/RAP/MAP cycle is repeated once and followed by yet another optional beacon and a Contention Access Phase (CAP). The CAP allows to send any traffic based on CSMA/CA or slotted ALOHA channel access. The narrow band PHY (optionally) also supports channel hopping. In this case the hub defines a pseudo-random hopping pattern, which it uses to switch the communication channel after staying on a channel for a fixed number of superframes. A parameter defines the minimum number of channels separated between two consecutive hops.

Figure 2.8: The IEEE 802.15.4-2006 superframe structure.

Although for a low-power wireless standard the IEEE 802.15.6 can be quite complex it is widely regarded a promising technology to enable intra-WBAN communication. So far no commercially available IEEE 802.15.6 transceivers exist, but recently a 802.15.6 transceiver was published which is reported to consume very low power (only 5 mW in Rx/Tx state) [152].

IEEE 802.15.4

In 2003 the task group 4 of the IEEE 802.15 working group released the first edition of the IEEE 802.15.4 standard [153] to enable wireless connectivity between "ultra-low complexity, ultra-low cost, ultra-low power consumption, and low data rate" devices. The standard has attracted strong interest in industry and academia and has been complemented by several other solutions, such as ZigBee [88], IETF 6LoWPAN [154] or Wireless Highway Addressable Remote Transducer Protocol (WirelessHART) [155].

Following the general IEEE 802 policy the standard covers PHY and MAC sublayer. The PHY services are structured into a data and management plane. The PHY data service allows to transport MAC Protocol Data Units (MPDUs) between peer MAC sublayer entities; the PHY management service allows to control the internal operating state of the transceiver (receive/transmit/off), to perform CCA, energy detection measurements within the current channel and to set PHY-specific attributes like the RF channel or transmit power. The standard specifies four PHYs in the 868/915/2450 MHz ISM bands with different modulation techniques. In the popular 2.4 GHz band the DSSS PHY achieves a data rate of 250 kbps using offset quadrature phase-shift keying (O-QPSK). One 4-bit symbol is mapped to a 32-chip sequence and chips are transmitted at a nominal chip rate of 2 Mchip/s resulting in a symbol rate of 62.5 ksymbol/s. The transmit power of a 802.15.4 transceiver is typically around 0 dBm.

The MAC services are structured into a data and management plane. The data service allows to transfer a MAC Service Data Unit (MSDU) to a peer device, which may include an acknowledgement from the peer device and/or several retransmissions. The management service is responsible for device configuration, periodic transmission of and synchronizing to beacons, enabling Personal Area Network (PAN) association and disassociation, employing security mechanisms and handling the Guaranteed Time Slot (GTS) mechanism.

The MAC sublayer supports different configurations and operating modes. One commonality is that every network has exactly one PAN coordinator, which is the primary controller responsible for PAN identifier and device address assignment and device synchronization. 802.15.4 PANs can either be nonbeacon-enabled or beacon-enabled. In a nonbeacon-enabled PAN frames are transmitted according to an unslotted CSMA/CA algorithm (nonpersistent CSMA): if the channel is detected idle the transmission can start immediately otherwise the device defers the transmission for a random time period uniformly drawn from an exponentially increasing back-

off interval. The destination device(s) either have to always listen or achieve synchronization with schemes beyond the scope of the standard.

In beacon-enabled mode coordinators periodically transmit beacons which mark the beginning of a superframe as depicted in Fig. 2.8. A beacon carries information about pending data and the current network configuration. Immediately after the beacon follows the CAP. During the CAP devices use a slotted variant of the CSMA/CA algorithm: a device must sense an idle channel twice before it may transmit and both, channel sensing and transmission must be performed on backoff slot boundaries. The CAP is followed by an optional Contention-Free Period (CFP), which is portioned in so-called GTS. GTSs are allocated dynamically and the corresponding time interval can be used exclusively to transmit packets in a contention-free fashion. The CFP is followed by an optional inactive period in which all nodes can sleep to preserve energy and achieve low duty cycles.

The initial IEEE 802.15.4 standard has been revised in 2006 with minor enhancements and clarifications [18], and once more in 2011 mainly with editorial changes to encompass additional PHY layers that have been specified in previous amendments, for example an UWB or Chirp Spread Spectrum (CSS) PHY in IEEE 802.15.4a-2007 [156]. In 2012 an amendment of the MAC was published as IEEE 802.15.4e [157] with the goal to "enhance and add functionality to the IEEE 802.15.4 MAC to better support the industrial markets". One of these extensions was the Timeslotted Channel Hopping (TSCH) mechanism, which is based on a pseudo-random hopping sequence: time is divided in a set of discrete time slots, which are large enough to exchange one data and acknowledgement frame between a pair of devices, respectively. A slot is associated with a directed link between two devices as well as certain channel (communication frequency). When a device has data to send to another device, it has to wait until the respective slot is active. Multiple slots are contained in one *slotframe*, and a slotframe continuously repeats in time. More details on TSCH are provided in Section 7.2.1.

The IEEE 802.15.4 standard has been considered a candidate for WBAN communication for many years. In the following section we discuss its role for WBAN applications.

Discussion

The three communication standards introduced in the previous sections are all considered candidates for WBAN communication. The IEEE 802.15.4 standard has received most attention so far, because the BLE and IEEE 802.15.6 standards have only recently been published (and BLE chips have recently started entering the markets, but commercially available IEEE 802.15.6 transceivers do not yet exist).

Previous studies have shown that IEEE 802.15.4 is well-suited for many WBAN applications [158–160], and several WBAN research projects described in Section 2.1.4 have successfully been building on IEEE 802.15.4 technology. Since it is also one of the most mature and applied low-power communication standards, the IEEE 802.15.4 standard has also been used to evaluate several algorithms presented in this thesis. The ideas these algorithms are built upon are, however, often not IEEE 802.15.4-specific, but exploit the characteristic operating conditions of a WBAN. Nonetheless the IEEE 802.15.4 has played an important role in the design and evaluation of our solutions, and therefore in the following we describe this implementation in more detail.

2.3.3   The TKN15.4 Implementation

The evaluation of the algorithms presented in this thesis is based on experiments that involve IEEE 802.15.4 hardware (PHY). A stable, open-source software 802.15.4 MAC implementation was, however, not available for our platform (and any other hardware) platform. Yet, some of our algorithms required modifications of MAC core functions. Therefore we implemented the IEEE 802.15.4-2006 MAC in the TinyOS 2 operating system. Our implementation is by now part of the official TinyOS 2 core, accessible via the project source code [161, $TOS-ROOT/tos/lib/mac/tkn154]. It covers the MAC functionality and includes the interfaces towards the layer above (e.g. network layer) and below (radio driver). The design and implementation of the radio driver (PHY), however, is platform/chip specific and thus not part of the implementation. In the following we briefly introduce the TinyOS 2 operating system and then describe the MAC implementation, which is also known as *TKN15.4*. More detailed information on the implementation can be found in [35–37].

TinyOS 2

TinyOS [162] is an operating system for wireless embedded sensor networks. It consists of a set of components and the *nesC* language [163] allows to combine these components to create an application. TinyOS 2 [164] is the second generation of the operating system that keeps many of the basic ideas of its popular predecessor while pushing the design in several key areas.

TinyOS 2 improves system reliability and robustness by redefining some of the basic TinyOS 1.x abstractions and policies such as initialization, the task queue, resource arbitration and power management. For example, in TinyOS 2, every task has its own reserved slot in the task queue and can be posted only once. These semantics lead to greatly simplified code (no need for task reposting on error) and more robust components. The same principle of compile-time allocation and binding is applied to many other aspects of the system: components allocate all of the state they might possibly need; and the invariants are explicitly reflected by the components and their interfaces, rather then being checked at runtime. This design principle limits the flexibility, but makes many Operating System (OS) behaviors deterministic.

TinyOS 2 supports code reusability and portability in several ways. It is based on a nesC version that uses the concept of generic components, which, similar to object-oriented programming, can be used to create multiple component instances, but instantiated at compile-time. It also supports network types at the language level: programs can declare data structures and primitive types that follow a cross-platform (1-byte aligned, big- or little-endian) layout and encoding. This allows services to specify platform independent packet formats without resorting to macros or explicit serialization. Finally, it uses a three-layer Hardware Abstraction Architecture (HAA) [165] that decomposes the functionality of an individual subsystem, for example timers, into three distinct layers. The HAA achieves code portability through platform independent interfaces at the top layer, but at the same time allows the applications to access a subsystem's full capabilities through the chip-specific second layer, if performance requirements outweigh the need for portability.

There are two main TinyOS-specific challenges for a platform independent 802.15.4 MAC implementation: first, TinyOS is not a real-time operating system, yet in beacon-enabled mode
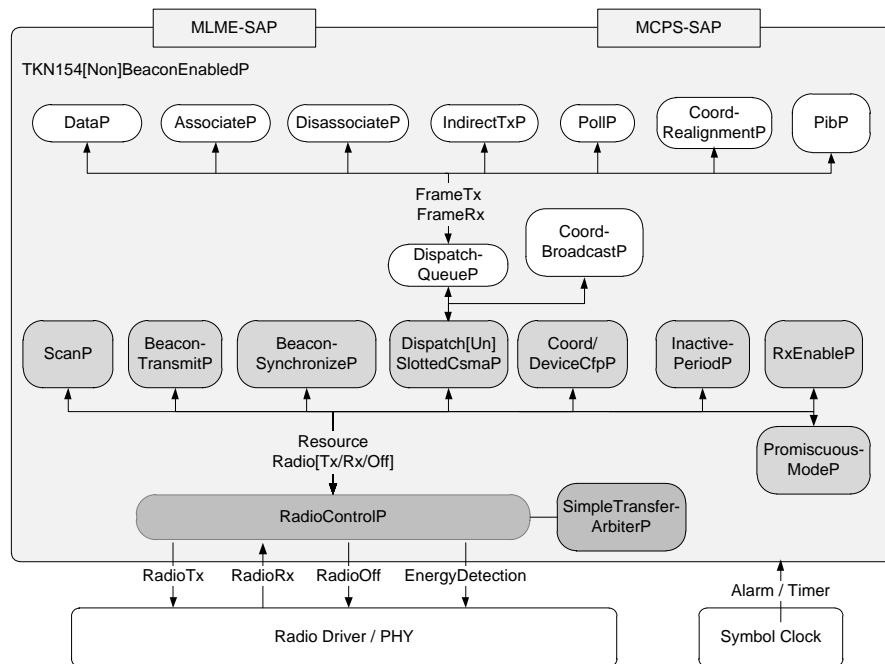
Figure 2.9: Architecture of the IEEE 802.15.4 MAC implementation *TKN15.4*.

some operations need to be accurately timed. TKN15.4 solves this problem by pushing most time-critical operations from the MAC to the chip-specific radio drivers, because the drivers typically operate in a asynchronous (interrupt) context and may better exploit the particular hardware characteristics, for example hardware-generated acknowledgements. Second, TinyOS 2 MAC protocols are traditionally implemented for a particular radio chip. TinyOS provides the *Active Message* layer to multiplex access to the radio above the link layer, but there exist no hardware abstractions for the lower layers. Since there are no established interfaces or guidelines in TinyOS 2 on how the radio hardware should be exposed or how MAC protocols are to be structured, one contribution of our work was the proposal of a set of platform independent radio interfaces to be provided by an IEEE 802.15.4-compliant radio chip abstraction in TinyOS 2.

Implementation

Fig. 2.9 provides an architectural overview of the TKN15.4 implementation, its main components and the interfaces that are used to exchange MAC frames between components. While this figure abstracts from the majority of interfaces and some configuration components, it illustrates one important aspect, namely how access to the platform specific radio driver (PHY) is structured. For the purpose of explanation, the MAC can be subdivided into three sub-layers. In the following we explain the components and interfaces on these layers, starting from the bottom.

On the lowest level (dark gray boxes), the *RadioControlP* component manages the access to the radio: with the help of an extended TinyOS 2 arbiter component it controls, which of the components on the level above is allowed to access the radio at what point in time. The use of

a TinyOS resource arbiter avoids inconsistencies in the radio driver state machine and is in line with the standard TinyOS 2 resource usage model: before a component may access a resource it must first issue a request; once it is signaled the corresponding *granted()* event by the arbiter the component can use the resource exclusively; and after usage the resource must be released.[14] The TKN15.4 implementation extends this model by allowing a component that owns the radio resource to dynamically transfer the ownership to a specific other component. This is necessary, because two properties make the standard arbiters suboptimal for beacon-enabled PANs: (1) the transition of the resource ownership from one component to another may take an arbitrary time, because it typically involves posting a task (the *Resource.granted()* event is signaled in a separate task) and (2) the resource arbitration model is based on the assumption that clients are independent components and arbitration is based on a fixed queuing policy (e.g., round robin or first-come-first-serve). The first is suboptimal, because an IEEE 802.15.4 superframe is represented by multiple components (c.f. Fig. 2.8) and as there is a continuous transition between the different intervals of a superframe the access to the radio should be transferred immediately between the involved components. For example, on a coordinator node immediately after a beacon has been transmitted the radio should be switched to receive mode. If the component responsible for managing the CAP is granted access to the radio too late (because there are other tasks served by the operating system before), then incoming frames may be lost. The second assumption property is unsuitable, because a superframe is dynamic and some parts of it are dependent on others. For example, a beacon defines whether the following CFP is present and missing an incoming beacon means that the entire superframe cannot be used by the device. Consequently the components that represent the superframe should be able to decide cooperatively at what time which of them should be active. Instead of treating them as independent clients that compete for access to the radio chip, it is more natural to let the parts of a protocol coordinate when and in what order they access the radio. Therefore we extend the standard Resource interface by a single command and a single event that enables immediate transfer of the resource from one client to another. The transfer does not involve posting a separate task and may override the default queuing policy.

Most of the components on the second level of Fig. 2.9 (medium gray boxes) represent the different time intervals in an 802.15.4 superframe: the *BeaconTransmitP* and *BeaconSynchronizeP* components are responsible for transmission/reception of the beacon frame, the *DispatchSlottedCsmaP* component manages frame transmission and reception during the CAP and the *CoordCfpP* and *DeviceCfpP* components are responsible for the CFP. In nonbeacon-enabled mode a superframe structure is not used and these components are replaced by the *DispatchUnslottedCsmaP* component, which is then responsible for frame transmission and reception in nonbeacon-enabled mode.

The CSMA/CA algorithm requires one (unslotted) or two (slotted) CCA to be performed one or two back-off slot boundaries before the actual transmission. Moreover, in case of a CCA failure the transmission has to be delayed for a random time period of 0 to 255 back-off periods (equals 0 μs to 5100 μs in the 2.4 GHz band). The transmission of an acknowledgements must start between 12 symbols (equals 192 μs in the 2.4 GHz band) and 32 symbols (512 μs in the 2.4 GHz band) after the reception of the last symbol of the previous data or MAC command frame.

---

[14]http://github.com/tinyos/tinyos-main/blob/master/doc/txt/tep108.txt

Since on a typical mote platform, these requirements can practically not be met by a platform independent MAC protocol [166], rather they should be pushed from the MAC to the PHY, ideally to hardware. For these reasons, TKN15.4 does not include the CSMA/CA algorithm: the *DispatchSlottedCsmaP* and *DispatchUnslottedCsmaP* components are only responsible for the initialization of the CSMA/CA parameters, but the CSMA algorithm is implemented and executed in the platform specific radio driver. In either beacon- or non-beacon-enabled mode the *ScanP* and *PromiscuousModeP* components are providing services for channel scanning and enabling/disabling promiscuous mode, respectively. The radio arbitration mechanism is used to coordinate the activities of the components on this level so that they do not overlap in time: typically a component is active only while it has exclusive access to the radio resource. It then performs a certain task (e.g., transmission of a beacon or performing a channel scan) and afterward either releases the resource or passes it on to some other component. This mechanism avoids race conditions when accessing the radio hardware. This pattern is supported by the extended arbiter component described above.

The components on the top level of Fig. 2.9 (white boxes) implement the remaining MAC data and management services, for example, PAN association or requesting (polling) data from a coordinator. These services typically utilize data and command frame transmission/reception based on the (un)slotted CSMA/CA algorithm and consequently the components are *wired* via a send queue, *DispatchQueueP*, to either *DispatchSlottedCsmaP* (in beacon-enabled PANs) or *DispatchUnslottedCsmaP* (in nonbeacon-enabled PANs). A component on this level typically provides a certain MAC MAC Sublayer Management Entity (MLME)/MAC Common Part Sublayer (MCPS) primitive to the next higher layer, it is responsible for assembling the particular data or command frame and it accepts and processes incoming frames of the same type. For example, the *DataP* component provides the MCPS-DATA primitive to the next higher layer to send a frame to a peer device. On receipt of the *MCPS-DATA.request* primitive *DataP* will assemble the data frame and enqueue it in the send queue *DispatchQueueP*. The *Dispatch[Un]SlottedCsmaP* component will eventually dequeue the frame, and manage its transmission, e.g. whenever the CAP has become active. Afterwards, it will signal a completion event to the *DataP* component, which in turn propagates a *MCPS-DATA.confirm* event back to the next higher layer including an appropriate status code that denotes whether the transmission was successful or not.

The next higher layer accesses all MAC services either via the *TKN154BeaconEnabledP* component (in beacon-enabled PANs) or via the *TKN154NonBeaconEnabledP* component (in non-beacon-enabled PANs). These configuration components are nesC facades [167] responsible for *wiring* the MAC components together, respectively. They allow to disable/remove certain MAC functionality by specifying empty placeholder [167] components.

The 802.15.4-2006 standard specifies 17 MAC and 6 PHY primitives. The TKN15.4 implementation slightly adapts the MAC primitives to better match TinyOS 2 design principles. TinyOS 2 does not support dynamic memory allocation, instead most allocation and binding is pushed to compile time, which makes code more robust and predictable. This also holds for the allocation of message buffers: a component that wants to send a packet is responsible for allocating a *message_t* buffer, which is the TinyOS 2.x message buffer abstraction that is used by

protocols on network layer and above (cf. TinyOS Enhancement Proposal (TEP) 111[15]). Consequently the implementation adapts the MAC primitives to support *message_t* and the TinyOS buffer swapping semantics. This is relevant for two MAC primitives, *MCPS-DATA* and *MLME-BEACON-NOTIFY*, which have to be slightly revised to pass *message_t* instead of pointers to the payload. Additionally and in contrast to the primitives specified in the IEEE 802.15.4 standard, these interfaces do not transport control information (source/destination address of the frame, etc.) explicitly. Rather, this information is already contained in a *message_t* and can be configured/retrieved via suitable accessory functions, which is in line with the TinyOS 2 concept of treating *message_t* as abstract data type. A drawback of using *message_t*, however, is that a message buffer always consumes the maximum MAC payload plus header size, which frequently results in suboptimal memory usage. All other MLME/MCPS primitives are adopted one-to-one from the 802.15.4 specification, with one exception: because it is usually more convenient for the caller the MLME GET/SET primitives consist of single commands (are non-split-phase), and the beacon payload for an outgoing superframe on a coordinator is accessed not via PAN Information Base (PIB), but through a separate interface.

To better meet the tight timing constraints in the beacon-enabled mode the PHY (radio driver) interfaces had to be modified. The MAC requires the radio driver to provide interfaces that push many time-critical operations from the MAC to the radio driver, because the latter can include platform- and chip-specific code and is thus in a better position to meet the tight timing constraints in beacon-enabled PANs. The MAC specifies the initial CSMA/CA parameters but the algorithm itself is implemented and executed in the radio driver. Also, the radio driver is responsible for performing any random back-offs and for the transmission of acknowledgements.

TKN15.4 Impact

The TKN15.4 implementation currently supports three hardware platforms — Telos [66], Shimmer2(r) [8] and MICAz [168] — and the implementation has been successfully used in several research projects beyond the work presented in this thesis: for example, TKN15.4 has been used to implement wireless process control loops [169], a wireless ECG System [170] or asynchronous network discovery techniques [171]. Others have extended TKN15.4 with a GTS service [172] and security features [173]. Basmer et al. have compared TKN15.4 to other IEEE 802.15.4 MAC implementations and found that it is "the most efficient implementation in case of runtime" [174]. Furthermore, often authors have reported that the implementation meets their expectations (e.g. "simulation and experimental curves exhibit the same trend and are very close to each other" [175], "nodes behave exactly as they are supposed to do" [176]).

---

[15]https://github.com/tinyos/tinyos-main/blob/master/doc/txt/tep111.txt

Chapter 3

Scope of the Thesis

There is a growing belief that many processes will become more (cost-)efficient once the context of physical objects becomes machine-accessible. This relates to processes not only in production, logistics or agriculture, but also in the healthcare sector. Due to the continuous miniaturization of electronic components and sensors it is already possible to accurately capture the physiological parameters of a person with tiny wearable devices. The devices must, however, be connected to a system that can store, process and present the data, because their own memory and computational resources are usually not sufficient. WBANs enable the communication infrastructure to deliver sensor data along the body to a more powerful device such as a cell phone. However, WBANs are a relatively new field of research and it is unclear if they can meet the stringent quality of service (QoS) requirements that some applications have. In particular, the reliable delivery of data is considered one of the most challenging tasks in WBAN applications [16, Chapter 1].

Communication Reliability — The wireless channel is an inherently unreliable transmission medium. Within fractions of a second the strength of the electromagnetic signal may drop by a factor 100 or more. Furthermore, other co-located wireless systems may emit interfering signals that can degrade communication performance. In WBANs, however, the situation is particularly harsh, because user mobility leads to very dynamic channel conditions. The movement of human limbs may result in significant variation in inter-node distance within short time. This will translate into varying path-loss (Sect. 2.2.2) and thus *average* received signal strength may change over short time. Furthermore, human limbs may (temporarily) block the path between sender and receiver causing shadowing and thus reducing signal strength. Antenna-to-body surface separation also has an impact: as the person moves the distance between the antenna and the body surface may change with often substantial impact on received signal strength [110]. Furthermore, in practice antenna radiation patterns are never perfectly omnidirectional, thus a change in the angle between sender and receiver antenna (e.g. introduced by turning a limb) may cause signal strength fluctuation. Moreover, the effects of multipath propagation may be significant: as the person moves the amount of signal components that are reflected and scattered from the environment (ground) and human body may vary strongly and the shape of the body may promote the existence of surface waves. In addition, interference from other co-located devices may be substantial: WBANs often operate in the unlicensed 2.4 GHz communication band,

47

which is also shared by other popular wireless communication technologies, notably WLAN, and in WLANs the transmission power may be 100 times higher than in WBANs. Yet, the interference situation in WBANs is different than in static networks: on the one hand, a WBAN may experience gradual increase in RF interference as the person approaches an interference source and a gradual decrease as the person withdraws from the interferer. On the other hand, it may happen that due to body shadowing the WBAN may (temporarily) be "shielded" from an interferer. These effects collectively influence the WBAN channel and thus reception quality. The latter is often expressed by the Packet Reception Rate (PRR), which is the ratio of the number of successfully (error-free) received packets to the number of transmitted packets. PRR thus reflects the fraction of error-free communication between two direct communication peers (i.e. without intermediate nodes), and it therefore represents a good notion of communication reliability on the level of the MAC. Since our focus is on the lower layers of the WBAN communication protocol stack, in this thesis *we consider PRR the main metric for communication reliability*.

Energy Efficiency — Often, a simple solution to improve PRR is to increase the radio's transmission power. However, not only does this result in higher levels of interference for co-located devices, it also increases energy consumption. The lack of wires implies that a WBAN devices are usually battery-powered. In many scenarios the batteries have to last for several weeks or months, the radio must therefore not only consist of extremely low-power hardware components, but the transmission power must be low as well (low transmission power is also desirable in order to "reduce interference to other devices and systems" IEEE 802.15.4 Standard [18, Sect. 6.9.5]). Since energy-efficiency is one central design goal in WBANs, energy is the second metric we consider in this thesis. More precisely, we focus on the energy that is spent for the delivery of packets originating from the next higher (with respect to the MAC) layer on the sender until they are signalled to the corresponding layer on the receiver side.

Latency — Applications may require that the time between collection and processing/visualization of data is bounded. For example, in health care applications this period may span from 50 ms for capsule endoscopes [59] over 500 ms for ECG sampling [60] to several minutes for blood glucose monitoring [61]. The latency of the delivery of a packet from sender to receiver is therefore last main performance metric we consider in this thesis. In the scope of this thesis latency is the time that elapses between the next higher layer passing the packet to the MAC layer on the sender side and the time the packet is signalled to the corresponding layer on the receiver side. Note, that we do not consider the latency of corrupt frames, but only of packets that are signalled to the next higher layer on the receiver side.

3.1 Goals

The previous metrics are often in conflict: increasing the radio's transmission power may improve reliability but also increase energy consumption; retransmitting packets based on an Automatic Repeat reQuest (ARQ) scheme may improve reliability but translate packet losses into latency. Within this thesis, we consider PRR as the core metric we seek to investigate/optimize; at the same time our solutions should not "considerably" deteriorate energy consumption or
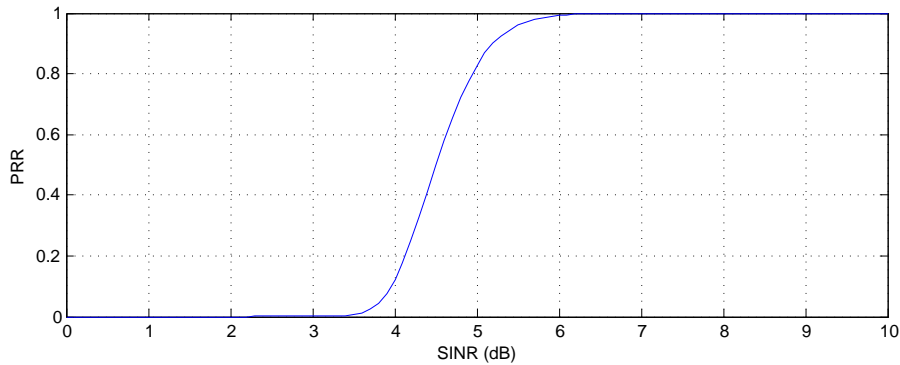
Figure 3.1: The CC2420 SINR/PRR relationship from TOSSIM.

latency. PRR is primarily related to the SINR. Recall that that the SINR is defined as:

$$SINR = \frac{P_{Signal}}{P_{Interference} + P_{Noise}} \tag{3.1}$$

The SINR is important, because it determines the probability of correctly decoding a symbol during the process of converting the analog signal to a digital sequence of symbols (demodulation). For example, the 2.4 GHz IEEE 802.15.4 PHY specification states that an SINR of 5 dB to 6 dB should typically be sufficient for good packet reception [18, Section E.3.2]. For the popular IEEE 802.15.4-compliant CC2420 radio [69] several authors have independently confirmed this statement through experiments [177, Figure 4] [29, Figure 12] [178, Figure 1]. For example, the empirical relationship between SINR and PRR reported by [29], which is also used in the TinyOS simulator TOSSIM [179], is shown in Figure 6.20. The function is monotonic and an SINR of 6 dB indeed results in high probability of successfully decoding a received packet. Therefore, essentially P_{Signal} and $P_{Interference}$ are the parameters that influence the reliability of intra-WBAN communication on the level of abstraction considered in this thesis (the P_{Noise} component in Eq. 3.1 is not addressed in thesis, because it depends on the electronic components of the receiver radio circuitry and is mainly influenced by thermal noise). The two effects that govern the dynamics of P_{Signal} and $P_{Interference}$ in a wireless system are (large and small scale) fading and interference, respectively. Fading represents the amount of attenuation the signal experiences at a given point in time due to path-loss, shadowing and multipath propagation, and interference represents any unwanted signals in the same portion of the spectrum as the carrier of the intended signal (Section 2.2.2).

Unfortunately, the fading and interference conditions under which WBANs have to operate are not well-understood: in contrast to the large body of empirical work that has been carried out in the WSN domain (e.g. [21, 23, 24, 29–31]), not many experimental studies have focused on the WBAN channel characteristics. However, "an understanding of radio propagation is essential for coming up with appropriate design, deployment, and management strategies for any wireless network" [180]. The first goal of this thesis is therefore to gain a better understanding of the WBAN fading and RF interference characteristics by experimentally analyzing the P_{Signal} and $P_{Interference}$ dynamics in different realistic WBAN scenarios. We focus on the 2.4 GHz

Chapter 4:	Channel measurements	⎫ Fading
Chapter 5:	Packet scheduling	
Chapter 6:	Channel measurements	⎫ RF Interference
Chapter 7:	Channel hopping &	
	Packet combining	

Table 3.1: Overview of the following technical chapters.

band, because for the foreseeable future this band is expected to continue playing a major role in the WBAN domain (Sect. 2.2.1). Based on the previous findings the second goal of this thesis is to investigate how communication reliability in WBANs can be improved by exploiting certain fading and RF interference characteristics. This analysis makes use of the three performance metrics introduced above. The topics we address are summarized in Table 3.1 and briefly explained in the following.

3.1.1 Analysis of Fading Characteristics

WBANs exhibit unique P_{Signal} dynamics (fading conditions), because "the movement of the human body has a dramatic effect on the strength of the received signal" [145]. The dynamics are influenced by the particular activity of the person, who is carrying the WBAN. In Chapter 4 we study the fading effects on different 2.4 GHz on-body communication links while the person is walking. Walking is an important activity in the context of WBANs as it is part of many medical, sports and personal fitness applications. We attach eight COTS WBAN nodes to a person and measure the fading effects on a millisecond-scale while the person is walking. We analyze the magnitude of P_{Signal} fluctuation, and in particular the regularity of P_{Signal} dynamics among subsequent gait cycles. Furthermore, we determine if the effects are influenced by changing the node position or attaching nodes to a different subject — we conduct experiments with five different subjects and consider eight different node locations. Our analysis shows that the magnitude of P_{Signal} fluctuation is often significant; it also reveals that while a person is walking P_{Signal} peaks follow a *regular* pattern. The effects that govern this are large-scale effects due to path-loss and shadowing, since small scale (multipath) effects are rather irregular.

3.1.2 Exploiting Periodic Fading Effects

In Chapter 5 we explore an approach that exploits the regularity of P_{Signal} oscillation within the WBAN communication protocol stack. The idea is to leverage signal strength dynamics by predicting when received signal strength is at (close to) its peak. Such packets will exhibit a higher SINR than packets sent at arbitrary points of time and thus render more reliable communication. To predict signal strength peaks our approach captures user mobility with the help of accelerometers. Through a large set of experiments involving multiple subjects, we show that when packets are transmitted at predicted signal strength peaks, *outdoors* unreliable links can often be turned into reliable links with PRR values well above 90 %. The gains in reliability

come at the cost of an increase in latency by on average one step period, but the additional energy spent by the devices is negligible.

3.1.3 Analysis of RF Interference Characteristics

WBANs often operate in unlicensed frequency bands and are therefore exposed to RF interference from other wireless technologies. To study the impact of realistic 2.4 GHz RF interference on WBAN communication, in Chapter 6 we introduce a custom measurement setup made of COTS WBAN nodes attached to five different body positions. Our setup enabled us to collect about half a billion RF noise samples in various urban environments. A statistical analysis of the dataset shows that — in comparison to the effects of fading — the overall impact of RF interference on WBAN link reliability is low. However, urban RF noise is about one order of magnitude more "bursty" (correlated in time) than random, i.e. uniformly distributed over time, RF noise. Therefore, even if average spectrum occupation is low, it may be necessary to apply RF interference mitigation techniques, in particular in applications with low latency bounds.

3.1.4 Mitigating the Effects of RF Interference

Several techniques to mitigate the effects of RF interference have been proposed in the past. In Chapter 7 we focus on two well-known strategies: we analyze to what extent *channel hopping* can reduce packet loss and in particular error burstiness, due to RF interference. We also analyze a *packet combining* approach that allows to recover an original packet from multiple copies that have been corrupted by RF interference. We find that, even when blacklisting notoriously busy channels, channel hopping can hardly beat a very simple strategy: static selection of IEEE 802.15.4 channel 26. We also show that *packet combining* can result in practical savings of 5-10 % of energy for intermediate (PRR below 0.9) WBAN links if only corrupted bits retransmitted.

3.2 Methods

Over the past years there has been a growing consensus in the WSN community that research involving wireless communication requires empirical evaluation. For example, the popular tree-routing protocol in TinyOS 2 has been evaluated in 12 different testbeds [32]. Empirical evaluation is necessary, because the propagation and interference characteristics are influenced by a large number of environment-specific parameters (layout and physical properties of objects, mobility of persons/objects, spectrum activity of co-located networks, etc.), which are hard to capture in models. This is especially true for WBANs, where the conditions are particularly dynamic and not well-understood. Therefore in this thesis the main method of evaluation is through experiments.

Our experiments are performed in realistic urban outdoor and indoor scenarios. To this end, we attach COTS WBAN hardware to various body positions of different subjects while they are performing common activities like walking along an urban shopping street. Our experiments can

51

Figure 3.2: WBAN platforms (from left to right): TMote Sky, Shimmer2 and Shimmer2r.

be grouped into two sets: one set of experiments measures the P_{Signal} and $P_{Interference}$ parameters to better understand the fading and interference conditions. A second set of experiments evaluates approaches that exploit certain fading and interference characteristics.

While this thesis puts a strong focus on experiments, it also relies on a few analytical and simulation models: for example, we derive a time-homogeneous discrete-time Markov chain model to estimate the performance gains of our packet combining algorithm before performing experiments (Sect. 7.3). And we formulate the task of finding a non-overlapping packet schedule as a Linear Programming (LP) problem (Sect. 5.4). These steps help us in properly designing our algorithms, while any performance evaluation of the algorithms is afterwards still performed through experiments. There is, however, one exception: we use the OMNET++/Castalia simulator [136] to evaluate the channel hopping algorithms described in Sect. 7.2. This is necessary, because the gains of a channel hopping algorithm are largely determined by the characteristic interference activity. In order to study these gains we need reproducible interference environments, but since we deal with uncontrollable urban environments reproducibility is not given. Therefore in this thesis channel hopping algorithms are evaluated in simulation, but the RF interference environment in the simulator is reproduced based on our $P_{Interference}$ data ("trace-driven simulation"). This allows us to compare channel hopping algorithms under identical RF interference conditions. The rest of this section describes our methods and a set of general assumptions in more detail.

3.2.1 Platforms and Node Positions

We consider a WBAN that consists of a set of body-worn nodes, which use IEEE 802.15.4 radios to communicate in the 2.4 GHz ISM band. IEEE 802.15.4 is one of the most mature and applied low-power communication standards and previous studies have shown its suitability for WBAN applications (Sect. 2.3.2). Our experiments are mainly conducted with the Shimmer2(r) platform [8], but a small number of experiments also involve the Tmote Sky platform [66]. The platforms are depicted in Fig. 3.2. They are quite similar in terms of hardware, as they use the

(a) Node positions   (b) Network topology   (c) Protocol stack
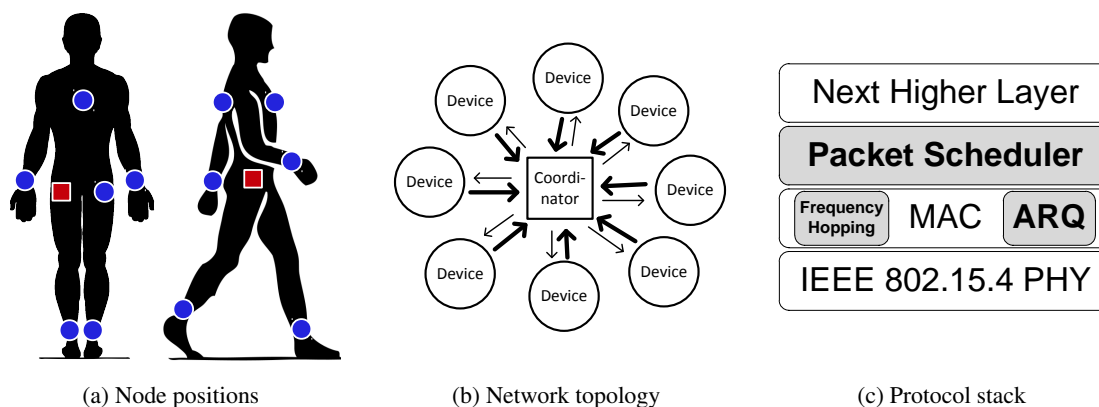
Figure 3.3: Properties of the WBAN under consideration.

same MCU (MSP430F1611 [68]) and IEEE 802.15.4-compliant radio (CC2420 [69]). However, Tmote Sky has an inverted-F microstrip antenna while Shimmer2r uses an Surface-Mount Device (SMD) antenna. Also, Shimmer2r has a smaller form factor and compact native enclosure and therefore the antenna-to-body surface distance is smaller than for the Tmote Sky, which also has a larger battery pack. The main reason why two different platforms were used is that the Tmote platform has been available for many years, but the Shimmer2 platform had become available only after the first study described in this thesis had been already conducted. Our experiments reveal, however, that although the platforms use the same radio chip, the wireless communication performance differs, which we attribute to different antennas and form factor (Sect. 6.2). The two platforms are described in more detail in Sect. 2.1.3.

In our experiments we use a maximum of 8 nodes that are attached to the human body at the positions shown in Fig. 3.3a: the wrist position is commonly used for measuring blood oxygen saturation, EMG, galvanic skin response or blood pressure (sometimes the electrodes are attached to a finger, but the node itself typically is located at the wrist) [181, 182]. The ankles are one of the positions for attaching motion sensors [183, 184]. The position at the back has been used to measure walking/running speed or detecting behavioral patterns [185, 186] and the chest is the preferred position to measure ECG or heart rate [72, 187]. Finally, the position next to the hips are suitable for attaching motion sensors to detect activity or posture changes [183, 188]. All nodes that are not carried in the pocket are strapped tightly to the body, either using adhesive tape or a sweatband.

In our configuration the nodes described so far all take the role of IEEE 802.15.4 *devices*. There is one additional node that takes the role of an IEEE 802.15.4 *coordinator*. This node is located in the right trouser pocket. We assume that in a real application this node is integrated in the user's cell phone, where the data may be analyzed or transmitted via the Internet to a remote server for storage, processing and analysis. The cell phone is, however, not part of our setup/investigation, because we are only interested in the intra-WBAN communication.

53

3.2.2 Network Topology and Traffic Pattern

In this thesis we assume that the coordinator represents the only node with whom all remaining WBAN devices communicate. The WBAN thus forms a star topology as depicted in Fig. 3.3b. As explained in Sect. 2.1.4 this is the most prevalent network topology in WBANs. The devices are responsible for data collection and transmission of the data to the coordinator, the application traffic is therefore uplink (devices → coordinator). The analysis in Sect. 2.1.3 has shown that many envisioned WBAN applications generate low data rates of no more than 150 byte per second (c.f. Table 2.1). Since a IEEE 802.15.4 packet has about 100 byte of application payload, we assume that a WBAN device generates no more than 1 or 2 packets per second. Since the majority of sensors shown in Table 2.1 are typically sampled periodically, we assume that application data is generated periodically.

3.2.3 Protocol Architecture

Apart from performing channel measurements in this thesis we analyze three mechanisms that are located either within the MAC layer or between MAC and next higher layer, respectively, as indicated by the dark gray boxes in Fig. 3.3c. One of our components represents a packet scheduler that can seamlessly be inserted between the MAC and the next higher layer. This component does not require a specific MAC protocol, it only requires that the next higher layer can control the transmission time of packets on the order of a few tens of milliseconds. Our second component defines a novel packet combining approach that extends the retransmission policy within a stop-and-wait ARQ scheme. This component requires modifications to a MAC's ARQ mechanism and tight control over the PHY (radio chip) during packet reception. It therefore needs to be integrated within the MAC protocol. Finally, we investigate a standard channel hopping technique for its suitability to mitigate the effect of RF interference in WBANs. This component must be able to control the selection of the radio channel for a radio event (packet reception/transmission) and is therefore integrated within the MAC. The three components are independent of each other and do not rely on a particular MAC protocol as long as the above stated requirements are met. Also, our components do not make any assumptions about the next higher layer (above the MAC). In our system the IEEE 802.15.4-2006 MAC nonbeacon-enabled mode (c.f. Sect. 2.3.2) is selected as the reference MAC protocol, because it is well-known and fulfills the above requirements. In practice, we use the *TKN15.4* implementation described in Sect. 2.3.3.

3.2.4 Experiment Scenarios

We use two sets of scenarios to conduct our experiments: the first set involves a person walking outdoors on the sidewalk of an urban street or indoors along a corridor of an office building. The related experiments are used to study the fading characteristics in a WBAN while the subject is physically active (Chapter 4) and to investigate the trade-offs between communication reliability and latency (Chapter 5). The experiments involve between 3 and 7 different subjects.

The second set of scenarios includes 12 different urban environments, in which the person carrying the WBAN is performing common activities. The environments allow for a number of

representative human activities as the experiments are conducted in: a park, a residential area (4x), a shopping street (4x), a university campus (2x), and an urban transportation system. In some of the scenarios the person is mobile (e.g. walking), in others the person is rather static, e.g. sitting at home in front of a television. Each experiment last on average around 60 minutes. The data are used to study the characteristics of urban RF interference (Chapter 6) and the results are also used in a simulator to investigate interference mitigation techniques (Chapter 7).

3.2.5 Channel Measurement Parameters

The main performance metrics used in this thesis have been introduced in the beginning of this chapter. However, these metrics are not applicable when performing channel measurements (Chapter 4 and Chapter 6). The goal of our channel measurements is to capture P_{Signal} and $P_{Interference}$ to study fading and external RF interference, respectively. Unfortunately in practice with our hardware we cannot directly measure P_{Signal}, instead we use Received Signal Strength Indication (RSSI) provided by the radio hardware. RSSI represents the power of the received signal (in dBm) averaged over the last 8 symbol periods (128 ms) [69]. When a signal is received while no interfering signals are present RSSI is indeed identical to P_{Signal} (plus P_{Noise} introduced by the electronic components in the receiver circuitry — however, P_{Noise} is comparably low and can safely be ignored). When a signal is received while interfering signals are present RSSI represents the sum of P_{Signal} and $P_{Interference}$ (plus P_{Noise}). When we measure P_{Signal} characteristics we therefore verify that no significant amount of RF interference is present by using the upper frequencies in the 2.4 GHz which are not overlapping with any standard IEEE 802.15.4 or IEEE 802.11 channels, and by monitoring the RF environment with a spectrum analyzer during an experiment.

$P_{Interference}$ is the sum of the power of any interfering signals. In this thesis we focus on external RF interference (environmental RF noise), i.e. signals that are emitted by devices external to the WBAN. To this end we only need to measure RSSI while no WBAN device is transmitting a packet. Note that technically what we measure is therefore not interference (there is no WBAN communication that can be interfered at that moment), but in this case RSSI represents the environmental RF noise power averaged over the last 8 symbol periods (128 ms) [69], or — whenever no interfering signals are present — the noise floor of the radio. In order to eliminate the chances of measuring internal RF interference generated by our own WBAN devices we apply synchronization mechanisms (by introducing a dedicated wired control channel in our measurement setup).

In order to characterize the fading and RF interference effects, we design a set of additional measures on top of the "raw" radio parameters:

- As a first indicator of the variability of the received signal strength we use the Interquartile Range (IQR), i.e. the distance between 25th and 75th percentiles, of the P_{Signal} samples collected during an entire experiment. The IQR holds 50% of the data centered around the median, this interval is therefore helpful to understand the magnitude of fading effects over larger timescales (Section 4.3).

- We define a new measure P_{Signal}-*gain(w)*, which is applicable for periodic human activities such as walking. It describes the difference between the mean P_{Signal} per activity

period and the mean P_{Signal} calculated over a window with size `w`, centered at the signal strength peak of the same period. It thus allows to describe the signal fluctuation on a per-period basis. By analyzing the corresponding distributions we can, for example, understand the magnitude of signal peaks on a smaller time scale. We can also derive a practical upper bound for how much can be gained by transmitting packets at signal strength peaks (Section 4.4).

- In order to describe the regularity of fading effects during periodic activities, we introduce a technique that allows to detect whether P_{Signal} peaks consistently occur at fixed offsets within the activity period, which is detail in Section 4.4.2.

- To quantify the amount of external RF interference in a WBAN, we use the empirical probability of $P_{Interference}$ being larger or equal than a threshold `THR`, i.e. `P (P`$_{\texttt{Interference}}$ `> THR)`. For example, we can set `THR` to a value of 5 dB below the sensitivity threshold of the radio, in this case only those interfering signals that may practically harm WBAN communication are considered (Section 6.3).

- We introduce a metric that captures the correlation of RF interference activity in the frequency domain: we use the conditional probability of $P_{Interference}$ being larger or equal than a threshold `THR`, given that the $P_{Interference}$ of a neighbouring channel (± 1 channel) was larger or equal than that threshold at the same time, i.e.
`P (P`$_{\texttt{Interference}}$$\geq$`THR | P`$_{\texttt{Interference@Ch-1}}$`>THR or P`$_{\texttt{Interference@Ch+1}}$`>THR)`,
where `P`$_{\texttt{Interference@Ch-1}}$ and `P`$_{\texttt{Interference@Ch+1}}$ represent the interference value measured on the channel above or below the channel on which `P`$_{\texttt{Interference}}$ was measured, respectively (Section 6.3).

- Finally, we introduce two additional metrics that capture the correlation of RF interference in the time and space domain. Analogous to the previous metric, these metrics make use of conditional probabilities (Section 6.3).

3.2.6 Assumptions and Limitations

We assume that the coordinator is less energy-constrained than the remaining WBAN devices, because it is integrated in the user's cell phone and may access the battery of the phone. The coordinator is also less constrained in computational and memory resources, because some of its tasks may be executed on the CPU of the cell phone. One central design goal is therefore to push the resource-expensive tasks in terms of energy, computation and memory to the WBAN coordinator and minimize the resource usage on the WBAN devices. This approach extends the lifetime of the devices and thus the overall lifetime of the WBAN. Furthermore, although adaptive transmission power schemes are a promising way of leveraging excessive link budget [189] (which some of the mechanisms analyzed in this thesis can generate), these techniques remain out of scope of this thesis. Instead, we assume that the transmission power is fixed, but we consider different values between 0 dBm and -25 dBm, which is in line with the IEEE 802.15.4 standard which states that "a transmitter shall be capable of transmitting at least -3 dBm. De-

vices should transmit lower power when possible in order to reduce interference to other devices and systems." [18, Sect. 6.9.5].

There are naturally a few limitations to this thesis. First of all, we focus on a single communication technology and frequency band. Signal propagation characteristics change when the communication frequency is altered, and the use of a different physical layer such as UWB may considerably impact communication reliability. Nonetheless we believe that our findings in Chapter 4 and 5 are also applicable for other narrowband 2.4 GHz physical layers, because they are primarily based on the on-body propagation characteristics of 2.4 GHz signals. This also includes the newly allocated MBAN bands, which are located just below the 2.4 GHz ISM band (c.f. Table 2.4) and can be expected to have similar propagation characteristics. On the other hand, the occurrence of RF interference is largely determined by the particular wireless systems occupying the spectrum. While the results in Chapter 6 are likely not representative for other bands, we believe that the experimental setup and measurement methods are still a valuable contribution. The interference mitigation algorithms studied in Chapter 7 are, however, also applicable in other frequency bands, especially our packet combining approach is not limited to the 2.4 GHz interference characteristics, but also applicable in other contexts. Finally, a natural limitation of conducting experiments is uncertainty due to limited data. Within our possibilities we have tried to reduce this uncertainty: we have covered a large set of scenarios, node positions and involved different subjects in our experiments.

Chapter 4

The Wireless On-Body Channel During Walking

In this chapter we study the fading effects on 2.4 GHz on-body communication links while the person is walking. Walking is an important activity in the context of WBANs as it is part of many medical, sports and personal fitness applications. Daily routines also involve significant amount of walking: according to Harvard Medical School on average adults take 6,000 to 7,000 steps daily [190], which corresponds to about 1 hour of walking per day (obviously there is a large variation over different types of people and professions). 10,000 steps per day is often considered a healthful level of physical activity and for children values between 11,000 and 13,000 steps per day are common (for details refer to [191]).

To characterize the on-body channel during walking the traditional methods that rely on fading distributions (Section 2.2.4) are ill-suited. During regular activities such as walking the distance, the amount of body shadowing and relative antenna orientation between a WBAN transmitter and receiver pair change *quasi-periodically* — periodic on a small scale of a few (tens of) seconds but less predictable at larger scale. The signal strength of an intra-WBAN communication link is therefore also expected to exhibit some kind of temporal correlation. Previous studies have, however, rarely focused their analysis on the time domain, but rather quantified the magnitude of body shadowing due to certain posture changes or analyzed summary statistics like fading distributions derived over long time intervals (several minutes or hours). An assessment of the fading effects on a smaller time scale (milliseconds) as a function of the user's movements is, however, important: it allows to better understand the channel conditions and adapt the corresponding mechanism in the WBAN communication protocol stack. On the other, it is more challenging to develop a measurement setup on top of commercial WBAN hardware that has the required fidelity.

The questions this chapter addresses are: how does walking affect the received signal strength pattern on a millisecond-scale and is this effect influenced by changing the node position or attaching nodes to a different subject? What is the magnitude of signal peaks and how regular do the peaks occur within a stride period (a stride period denotes the time from initial contact of one foot to the following initial contact of the same foot)? In our analysis we put an extra focus on the stride period, because it corresponds to the fundamental frequency of the user activity.

To answer the above questions we designed a measurement setup that allows us to capture both, the human walking pattern (with the help of motion sensors) and the low-level metric P_{Signal} defined in Section 3.2.6 with high resolution. This allowed us to analyze the temporal behavior of P_{Signal} for several subjects and node positions in relation to their movement. Note that the methods and analysis have also been published in the following two articles: [38, 192].

4.1   Related Work

Several studies have focused on the effects of human body shadowing on RF communication. One common method is to evaluate the attenuation caused by people crossing a communication link between a stationary transmitter-receiver pair: for example, Kara and Bertoni examined the effects of body shadowing caused by people crossing a 2.4 GHz band link between an RF signal generator and spectrum analyzer and showed that the human body can result in attenuation of up to 20 dB [193]. Similar shadowing effects have been reported in [194] for 802.11 radios, and the authors also conclude that the influence of body shadowing is higher than other factors such as the type of ground. In the 10 GHz band the attenuation caused by human obstruction has been reported to reach 25 dB, depending on the distance between human body and transmitter [195]. At higher frequencies the effects of human shadowing are typically stronger, for example, a study comparing the 900 MHz and 60 GHz bands shows a difference of approximately 10 dB [196]. Bahillo et al. have come to a similar conclusion when comparing the cellular Global System for Mobile Communications (GSM)-900 band with transmissions at 100 MHz and 1800 MHz [197].

While these studies are based on a stationary transmitter-receiver pair, other authors have investigated scenarios where at least one node is carried on the body (all following references related to the 2.4 GHz band): for example, Miluzzo et al. focused on IEEE 802.15.4 person-to-person communication and showed that not only can the body result in severe attenuation, but the placement of the node on the body also "has a significant effect on the performance of the communications system" [116]. For on-body communication (sender and receiver are attached to the same person) others have investigated the impact of the distance between antenna and body surface and shown that a difference of 4.5 cm can translate in signal strength variations of up to 20 dB [110]. Once communication alternates between LOS and NLOS, for example when the person changes the posture, the signal strength for on-body propagation typically experiences significant variations (> 10 dB [198]). On-body channel measurements performed by D'Errico and Ouvry indicate that the slow fading component is to a large extent influenced by the amount of body shadowing [199] and propose a relaying strategy to mitigate channel time variations [200]. Miniutti et al. conducted an extensive measurement campaign in an office environment and showed that the on-body channel is mainly characterized by slow, flat fading [201]. Depending on the subject's activity the variation in path loss ranged from 7.1 dB (standing still) up to 46.8 dB (running). Path loss was particularly dynamic when communication alternated between LOS and NLOS and the variation was "repetitive and consistent with the speed of movement of the test subject" [201].

The last study indicates that regular human movements may manifest in quasi-periodic on-body signal strength behavior. The main limitation of the previous studies is, however, that this relationship is not explored in detail. The measurement setups used in these studies allowed to

capture on-body channel characteristics with high fidelity, but they typically did not allow to precisely track the human movement. In the following we present an experimental setup that allows to measure both, human movement (with the help of motion sensors) as well as on-body signal strength fluctuation (with the help of COTS WBAN radios).

4.2 Setup and Scenario

Our WBAN setup consists of eight Shimmer2r [8] nodes, which integrate the Texas Instruments MSP430 MCU and the IEEE 802.15.4-compliant CC2420 transceiver [69] (Section 2.1.3). In all experiments the eight nodes were positioned on the body as described in Chapter 3 and depicted in Figure 3.3a. One node had the role of the transmitter, the other nodes set their radios to receive mode. The transmitter was located in the right trouser pocket of the subject, because in real life this node might be integrated in the user's cell phone, which is typically carried in the pocket. The receiver nodes were placed on the left and right ankle, in the left trouser pocket, on the left and right wrist, in the center of the chest, and in the center of the back. Nodes were strapped tightly to the body, either using adhesive tape or a sweatband resulting in a body surface-to-antenna separation of about $1\,cm$.

Instead of exchanging packets the transmitter was emitting a continuous IEEE 802.15.4 modulated carrier signal using a transmission power of $0\,dBm$. Emitting a continuous carrier signal is one of the testing features of the CC2420 radio, it allows to send random data in an infinite transmission [69]. The other seven nodes were passively measuring P_{Signal} (RSSI) at a sampling frequency of $1000\,Hz$.[1] In addition to emitting a continuous carrier signal, the transmitter node periodically measured its 3-axial accelerometer with a sampling frequency of $200\,Hz$. In each experiment the subject that carried the WBAN was continuously walking outdoors on the sidewalk of an urban street, i.e. a typical pedestrian footpath in a downtown urban area. Subjects were walking at regular walking speed of approximately $100\,steps/min$. Five different male subjects aged between 23 and 35 years each performed an experiment of approximately 15 minutes and collected a total of about 30 million P_{Signal} samples.

We took careful precaution to minimize the influence from external RF interference on our setup. Our transmission frequency was set to $2482\,MHz$, which is at the upper end of the $2.4\,GHz$ ISM band. This frequency does not overlap with any of the standard IEEE 802.15.4 channels and it is also spaced sufficiently far apart from the highest available IEEE 802.11 center frequency, which was WLAN channel 13 at center frequency $2472\,MHz$. Before performing the experiments, we verified with a USB spectrum analyzer attached to a laptop that this part of the spectrum was not used on our path. During each experiment a subject was carrying an additional *sniffer* node which was sampling $P_{Interference}$ at a channel $5\,MHz$ below the coordinator's transmission frequency. Since IEEE 802.15.4 has a channel width of $2\,MHz$ the *sniffer* could not detect energy from our transmitter, but it would detect energy from other interfering

---

[1]As explained in Section 3.2 we assume that application data is sent in the device→ coordinator direction. In this section we measure signal strength only on the coordinator→ device link, because measuring P_{Signal} in the reverse direction with $1000\,Hz$ would require tight (out-of-band) time synchronization between the devices and a very short (not IEEE 802.15.4-compliant) packet format. Instead we assume that links are reciprocal, i.e. instantaneous signal strength is comparable in both directions.

Figure 4.1: Ten seconds of the acceleration (top) and RSSI signal (others) while subject 1 was walking outdoors.

sources, such as a strong IEEE 802.11 sender on channel 13. After each experiment we checked the data from the sniffer, which confirmed that normally less than 0.2 % of the samples were significantly high, i.e. larger or equal than the radio sensitivity threshold of -94 dBm (only in one experiment the sniffer reported 0.85 % RF noise). Thus, we believe that the impact of external RF interference on our results is negligible.

4.3 Signal Strength Variation

Figure 4.1 and Figure 4.2 show a 10-second snapshot of P_{Signal} during the experiments performed by subject 1 and subject 2, respectively. The figure also shows the 200 Hz acceleration time series collected on the transmitter node located in the right trouser pocket (top graph, respectively). It can be seen that the acceleration graphs have a period of about 1.2 s, which matches the *stride period* of the subject. The P_{Signal} time series at the left and right wrist resemble a noisy sine curve with a peak-to-peak amplitude of nearly 30 dB and 15 dB, respectively. Interestingly, the "left wrist" position of subject 2 (Figure 4.2) shows two peaks within one stride period, which likely correspond to the position in front and at the back of the torso. The P_{Signal} graphs for the other node positions are noisier and have a smaller, but still substantial, range. All P_{Signal} graphs exhibit a period identical to the subject's stride period as can be seen by the solid gray lines overlaid on the individual P_{Signal} graphs: these lines represents a simple (unweighted) moving average over a 500 ms window. The smoothed curves for the left

Figure 4.2: Ten seconds of the acceleration (top) and RSSI signal (others) while subject 2 was walking outdoors.

and right wrist have an offset of roughly 1/2 period, and the smoothed signals captured on the left pocket and chest are almost in phase. The shape of the curves (number of peaks and range) as well as the amount of noise is an individual characteristic of the subject. But for all subjects the period of the smoothed P_{Signal} signal generally matched the stride period extracted from the acceleration signal. The P_{Signal} of the nodes located on the back and in the left pocket, however, showed smaller fluctuations and it was also noisier than other nodes.[2]

Figure 4.3 summarizes the results of two 15-minute experiments performed by the two subjects in a boxplot, respectively. The edges of the boxes represent the 25th and 75th percentiles and the central red mark corresponds to the median P_{Signal} value. A box thus represents the IQR which holds 50 % of the data. It can be seen that there there is often a large fluctuation (IQR) of more than 10 dB for the nodes on the extremities and a smaller spread for the nodes on the chest or in the left pocket. The node located on the back, however, shows little variance and many readings are close to the noise floor. The results for the other three subjects (not shown) are similar.

As explained in Section 2.2.4 the on-body wireless channel is often modeled with the help of statistical distributions that describe the signal amplitude. Various authors have derived WBAN fading distributions, typically by determining the best fit out of well-known distributions like Lognormal or Weibull with their empirical data [122]. For comparison we have determined the parameter estimates for different fading distributions based on MLE for our measurements. The

---

[2]Note that the noisiness of the signal is also influenced by how well nodes are fixed to the body - recall that in our case they were tightly attached using adhesive tape or a sweatband.
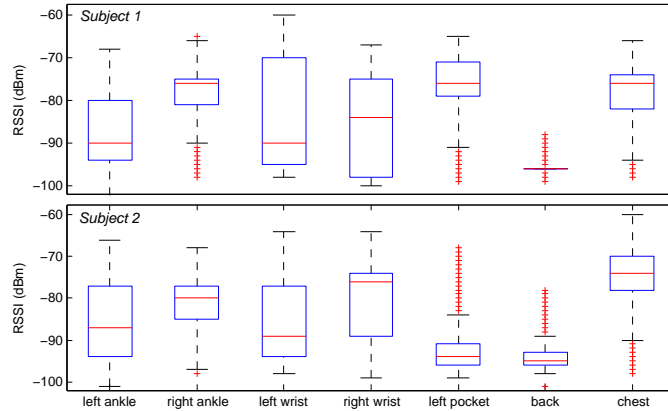
Figure 4.3: Summary of the P_{Signal} fluctuation captured during a 15-minute walking experiment for two subjects.

Position	Lognormal (σ)	Weibull (λ, k)	Gamma (k, θ)	Mean (dBm)
left ankle	1.69	2.41,0.59	0.47,8.16	-87.80
right ankle	1.47	1.86,1.11	0.99,1.82	-78.88
left wrist	2.77	4.20,0.37	0.26,61.03	-83.85
right wrist	2.48	3.44,0.45	0.33,23.81	-85.39
left pocket	1.34	1.88,0.91	0.87,2.25	-76.01
back	0.19	1.11,3.31	25.52,0.04	-95.97
chest	1.51	1.97,0.91	0.82,2.50	-78.67

Table 4.1: MLE parameter estimates for different fading distributions based on our outdoor walking experiments.

results for one representative subject are summarized in Table 4.1.[3] Since the fading distributions abstract from temporal correlation in the signal, we will, however, not consider them any further in this thesis.

4.4 Signal Strength Peaks

Figure 4.1 suggests that when a person is walking P_{Signal} features some kind of regularity and Figure 4.3 indicates that for several node positions the magnitude of P_{Signal} fluctuation is quite significant. In this section we analyze this more systematically. But rather than applying techniques like auto-correlation, LCR or ADF (Section 2.2.4) on the whole dataset we partition the data into intervals corresponding to the fundamental frequency of the user's movement pattern, i.e. the stride period. Furthermore, we focus our analysis on evaluating how a real system might

---

[3]Note that for the analysis we have normalized the data by removing the local mean and report it in a separate column — this implies that μ parameter for lognormal distribution is zero.

Figure 4.4: Magnification of the "left ankle" graph from Figure 4.1.

benefit from the aforementioned effects. The largest gains can generally be expected when packets are transmitted while P_{Signal} is high relative to the link's average P_{Signal}. In this case packets are more likely to arrive without errors as they have a larger margin towards the radio's sensitivity threshold, and they are less susceptible to environmental RF interference. In other words, the SINR of such packets is higher and the transmission will more likely be successful. We therefore focus our analysis on the properties of the P_{Signal} peaks. Note that at this stage the basis for our analysis is the P_{Signal} metric (we will focus on higher layer metrics in Chapter 5).

In the following we address two questions: how much can practically be gained by transmitting packets at P_{Signal} peaks rather than at random points in time? And how regular are P_{Signal} peaks, or stated differently: given some knowledge about the past behavior of P_{Signal}, how well can the time of future P_{Signal} peaks be estimated? In this analysis we assume perfect knowledge about the walking pattern, i.e. the beginning and duration of the last stride is always known (we derive this information from the acceleration signal in our traces using a technique explained in Chapter 5.3).

4.4.1 Magnitude

To get an indication of the potential P_{Signal} gain that can be achieved by sending packets at P_{Signal} peaks during a stride, we compared the mean P_{Signal} per stride period to an unweighted moving average of P_{Signal} computed over the same interval. We call the maximum difference `P`$_{Signal}$`gain(w)`. This concept is illustrated in Figure 4.4. `P`$_{Signal}$`gain(w)` thus represents the average increase in P_{Signal} we can expect if a system is able to send a packet (uniformly distributed) within an interval of $\pm w/2$ centered at the time of the P_{Signal} peak of the smoothed signal, where w is the size of the moving average window. We consider a smoothed signal rather than the actual P_{Signal}, because in practice a system will hardly be able to accurately predict the time of the P_{Signal} peak within the upcoming stride period: P_{Signal} in a WBAN may be noisy and the stride period may change, because the subject slows down or speeds up. Furthermore,
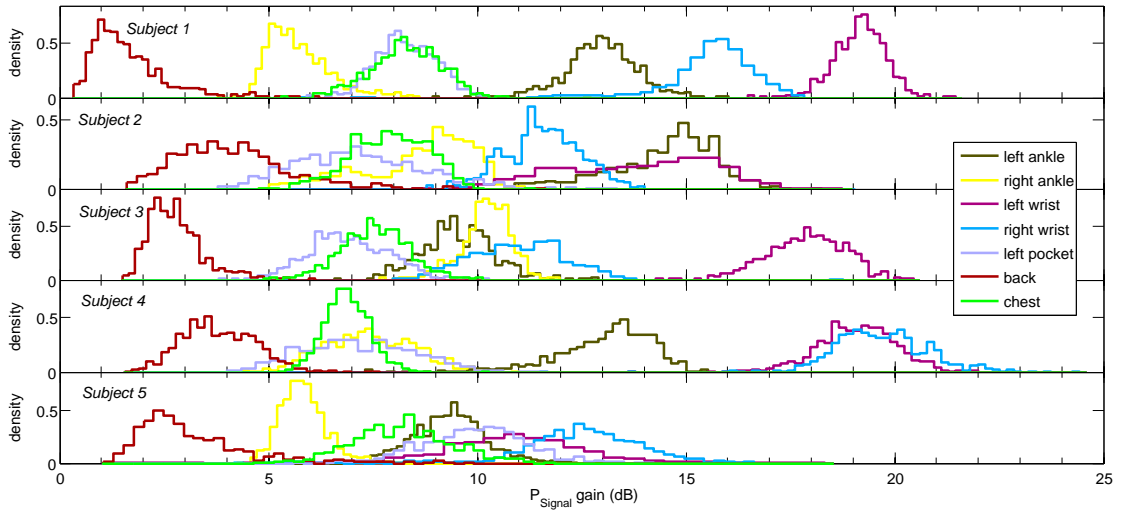
Figure 4.5: Normalized histograms of P_{Signal} gain if a packet is sent within $\pm 25\,\text{ms}$ of the P_{Signal} peak.

packet transmission time may not be controlled precisely.[4] And a packet transmission will take some time, for example, the airtime of maximum-sized IEEE 802.15.4 DATA frame (133-byte PHY Protocol Data Unit (PPDU)) is 4.256 ms, the airtime of an Acknowledgement Message (ACK) frame is $352\mu s$ and the radio turnaround time is typically $192\mu s$. Therefore, practically P_{Signal} should be stable at high level for at least 15 ms.

In short, a future P_{Signal} peak may not be reached exactly, but only within an interval of $\pm w/2$. To account for this fact we explored two values, $w = 500, \text{ms}$ and $w = 50\,\text{ms}$. We first segmented our traces into stride periods (we used the algorithm described in Chapter 5 to derive the stride period from the acceleration signal). Afterwards for each node and each experiment we calculated the per-stride $P_{Signal}\texttt{gain(w)}$. The corresponding distribution represents the increase in P_{Signal} we expect to experience if a packet transmission had occurred at a time $\pm 25\,\text{ms}$ and $\pm 250\,\text{ms}$ (uniformly distributed) around the P_{Signal} peak, respectively.

Figure 4.5 and Figure 4.6 show the results: each graph represents a histogram of P_{Signal} gain for one node position. The histograms are normalized to display the relative frequencies of $P_{Signal}\texttt{gain(w)}$, with the total area equaling 1, respectively. The nodes located on the wrists and the left ankle show the highest gain, between 10 dB and 20 dB for $w = 50\,\text{ms}$ and between 5 dB and 15 dB for $w = 500\,\text{ms}$. For $w = 50\,\text{ms}$ even the locations chest, right ankle and left pocket show an average $P_{Signal}\texttt{gain(w)}$ of at least 5 dB (at $w = 500\,\text{ms}$, however, the gain for these positions is often below 5 dB). The node located on the back of the subject always shows the least $P_{Signal}\texttt{gain(w)}$, because signal strength hardly ever exceeds the radio's noise floor. The results indicate that when packets are transmitted within $\pm 25\,\text{ms}$ of the P_{Signal} peak, the $P_{Signal}\texttt{gain(w)}$ is often substantial. This gain could, for example, be leveraged by reducing radio output power: on our platform a reduction from 0 dBm to -10 dBm

---

[4]CSMA protocols like the IEEE 802.15.4 MAC typically allow the next higher layer to control the transmission time quite accurately, on the order of 10 ms.

Figure 4.6: Normalized histograms of P_{Signal} gain if packet is sent within $\pm 250\,\text{ms}$ of the P_{Signal} peak.

translates into transmitter energy savings of $35\,\%$ [69]. Alternatively, the gain might translate into better packet reception rates (which we explore in Chapter 5).

4.4.2 Regularity

Our notion of regularity is related to the time of the P_{Signal} peak within the stride period. Regularity is present if the occurrence of the P_{Signal} peaks follows a certain "rule". Theoretically this "rule" may be arbitrary complex. In the following we propose a simple algorithm that predicts P_{Signal} peaks based on a set of past P_{Signal} observations and the duration of the current stride period. By showing that this algorithm can indeed consistently estimate P_{Signal} values significantly above a link's average P_{Signal} we show that P_{Signal} peaks are indeed regular.

Our algorithm estimates a future P_{Signal} peak by adding (integer multiples of) the current stride period duration to the time of the last-known P_{Signal} peak.[5] One option to approximate the latter is to let the sender transmit a chain of probing packets or emit a continuous carrier signal for the duration of a stride period. During this time other devices could sample P_{Signal} and afterwards derive the peak using a maximum identification algorithm. However, this approach is rather inefficient, because it consumes a significant amount of energy as the radios need to be enabled for an entire stride period. Instead, we assume that in a practical system N observations are made by sampling P_{Signal} quasi-regularly with a frequency

$$\frac{1}{p\left(\frac{1}{N} + 1\right)} \tag{4.1}$$

where p is the person's stride period. For each P_{Signal} sample we convert the sampling time to the current offset within the subject's stride period as illustrated in Figure 4.7 (in this

---

[5]We explain in Chapter 5 how we derive the stride period from the acceleration signal.
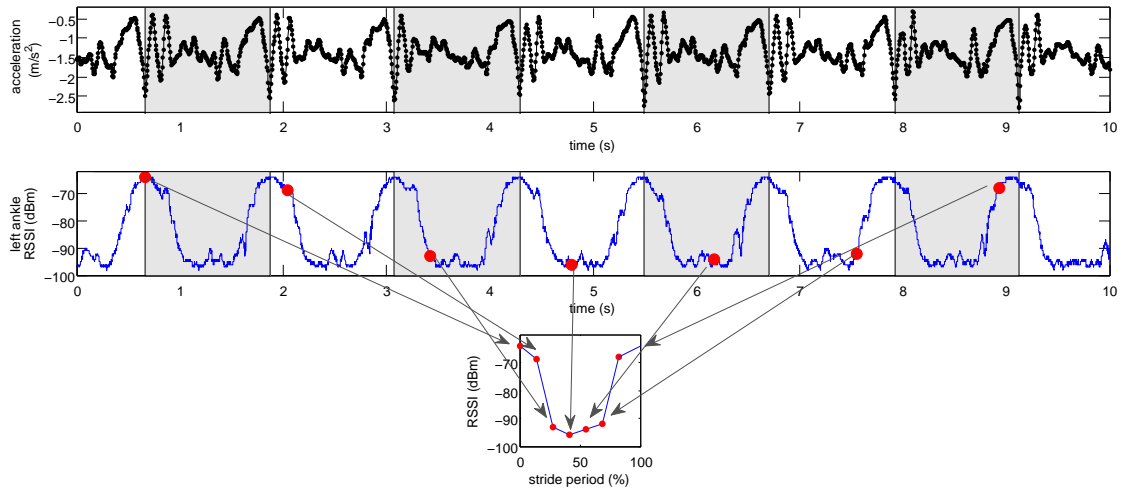
Figure 4.7: Approximating P_{Signal} by mapping samples onto the stride period interval.

example the subject's walking speed was about $1.2 \, \mathrm{m/s}$). This assumes that the subject's walking pattern is quasi-periodic and the stride period and phase offset are known. As the number of samples increases a more comprehensive representation of P_{Signal} per stride (Figure 4.7 bottom) may be approximated. The P_{Signal} peak then simply corresponds to the maximum value in the approximated signal. Since time is expressed relative to the duration of the stride period the approach adapts to de/acceleration in the subject's step frequency.[6]

The number of samples used to approximate the shape (and thus peak) is a parameter that likely influences the prediction performance of our algorithms. To analyze this relationship we processed the traces from our five subjects as follows: we considered $N = 1..30$ consecutive strides to extract N P_{Signal} samples using the quasi-uniform sampling process described previously and visualized in Figure 4.7.[7] From the N P_{Signal} samples the maximum was identified and the respective phase offset within the stride was used to predict the P_{Signal} peak in stride N+1. We then compared the mean P_{Signal} in a $\pm 25 \, \mathrm{ms}$ window around the predicted P_{Signal} peak time with the mean P_{Signal} in a $\pm 25 \, \mathrm{ms}$ window around the optimum P_{Signal} of stride N+1. Like in Section 4.4.1 a window of $50 \, \mathrm{ms}$ was selected to account for the fact that in a practical system the time of transmission may not be controlled precisely. We call the difference between the two values *estimation error*. Based on our traces we calculated the estimation error for $N = 1..30$ per node over all strides and all subjects.

Figure 4.8 shows the mean estimation error for one subject. With a single P_{Signal} sample the error may be as large as $18 \, \mathrm{dB}$, because the sample is uniformly distributed over stride period

---

[6]One question is what approximation (interpolation) function should be used to estimate the shape (and thus peak) of the comprehensive per-stride representation of P_{Signal} (Figure 4.7 bottom). We considered approximation of P_{Signal} by piecewise linear interpolation and by polynomials of different degrees using least squares fitting. However, once we discovered that polynomials up to degree 7 generally resulted in worse predictions (larger error, see below) we selected the piecewise linear interpolation method shown in Figure 4.7 (bottom). This reduces the task of detecting the P_{Signal} peak to finding the maximum P_{Signal} value in the sample space.

[7]Note that we added a small random offset (uniformly from $0..p/N$, where p = stride period duration) to each sample time to avoid systematic overlap with the beginning of the gait cycle.
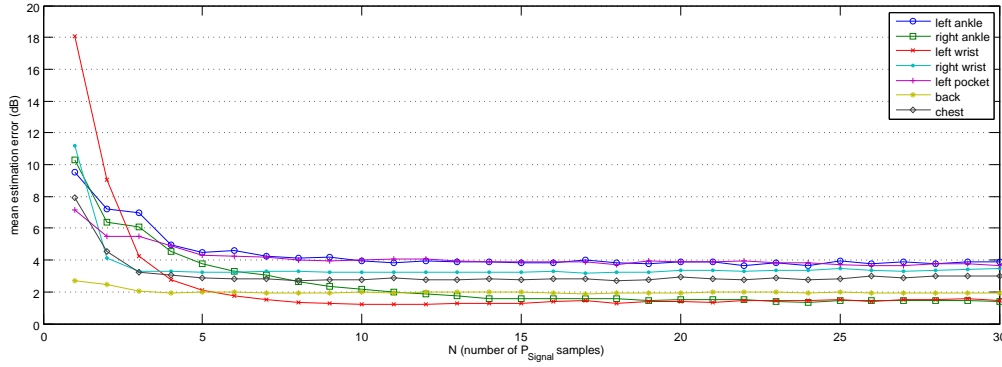
Figure 4.8: Mean error when predicting an P_{Signal} peak based on N P_{Signal} samples obtained with the technique shown in Figure 4.7.

Figure 4.9: Estimation error when predicting a P_{Signal} peak with 15 P_{Signal} samples.

N+1. With two samples the estimation error may still reach up to $9\,\mathrm{dB}$ and the variance is often quite large ($\sigma \approx 5\,\mathrm{dB}$, omitted from Figure 4.8). With more than 10 samples the error reduces to at most 4 dB with $\sigma \leq 3\,\mathrm{dB}$. Interestingly, increasing the number of P_{Signal} samples beyond 10 decreases the error only marginally; also the variance of the error does not decrease significantly. We attribute this to the fact that in two consecutive strides the human gait cycle never exhibits a perfect match. Small differences in limb movement or transient de/acceleration translate into non-identical P_{Signal} patterns. A prediction based on past observation therefore always include some residual error. We found that in our experiments this residual error was always reached after no more than 10-15 P_{Signal} samples independent of the particular subject and node position. *One important intermediate result is therefore that our algorithm needs to collect no more than 15 P_{Signal} samples before estimating future peaks.*

The estimation error based on 15 P_{Signal} samples for all subjects and node positions is shown in Figure 4.9. The mean estimation error, over all subjects and node positions, is $2.48\,\mathrm{dB}$, the corresponding standard deviation is $2.38\,\mathrm{dB}$. For some positions such as the left wrist position the error can be as low $0.6\,\mathrm{dB}$, while for the left pocket position the estimation errors are often as large as $5\,\mathrm{dB}$, because the corresponding P_{Signal} fluctuates quickly and is quite noisy.

Since the estimation error is defined relative to the optimum P_{Signal} gain, this does not allow
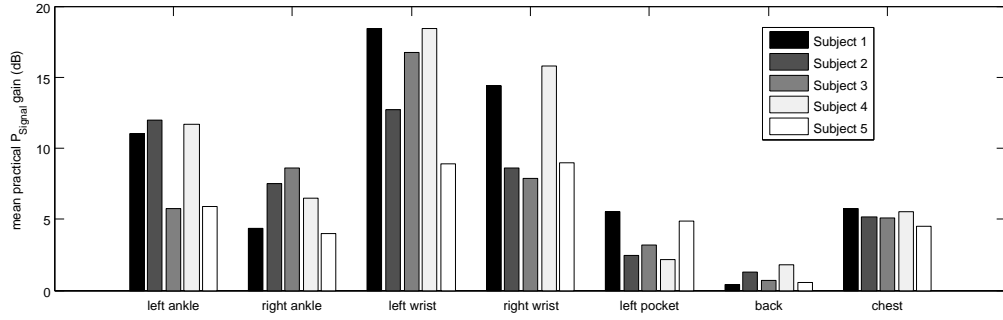
Figure 4.10: Difference between the optimum P_{Signal} gain and estimation error ("practical P_{Signal} gain").

to draw immediate conclusions about the *regularity* of P_{Signal} peaks. Recall that regularity means that the occurrence of a P_{Signal} peak — or at least P_{Signal} value significantly above a link's average P_{Signal} — follows a certain "rule". To this end we define a metric which represents the difference between the P_{Signal} peak predicted by our algorithm and the mean P_{Signal} of the corresponding stride and call this metric "practical P_{Signal} gain". Figure 4.10 shows the expected practical P_{Signal} gain for all subjects and node positions (combining the results shown in Figure 4.5 and Figure 4.9). The values range between 0.4 dB and 18.4 dB with a mean of 7.35 dB. For all subjects the node located at the back of the subject shows values of around 1 dB only, because at this position the P_{Signal} signal is usually close to the P_{Signal} noise floor. However, most node positions render substantial gains above 5 dB. These results indicate that while a person is walking P_{Signal} peaks indeed follow a *regular* pattern. The effects that govern this are large-scale effects due to path-loss and shadowing, since small scale (multipath) effects are typically rather irregular.

4.5   Summary

In this chapter we have investigated the fading effects on a 2.4 GHz on-body communication link while the person is walking. Several factors are responsible for the large P_{Signal} fluctuations we have measured:

- changes in relative node distance affecting path-loss,

- the amount of shadowing introduced by the body and the alternation between LOS and NLOS,

- changes in antenna-to-body surface distance (c.f. [110]),

- multipath fading (scattering, reflections from ground or body surface, diffraction on body parts),

- relative node orientation, which in conjunction with the imperfect antenna radiation pattern may result in different signal strength levels at the same distance.

Our evaluation showed that while a person is walking the magnitude of P_{Signal} fluctuation is position-dependent, but often significant. With the sender located in the right trouser pocket, we observed P_{Signal} IQRs ranging from 0-2 dB for the node located on the back of the subject, to 10 dB for the ankles up to 25 dB for a wrist node. These results match previous results obtained with other mote platforms [192] [202].

In a further analysis we partitioned our dataset into intervals that match the fundamental frequency of the user's movement pattern, i.e. stride period. Our analysis showed that the magnitude of P_{Signal} changes during a stride is similar to the fluctuation we had observed on larger time scale. It also revealed that while a person is walking P_{Signal} oscillates and peaks follow a *regular* pattern. The effects that govern this are large-scale effects due to path-loss and shadowing, since small scale (multipath) effects are rather irregular. The regularity implies that past observations allow to estimate future peaks (assuming the user continues the activity). Our analysis showed that 15 uniformly-spaced P_{Signal} samples are typically sufficient to predict a P_{Signal} peak and more samples do not decrease the estimation error. There may be a small residual estimation error, but with this approach substantial P_{Signal} gain of on average 7.5 dB can be expected in practice. An exception is the node located on the back, whose P_{Signal} was typically too close to the radio noise floor to render any practical gains.

From the results it is, however, not immediately clear how P_{Signal} fluctuations affect higher-level metrics like packet reception rate (PRR). Figure 4.3 indicates that during a stride period the absolute P_{Signal} values may drop to the noise floor of the radio, where even a small change in P_{Signal} translates into a substantial difference in packet delivery performance: "a variation in RSSI as small as 1.5 dB can change a good link to a bad one if the link is operating near the noise floor" [29]. Therefore, if a system is able to transmit packets close to the P_{Signal} peaks a significant increase in packet delivery performance may be possible. In the following Chapter 5 we explore this question in detail.

Chapter 5

Opportunistic Packet Scheduling

The previous chapter has shown that while a person is walking the received signal strength on a 2.4 GHz on-body communication link often fluctuates significantly. Instantaneous received signal strength may be noisy, but the peaks of the smoothed signal often follow a *regular* pattern that corresponds to the fundamental frequency of the user's movement pattern, i.e. the stride period. In this chapter we explore an approach that exploits this phenomenon within the WBAN communication protocol stack. The idea is to leverage signal strength fluctuations by transmitting packets when received signal strength is at (close to) its peak. These packets will exhibit a higher SINR than packets sent at arbitrary points of time and thus more reliable communication can be achieved.

In the following, after discussing related work, we introduce the software architecture of our packet scheduler. Afterwards we describe how we track the movement pattern of the person with the help of accelerometers. In combination with the algorithm described in Section 4.4.2 this information allows us to estimate received signal strength peaks per device $\rightarrow$ coordinator link. In order to avoid packet collision among different links within the same WBAN we formulate the task of finding a non-overlapping packet schedule as a LP problem and show that it can be solved on a sensor node within short time (<100 ms). Finally we explain the design and implementation of our packet scheduler and present results from an extensive experimental evaluation.

5.1 Related Work

Sensor network communication protocols usually do not take human mobility directly into account. For example, state-of-the-art sensor network MAC protocols may utilize sophisticated probing of the communication peer (such as A-MAC [203]) or adjust to dynamic external RF interference environments (such as BuzzBuzz [204]), but they are not designed to exploit the effects of regular channel fluctuations. Rather, these MAC protocols are often configured with constant values for the listening (sender-initiated Low-Power Listening (LPL) protocols) or probing (receiver-initiated Low-Power Probing (LPP) protocols) periods. As explained in Section 2.3.1 also WBAN MAC protocols typically do not consider periodic movement patterns; instead a common design goal is to support different types of WBAN traffic or packet priorities by (dynamically) structuring time into different contention-free or contention-based intervals and map-

ping certain traffic classes onto them. These time slots are, however, allocated independent of the channel conditions (Section 2.3.1).

Opportunistic scheduling exploits the time-varying channel conditions of communication links. Cellular Orthogonal Frequency Division Multiplexing (OFDM) systems are an example for opportunistically scheduling users to frequencies and time slots. Regarding the spatial dimension, opportunistic selection of the next node on a multi-hop path can mitigate the effects of varying channel conditions [205]. In this chapter, however, we focus exclusively on the time domain. In this context opportunistic schemes have been proposed for sensor networks: for example, some protocols estimate coherence time and exploit link burstiness by sending packets back-to-back as long as channel conditions are good [206]; others chose routes that minimize worst case burst lengths [207]. While these approaches exploit temporal correlation in link quality, they do not directly target periodicities caused by, for example, human mobility.

Mobility has been exploited in the domain of delay tolerant networks, where continuous end-to-end connectivity cannot be assumed. For example, Jain et al. have proposed a system where mobile nodes act as forwarding agents as they move through a field of sensors [208]. *Bartendr* [209] is a recent proposal that builds on the observation that the variations in cellular signal strength can result in significant short-term differences in energy consumption. The authors suggest to defer mobile phone data communication while the signal is weak and demonstrate significant performance gains. The approach is similar to the one we present in this chapter: like our approach, but unlike many other opportunistic schemes, the channel state is *predicted* rather than reacted upon. Bartendr tries to forecast time intervals that yield high P_{Signal} values based on the user's location and direction of travel (e.g. relative to the cell phone tower). It works on a larger time scale than our approach and it is based on the characteristics of the cellular handover process as it assumes communication between a mobile peer and a stationary basestation.

The general idea of leveraging excessive link budget in WBANs has been the focus of studies on adaptive power control in WBANs [210–212]. Their goal is to reduce the transmission power and thus energy consumption of on-body communication in order to maximize network lifetime. After probing the communication link any excessive link budget is leveraged through the reduction of output power. Since the proposals do not specifically target periodic changes, this requires either continuous probing and adaptation with high control overhead, or may result in sub-optimal power allocation and packet loss if path-loss variations are not detected in time. Others have investigated how power control can reduce RF interference: Kazemi et al. [213] propose a machine-learning approach to discover the RF environment and derive the best policy to reduce energy consumption. Their approach is evaluated in simulation based on the IEEE 802.15.6 channel model, which shows the algorithm's potential to increase network lifetime. Smith et al. [214] propose to couple a channel prediction algorithm that considers the partial periodicity of a WBAN channel with adaptive power control. Their evaluation is based on empirical data from several WBAN channel sounder experiments and indicates that their scheme can improve reliability and power consumption in comparison to transmission policies without channel prediction.

The approach we describe in this chapter is different: we perform *context-based prediction of the WBAN channel state*. The prediction is based on "out-of-band" (with respect to communication) information from a tri-axial accelerometer. This is effective, because the acceleration
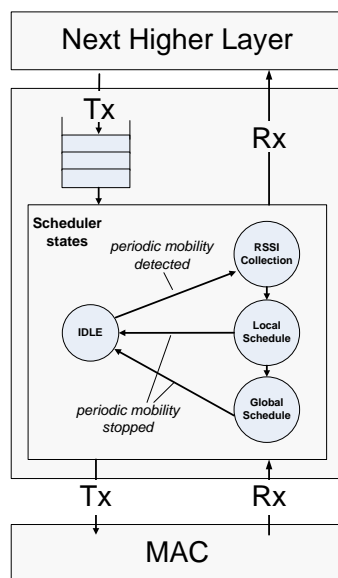
Figure 5.1: Scheduler architecture.

signal quickly reflects human movements, which in turn have significant impact on the state of the WBAN channel. Accelerometers are low-cost and have a very small form-factor, therefore they are often available on many WBAN platforms (Section 2.1.2). Furthermore, the sampling process consumes only a very small amount of energy.

5.2 Scheduler Architecture

The task of a scheduler is to coordinate access to a shared resource. This may involve different dimensions such as the time or frequency dimension. In our case the shared resource is the wireless channel and to coordinate means to decide when packets are transmitted, i.e. in this chapter we focus on the time domain. The goal of our scheduler is to control the packet transmission time per device $\rightarrow$ coordinator link so that packets arrive when received signal strength is at (or close to) the peak.[1]

Traditionally, a packet scheduler is part of a communication protocol, e.g. integrated within the MAC protocol. This ensures tight control over the transmission process. However, this integration makes the design inflexible: the scheduler cannot be easily transferred to a system that uses a different MAC protocol. Therefore our scheduler represents a small software component that is located between the MAC and next higher layer protocol as shown in Figure 5.1. The advantage is that our scheduler implementation works on top of different MAC protocols — on the other hand, we require that the packet processing jitter within the MAC is no more than about 50 ms as explained below.

---

[1]Recall that our network model consists of a set of devices connected via star topology to a central coordinator node. The coordinator takes the role of the sink (receiver) and the devices represent the data sources (senders).

On the transmit path our scheduler component intercepts packets from the next higher layer to queue and possibly defer their transmission to some later time. On the receive path it snoops on received packets and extracts metadata statistics (RSSI and timestamps) before they are forwarded to the next higher layer. Our scheduler can be realized in different variants: nodes may either make all decisions regarding packet scheduling autonomously, based on their local observations. Or nodes cooperate and exchange information (like the current stride period), to avoid redundant signal processing or to agree on a global collision-free packet schedule. In the latter case nodes need to exchange some small amount of control traffic, which may either be piggybacked on existing traffic (e.g. beacons, DATA packets or ACKs) or transmitted as separate control packets.

There are four main scheduler states as shown in Figure 5.1. Initially and whenever no periodic mobility has been detected the scheduler is in *Idle* state. In this state the scheduler is inactive and its queue is always empty. Once periodic mobility is detected (either locally or remotely and signalled via control information), the scheduler leaves the *Idle* state and enters the *RSSI Collection* phase in which it collects RSSI information to capture signal strength fluctuations and derive the P_{Signal} peak. The previous Section 4.4.2 showed that about 15 RSSI samples (uniformly-spaced in time) are sufficient to perform this task. Once 15 RSSI samples have been collected, the scheduler may enter the *Local Schedule* state. In this (optional) state the scheduler computes the optimal transmission time based on the locally available RSSI information. Packets on the transmit path are now delayed until a point in time that promises higher than average signal strength from a local perspective. Finally, the scheduler may enter the *Global Schedule* state. This state assumes that one entity (typically the coordinator) has obtained RSSI information for each device $\rightarrow$ coordinator link and calculated a non-overlapping packet schedule for all devices.

Recall that our scheduler is only one of three components in our overall system model (Figure 3.3c) and that the scheduler is decoupled from the other two components, which will be introduced in later chapters of this thesis. However, the scheduler must be able to queue packets and influence the packet transmission time. To this end we assume that the transmit primitive of the MAC layer may introduce some jitter before a packet is transmitted, but we require it to be reasonably low (around 50 ms based on our observations in Section 4.4.1). We believe in practice this is often feasible. For example, in our implementation we use the IEEE 802.15.4 MAC and configure the number of retransmissions and CCA backoffs to zero, which reduces the packet transmit operation to 2.56 ms plus a few ms CPU processing time.

The architecture shown in Figure 5.1 is still quite generic. In Section 5.5 we explain how we instantiate this architecture in a concrete implementation. But beforehand we describe the tasks that are performed by the scheduler in more detail.

5.3 Gait Monitoring

The state transitions within our scheduler are related to the mobility of the subject (Figure 5.1). While a subject is mobile our scheduler must be able to track the stride period and its phase, because packet scheduling is performed relative to the beginning of a stride (recall that the stride period denotes the time from initial contact of one foot to the following initial contact of the

Figure 5.2: The acceleration signal while walking outdoors and the corresponding AMDF.

same foot). In the following we explain how our scheduler derives this information from an accelerometer signal.

5.3.1 Mobility and Stride Period Detection

Several algorithms for human stride frequency detection based on an accelerometer signal have been proposed and evaluated in the past [215, 216]. We adopted one that is based on an ACF analysis, because in comparison with other approaches the computational complexity is acceptable [216]. However, inspired by previous work on pitch detection in human speech, and in order to further reduce computational complexity, we use a variation of the ACF called Average Magnitude Difference Function (AMDF) [217], which, in contrast to the ACF, has no multiply operation in the summation. The AMDF is defined as:

$$R(\tau) = \frac{1}{L} \sum_{j=1}^{L} |S_j - S_{j-\tau}| \tag{5.1}$$

where L is the length of the segment of the sampled acceleration signal (within the portion of a sufficiently large analysis window), S_j are the accelerometer samples, $\tau = \tau_{min}, \tau_{min} + 1, \ldots \tau_{max}$ is the delay value and τ_{min} and τ_{max} represent the minimum/maximum delay shift.

We select $\tau_{min} = 200\,\mathrm{ms}$ and $\tau_{max} = L = 2\,\mathrm{s}$ as the minimum detectable stride period, i.e. these values represent the upper and lower stride periods we consider, respectively. Before we calculate the AMDF we always apply a moving average filter with a window of $100\,\mathrm{ms}$ to smooth the signal.[2] From the resulting AMDF the stride frequency can then be derived by locating the minimum. For example, Figure 5.2 shows the AMDF computed over the x-axial acceleration signal. Here the AMDF shows a local minimum at about $1.2\,\mathrm{s}$, which our algorithm would select as the current stride period. We use two criteria to decide if the there is relevant

---

[2]We considered other filters, such as a 2nd order Butterworth lowpass filter but rejected it when our implementation revealed that it is computationally too expensive for our hardware.

periodic behavior at all: before performing the AMDF computation the range of the acceleration time series must be above a certain threshold (depending on hardware characteristics such as accelerometer sensitivity) and the AMDF minimum must be smaller than 50 % of the AMDF maximum. Both criteria were selected via trial-and-error in experiments with several subjects (further details on the parameterization for our platform are reported below).

5.3.2   Stride Period and Phase Tracking

Once mobility has been detected and the initial stride period has been calculated with the technique described in the previous section a peak in received signal strength has to be found. In this chapter we use the term Opportune Transmission Window (OTW) to describe a time interval that yields high P_{Signal} values relative to the average P_{Signal} of a WBAN link, for example, the time interval of 0.1 s to 0.4 s for the left wrist position in Fig. 4.1. Intuitively, an OTW corresponds to a certain time interval within the human gait cycle, for example, a period of temporary LOS connection between left foot and right pocket when the foot has swung to the front.

To measure the first OTW we adopt the technique introduced in Section 4.4.2, i.e. we periodically obtain RSSI samples by exchanging packets with a rate defined in Equation 4.1 (this concept is illustrated in Figure 4.7). Note that rather than exchanging dedicated control packets our scheduler uses application traffic to obtain RSSI samples (the details are explained in Section 5.5). In order to perform this task (and afterwards be able to estimate future RSSI peaks) the phase offset within the subject's stride (or relative to some other reference point within the gait cycle) has to be known. Tracking the gait cycle over longer time based using periodic computation of the AMDF is error-prone, because its precision is bounded by the acceleration sampling frequency and an analysis has shown that small errors can accumulate to unacceptable drift within a few tens of seconds. Instead, we suggest to track the phase offset and stride duration by performing periodic pattern matching between the current acceleration signal and the acceleration signal obtained during a previous stride. The technique described in the previous Section 5.3.1 is thus only used to detect relevant mobility and then bootstrap the pattern matching algorithm described in the following with the initial stride period value.

An inspiring class of solutions from the area of speech processing is based on Dynamic Time Warping (DTW). DTW is an algorithm that evaluates how similar two sequences are, assuming that the sequences can vary in speed and accelerations and decelerations may occur. Given two sequences, R (of length n) and Q (of length m), the DTW algorithm constructs an n-by-m matrix, where each element (i,j) represents the Euclidean distance d(i,j) between the two points r_i and q_j. A warping path is a set of matrix elements that defines a mapping between R and Q given some constraints are met (the explanation on DTW follows [218]):

- *Boundary conditions*: the warping path starts at element (1,1) and finishes at the opposite corner of the matrix at element (n,m)

- *Continuity*: the path consists of adjacent cells (including diagonally adjacent cells)

- *Monotonicity*: the path must be monotonic in time, i.e. if element (i_2,j_2) follows element (i_1,j_1) then $i_2 - i_1 \geq 0$ and $j_2 - j_1 \geq 0$.
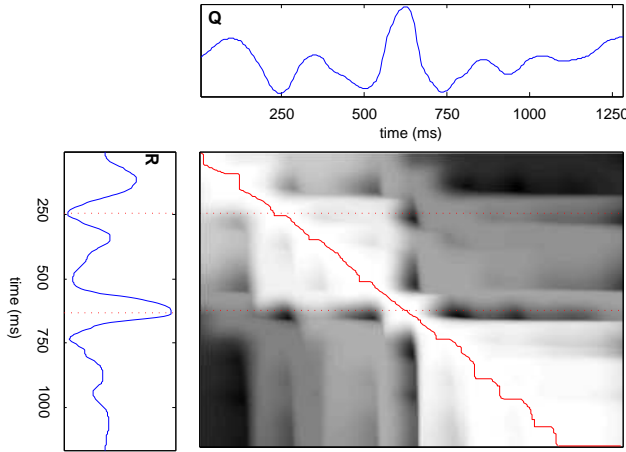
Figure 5.3: An example distance matrix and warping path (red solid line) for two smoothed acceleration sequences. Dashed red lines represent two anchor positions.
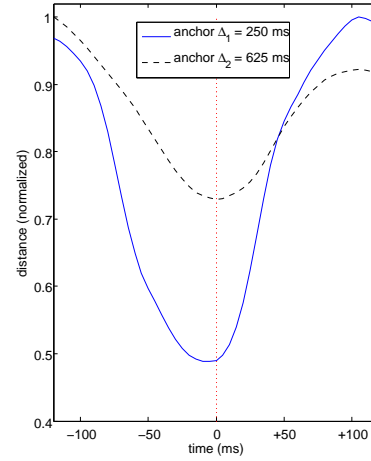
Figure 5.4: Distance (normalized) between the two anchors in R shown in Figure 5.3 and corresponding points in Q ± 100 ms.

From all possible warping paths $W = w_1, w_2, ..., w_k, ..., w_K$ with $max(m, n) \leq K \leq m + n + 1$ the path with the least cost is sought:

$$DTW(R, Q) = min \left\{ \sqrt{\sum_{k=1}^{K} w_k / K} \right. \tag{5.2}$$

To find this path dynamic programming can be used to evaluate a recursive function that defines the cumulative distance $\gamma(i, j)$ as the distance d(i,j) found in the current elements plus the minimum of the cumulative distances of the valid adjacent elements:

$$\gamma(i, j) = d(r_i, q_j) + min \left\{ \gamma(i - 1, j - 1), \gamma(i - 1, j), \gamma(i, j - 1) \right\} \tag{5.3}$$

An example of a cumulative distance matrix for two acceleration sequences is shown in Figure 5.3. Our task could be solved as follows: let the reference sequence R represent the acceleration signal captured during the initial stride period computed with the technique described in Section 5.3.1. The sequence R thus represents the acceleration signal for one full gait cycle. Assume that Q represents the acceleration signal immediately following sequence R. For any sample in Q the warping path represents a mapping into the reference sequence R (the warping path does not represent a bijection, i.e. some interpolation may be necessary). The path can thus be used to determine the current phase offset with respect to the gait cycle. The point in time where the warping path first reaches the last row of the matrix defines the beginning of the next acceleration sequence Q. It also corresponds to the current stride period. By continuously

sampling subsequent versions of Q we can thus always map the current time onto the gait cycle. To account for gradual changes in walking speed we may also replace R by Q and thus achieve better adaptability to de/acceleration. However, DTW has a time complexity of $\mathcal{O}(nm)$ and requires $n \cdot m$ matrix elements, which exceeds the RAM resources of our nodes by about a factor ten.

Instead of computing the full matrix we suggest to consider only a set of N anchor (extreme) points in the reference sequence R at associated time $\Delta_1, \Delta_2, ...\Delta_N$, e.g. the two extrema at $\Delta_1 = 250\,\text{ms}$ and $\Delta_2 = 625\,\text{ms}$ in graph R in Figure 5.3. We propose to select only those extrema in R which have large distances to the adjacent extrema (e.g. at least 10 % of the stride period). Given an anchor in R at Δ_i we extract from Q the data in a small interval around Δ_i, e.g. $Q[\Delta_i \pm 100\,\text{ms}]$. For each sample in this interval we then compute the Euclidean distance to the respective anchor value in R. In the resulting vector we localize the minimum value. If the minimum occurs at the center of the vector the stride period has not changed. Otherwise the offset to the center of the vector represents the current increase or decrease in the stride period.

Figure 5.3 shows this approach on an example: we have selected the two extrema at $\Delta_1 = 250\,\text{ms}$ and $\Delta_2 = 625\,\text{ms}$ of R in Figure 5.3 as anchors. We use a window of $\pm 100\,\text{ms}$ and thus compute the Euclidean distance between $Q[150..350]$ and the first anchor value $R[\Delta_1]$; and between $Q[525..725]$ and the second anchor value $R[\Delta_2]$. As shown in Figure 5.4 the vector resulting from the first operation has a minimum at about -5 ms, indicating a slight increase in walking speed, i.e. decrease in stride period. The second vector has a minimum at its center, indicating that the stride period has not changed with respect to R.

Initially, we (arbitrarily) select one anchor in R and consider it the beginning of the gait cycle (although it may not represent the actual beginning of the person's stride). We predict the next anchor by adding to the anchor in R the current stride period. We then compute the distance between the anchor in R and a certain interval around the predicted next anchor. We adjust the predicted anchor in time such that we achieve minimum distance to the anchor in R (Figure 5.4). The time difference between the two anchors is then considered the new stride period. Finally, we replace the old anchor in R by the new predicted anchor and let the process repeat after the next stride period. Whenever we need to map the current time onto the gait cycle we count the time elapsed since the beginning of the current gait cycle (last anchor) and divide it by the current stride period. Since we continuously update the current stride period our approach adapts to de/acceleration in walking speed. The approach will, however, fail if the increase or decrease in speed per stride period is larger than the (implementation-specific) interval we place around the predicted anchor; or, in case the acceleration signal is distorted by sudden movements or signal artifacts. Our solution is to use additional (intermediate) anchors to increase the algorithm's robustness, e.g. by applying a majority vote over the results from a set of anchors. Finally, we use the following criteria to decide when to stop tracking and return to idle state (Section 5.3.1): (a) one (or more) distance vectors (Figure 5.4) do not have a minimum within the considered interval, or (b) the range of the values in Q drops below a certain threshold (depending on hardware characteristics such as accelerometer sensitivity).

The acceleration sampling rate, the type of filter applied to the acceleration signal, the number of anchor points and the size of the window sampling are implementation-dependent parameters. Based on experiments with our hardware platform we found the following configuration
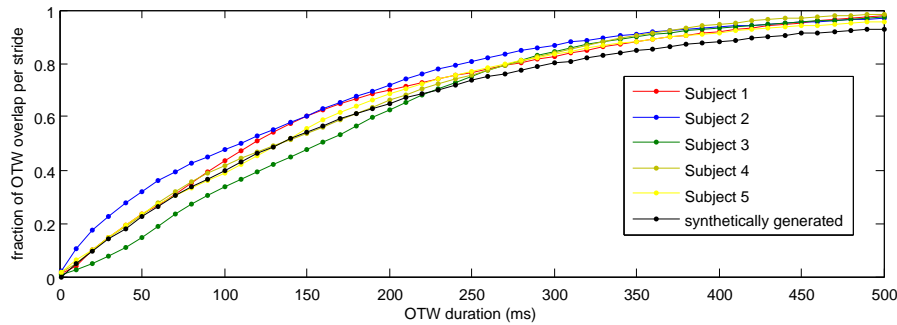
Figure 5.5: Fraction of OTW overlap per stride period when OTWs are computed on a per-link basis.

suitable: 200 Hz acceleration sampling frequency, applying a moving average filter with a window size of 100 ms, 1 anchor per stride period, and an acceleration window size corresponding to 30 % of the stride period.[3] This configuration was used in the analysis reported in Section 4.4 and in our implementation described in Section 5.5.

5.4 Non-Overlapping Packet Scheduling

We have so far ignored the case that the OTWs of different WBAN links may overlap in time. Since OTWs are defined relative to the gait cycle, if the OTWs belonging to two different devices have a similar phase offset (as nodes "left ankle" and "right wrist" in Figure 4.1), packets from these devices will *always* have a high chance to collide. In the following we analyze this problem and propose a global packet scheduling algorithm that avoids intra-WBAN packet collisions.

5.4.1 Overlap Analysis

In the evaluation reported in Section 4.4.2 we had used our traces to compute the time that our algorithm would select as OTW per stride and node position. As explained in Section 4.2 before each experiment the nodes had exchanged control packets to achieve clock synchronization. By comparing the predicted OTWs for the 7 different node positions we are thus able to analyze to what degree OTWs overlapped in time.

The amount of overlap depends on the duration of the OTWs: for an OTW lasting 10 ms the fraction of overlap in our experiments was around 3 % as shown in Figure 5.5. At 50 ms OTW duration the value increases to 15 % and when OTWs are lasting 250 ms we can expect 75 % overlap (recall that the results are based on 7 receiver nodes). We also compared the results with an abstract scenario in which the packet transmission time for 7 nodes are uniformly distributed over the interval of a 1200 ms stride period (graph "synthetically generated" in Figure 5.5). Interestingly, the amount of collision matches the amount of packet collisions in our system; the gait seems to naturally spread the signal strength amplitudes over the stride period. There is,

---

[3]Note that in our implementation the acceleration signal does therefore not need to sampled continuously, but only during a ± 15 % interval around the next predicted anchor.

however, an important difference: in a system with randomized transmissions, collisions will (on average) affect all devices equally. Instead, in our system packet collision will continuously affect the same devices: since OTWs are defined relative to a global event (gait cycle), if two OTWs overlap once, they will *always* overlap. In this case the corresponding nodes will continuously suffer from a collision probability that is significantly higher than in a system with random transmission time. Therefore, in this section we propose a mechanism that resolves overlapping OTWs.

5.4.2 Centralized Packet Scheduling

We formulate the task of finding a non-overlapping packet schedule based on the RSSI samples that nodes have obtained during the RSSI collection phase as a LP problem. To this end we assume that stride periods are divided into equally-sized time slots. We also assume that a packet transmission (DATA+ACK) fits within the boundaries of a single slot. Let N be the number of WBAN devices, and let T be the number of slots available per stride period. Let W_n be the number of slots per stride period that should be reserved exclusively to node n, depending on its individual bandwidth requirement. Let x_{nt} be a binary variable that defines if node n ($1 \leq n \leq N$) may use slot t ($1 \leq t \leq T$) for its packet transmission or not:

$$x_{nt} = \begin{cases} 1 & \text{if node n may transmit a packet in slot t} \\ 0 & \text{otherwise} \end{cases} \tag{5.4}$$

x_{nt} are the decision variables of our LP program. Let r_{nt} be the RSSI value that node n has measured in slot t during the RSSI collection phase and note that in the following we use r_{nt} as dimensionless quantity. We obtain the values for those slots which have not been sampled during the probing phase by means of (linear) interpolation. We define the objective function of our linear programming problem as follows:

$$\text{Maximize} \sum_{n=1}^{N} \sum_{t=1}^{T} f(r_{nt}) x_{nt} \tag{5.5}$$

$f(x)$ is a function that converts r_{nt} to a positive value (e.g. $f(x) = x + 100$, assuming -99 is the lowest practical RSSI value), otherwise the solution is an empty schedule. There are two sets of constraints in our LP program. Every node should be able to transmit all its packets during the stride period, thus

$$\sum_{t=1}^{T} x_{nt} \leq W_n \quad (n = 1 \cdots N), \tag{5.6}$$

and we do not consider schedules in which any two transmissions overlap:

$$\sum_{n=1}^{N} x_{nt} \leq 1 \quad (t = 1 \cdots T). \tag{5.7}$$

Note that with $f(x) = x + 100$ our objective function maximizes cumulative RSSI, independent if links are strong (high RSSI) or weak (low RSSI). Yet, our primary optimization goal may be to achieve high reliability (high PRR) on the weak links, because strong links (high RSSI) tend to be more reliable in the first place [219]. Only as a secondary goal we may want to maximize RSSI among the strong links. A weak link should then be assigned a slot that improves the RSSI of the link, even if the same slot would yield bigger relative improvements for a stronger link. Let Δ be the boundary between weak and strong links, for example, Δ may be set to -85 (corresponding to a threshold of -85 dBm). Let P_{min} represent the minimum RSSI value of a packet, for example, we may set P_{min} to -100 (note that P_{min}, Δ and r_{nt} are all used as dimensionless quantity and we assume $r_{nt} < 0$). We may then define $f(x)$ as follows:

$$f(x) = \begin{cases} |P_{min}| & \text{if } x = P_{min} \\ f(x-1) + |x| & \text{if } P_{min} < x \leq \Delta \\ f(|\Delta|) + x - \Delta & \text{otherwise} \end{cases} \tag{5.8}$$

The first part of the equation turns any RSSI reading that equals P_{min} into its absolute value. This ensures that the objective function coefficients can be used in a maximum problem. The second and third part of Equation 5.8 make sure that any improvement of a weak link ($< \Delta$) will yield larger cumulative gain in the objective function Equation 5.5 than any improvement of a stronger link ($\geq \Delta$). Note that Equation 5.8 may in practice be realized by a small (precomputed) lookup table and in particular Equation 5.8 does not change the quality of our problem.

The constraint matrix of our LP problem, i.e. the coefficients on the left side of the two constraint set, has dimension $(N+T) \times (NT)$ and represents a classical assignment problem. It has the nice combinatorial property of being *totally unimodular*. To see this, note that a matrix is totally unimodular when ([220]):

1. all coefficients $aij \in \{+1, -1, 0\}$,

2. each column contains at most two nonzero coefficients, and

3. there exists a partition of the set of rows such that (i) if a column has two entries of the same sign their rows are in different sets; and (ii) if a column has two entries of different sign their rows are in the same sets.

Our constraints use either 0 or 1 as coefficients. Furthermore, the left hand side of coefficients from Equation 5.6 result in a matrix with columns that have exactly one 1 as entry. Equation 5.7 results in a set of N identity matrices of dimension $T \times T$ joint horizontally. Thus, every column has exactly two entries with a value of 1 and a partition that satisfies above requirement lies between the two constraint sets. Therefore our matrix is totally unimodular.

	x_{11}	x_{12}	x_{13}	x_{21}	x_{22}	x_{23}	
s_1	1	1	1	0	0	0	W_1
s_2	0	0	0	1	1	1	W_2
s_3	1	0	0	1	0	0	1
s_4	0	(1)	0	0	1	0	1
s_5	0	0	1	0	0	1	1
	-32	-38	-31	-23	-18	-21	0

$$\Downarrow$$

	x_{11}	s_4	x_{13}	x_{21}	x_{22}	x_{23}	
s_1	1	-1	1	0	-1	0	W_1+1
s_2	0	0	0	1	1	1	W_2
s_3	1	0	0	1	0	0	1
x_{12}	0	1	0	0	1	0	1
s_5	0	0	1	0	0	1	1
	-32	38	-31	-23	20	-21	38

Table 5.1: Two simplex tableaux demonstrating the effect of a pivot operation around the circled entry for a configuration N=2, T=3, $r_{11} = -68, r_{12} = -62, r_{13} = -69, r_{21} = -77, r_{22} = -82, r_{23} = -79$.

Since the constraint matrix is totally unimodular our linear program has only integer solutions [220] and may be solved with the *simplex* algorithm [221], a popular method for solving LP problems. Before using the simplex algorithm a problem is first turned into standard form by adding for each inequality constraint a new *slack variable* s_i to turn the constraint into an equality constraint. Then a simplex tableau can be created. And afterwards the simplex method consists of applying a sequence of *pivot* operations, which change the linear program from one feasible solution to another feasible solution by substituting one of the independent with a dependent variable, until an optimum solution is found (if one exists).

An illustrative example of a simplex tableau for our LP problem is shown in Table 5.1 for a configuration N=2 nodes T=3 time slots. The examples assumes node 1 has sampled RSSI in the three slots with values -68 dBm, -62 dBm and -69 dBm; and node 2 has obtained values -77 dBm, -82 dBm and -79 dBm. Using $f(x) = x + 100$, the objective function is thus to maximize $32x_{11} + 38x_{12} + 31x_{13} + 23x_{21} + 18x_{22} + 21x_{23}$. After adding slack variables s_i the five constraints are: $s_n + \sum_{t=1}^{T} x_{nt} = W_n$ for $n = 1, 2$ and $s_{2+t} + \sum_{n=1}^{N} x_{nt} = 1$ for $t = 1, 2, 3$. The problem can be written in tableau form as shown at the left side of Table 5.1 (refer to [222] for a compact introduction to the simplex algorithm and tableau construction). The pivot operation consists of replacing one of the independent variables x_{nt} by the dependent variables s_i, and substituting the value of x_{nt} into the other equations. For example, in Table 5.1 we have replaced x_{12} by s_4. The rules for the pivot operation are [222]: the pivot element turns into its reciprocal; entries in the same row as the pivot element are divided by the pivot element; entries in the same column as the pivot element change sign and are divided by the pivot element; and from all other elements we subtract the value of the product of the corresponding entries in the same row and column as themselves and the pivot, divided by the pivot.

The selection of the pivot element is to some degree arbitrary. In our implementation we select the column that corresponds to the largest objective function coefficient, i.e. the smallest value in the last row of the tableau. If the value is not negative the algorithm has found the solution and terminates, assuming there exists a solution that is feasible and bounded. Otherwise, from all positive matrix entries in the same column we select the row that yields the lowest value in the vector corresponding to the right hand side of the equation system (minimum ratio test, refer to [222] for details on the update rules). At least one such entry must exist, otherwise the problem is unbounded feasible.[4]

In our scenario a typical configuration may be N=7 nodes and T=40 slots, resulting in a total of 13160 matrix entries. If an entry is represented by a signed 8-bit integer, the matrix' memory requirements would exceed the $10\,\text{kB}$ RAM available on our platform. However, pivot operations preserve the matrix' property of being totally unimodular [223]. Thus every entry in the constraint matrix can be encoded by only 2 bits. In this case the RAM requirements of the matrix in our example configuration reduces to a manageable 3290 byte.

The pivot operation may require updating a considerable number of matrix entries (in the worst case all matrix entries except one row). If the pivot operation rules listed above were applied iteratively to each matrix element the computational costs for the platforms we target ($8\,\text{MHz}$ 16-bit CPU) are prohibitive. However, since the simplex algorithm always selects a non-zero pivot element (which must therefore be 1 in our case) the rules stated above become simpler: entries in the same row as the pivot element remain unchanged; entries in the same column as the pivot element change sign; and for all other entries depending on whether the corresponding element in the pivot column has a value of 1, 0, or -1 the (corresponding entry in the) pivot-row is subtracted, the row remains unchanged, or the (corresponding entry in the) pivot-row is added, respectively. Practically, the last set of rules may still be rather inefficient, because CPUs typically do not natively support signed 2-bit integer operations. Therefore, we suggest to replace the constraint matrix by two matrices that represent the positive/negative sign of the matrix element as bits, respectively. As a consequence the set of addition/subtraction operations, can be expressed by a combination of basic boolean algebra operators, on conjunction $\wedge$, disjunction $\vee$ and complement function $\neg$, as explained in the following.

Let $x \in \{+1, -1, 0\}$ be a matrix element encoded by two bits, x_p and x_n, as follows:

$$x_p = \begin{cases} 1 & \text{if } sgn(x) = 1 \\ 0 & \text{otherwise} \end{cases} \qquad x_n = \begin{cases} 1 & \text{if } sgn(x) = -1 \\ 0 & \text{otherwise} \end{cases} \qquad (5.9)$$

where *sgn* is the sign function defined as

$$sgn(x) = \begin{cases} 1 & \text{if } x > 0 \\ 0 & \text{if } x = 0 \\ -1 & \text{if } x < 0 \end{cases} \qquad (5.10)$$

Let x be one of the matrix elements we would like to update with the pivot operation and let the pivot element be located at a different row and column than x. Let c be the matrix element

---

[4]Our problems are always feasible and bounded as long as the configuration is consistent.

located at the same column as the pivot element, and at the same row as x. Let r be the matrix element located at the same row as the pivot element, and at the same column as x. The pivot operation is identical to the following statement: if $c = 1$ then $x := x - r$; else if $c = -1$ then $x := x + r$; otherwise do nothing. This is equivalent to the following two assignments:

$$x_p := \begin{cases} (x_p \wedge \neg r_p) \vee (\neg x_n \wedge r_n) & \text{if } c_p = 1 \\ (x_p \vee r_p) \wedge (\neg x_n \wedge \neg r_n) & \text{if } c_n = 1 \\ x_p & \text{otherwise} \end{cases} \qquad (5.11)$$

$$x_n := \begin{cases} (x_n \wedge \neg r_n) \vee (\neg x_p \wedge r_p) & \text{if } c_p = 1 \\ (x_n \vee r_n) \wedge (\neg x_p \wedge \neg r_p) & \text{if } c_n = 1 \\ x_n & \text{otherwise} \end{cases} \qquad (5.12)$$

The equivalence is demonstrated by the truth tables shown in Table 8.1 and Table 8.2 in Appendix A. The advantage of using Equation 5.11 and 5.12 is that we can apply the pivot operation to as many entries in parallel as the CPU data registers are wide (16 on our platform). As a consequence our class of LP problems can be solved by a sensor node within $100 \, \text{ms}$ as reported in Section 5.5.2.

5.5 Implementation

In Section 5.2 we introduced our scheduler architecture and explained that the scheduler may be instantiated in different ways. For example, not all states shown in Figure 5.1 must be implemented on all devices: the coordinator may compute a global schedule based on the incoming application traffic, thus devices transition immediately from *Idle* to *Global Schedule* state. Alternatively, only a subset of devices might implement a scheduler variant that calculates a local schedule based on the RSSI in ACK packets and their own local accelerometer (note that such devices could seamlessly integrate into any CSMA-based WBAN MAC). Our suggestion is to assign most signal processing tasks to the coordinator, because we assume that the coordinator has more energy resources than the devices. For example, the coordinator may be integrated in the user's cell phone with access to a larger battery capacity; or it may have more resources allocated due to its prominent role of connecting the WBAN to the distribution network. Thus we envision the coordinator (sink node) to perform most of the communication and signal processing tasks and distribute the information in its network, e.g. via beacons or as ACK packet payload. Therefore only the coordinator is required to have an accelerometer, but none of the devices.

We call our reference design Centralized Opportunistic Packet Scheduler (COPS). COPS initially sets the scheduler components on all nodes to *Idle* state (Figure 5.1). Once the coordinator has detected periodic mobility it enters the *RSSI Collection* state, which is signalled via a customized ACK to the devices. The ACK payload includes the stride period and the beginning of the gait cycle, expressed relative to the packet transmission time. The scheduler instances on

the devices then also transition to the *RSSI Collection* state and start to control the DATA transmission time to achieve uniform transmissions over the stride period as shown in Figure 4.7. Once the coordinator has received enough DATA packets (RSSI samples) it starts computing the global schedule using our variant of the simplex algorithm described in Section 5.4.2. COPS splits a stride period into 40 equally-spaced slots and once the coordinator has received from a device at least one sample in every third slot it begins to compute the global schedule (this corresponds to at least 13 quasi-uniform samples, which rendered good results in Section 4.4.2). The global schedule is then distributed to the devices via the ACK payload and their scheduler instances enter the *Global Schedule* state. To account for lost ACK packets, a global schedule is sent to a device as long as the device has not set a flag in its DATA packet that acknowledges the adoption of the schedule. If the coordinator has not received at least 2 packets from a device during the RSSI collection phase, the device will not be assigned a slot in the global schedule (but use random transmission times). As long as the coordinator detects periodic mobility it continuously distributes the stride period and beginning of the gait cycle in its ACK packets. Once the coordinator detects the end of periodic mobility, it enters the *Idle* state and distributes this information to all devices by sending standard ACKs (without payload). When devices receive such an ACK (or when they have not received an ACKs for a certain number of DATA packets), their scheduler enters the *Idle* state.

We have implemented our reference design COPS in the TinyOS 2 [164] operating system. As MAC protocol we have selected the TinyOS 2 IEEE 802.15.4 MAC protocol [35, 36] in nonbeacon-enabled mode, because it represents a classical CSMA protocol. In order to have maximum control over the transmission time we configured the number of retransmissions and CCA backoffs to zero, which reduces the packet transmit operation to 2.56 ms plus a few milliseconds of CPU processing time. In our system the task of the MAC protocol essentially reduces to transmitting and receiving DATA/ACK packets: the devices send DATA packets and the coordinator radio is set to receive mode unless it transmits an ACK packet. We make one modification to the standard IEEE 802.15.4 DATA transmission primitive: we allow the next higher layer (our scheduler) to add a payload to ACK packets. The scheduler component on the coordinator node uses this primitive to insert the current stride period and beginning of the gait cycle (expressed relative to the ACK timestamp) into the ACK packets.

5.5.1 Scheduler Components

Our TinyOS 2 implementation consists of the following components:

- the scheduler core component that manages the transmit queue and controls the packet transmission time;

- a component responsible for sampling the acceleration signal at 200 Hz and applying a moving average filter with a window size of 100 ms;

- a component for AMDF calculation and acceleration pattern matching;

- our variant of the simplex algorithm described in Section 5.4.2;

- helper components for heap memory management and data logging on the SD card.

Component	Flash (B)	RAM (B)
Scheduler core	4682	192
Acceleration signal sampling and filtering	746	203
AMDF calculation	888	47
Acceleration signal pattern matching	4922	40
Management of RSSI probing results	4284	287
Simplex algorithm	1638	36
Memory management	48	6779

Table 5.2: Memory footprint of the main system components.

As explained above the acceleration signal processing and simplex algorithm components are only required on the coordinator node. The device implementation essentially consists of a scheduler core component. The implementation of the coordinator almost exhausts the 10 kB RAM resources of our platform. The reason is that multiple seconds of the 3-axial acceleration signal are kept in memory to compute the stride period based on the AMDF. In addition the simplex algorithm required a large portion of the RAM. However, since the AMDF component and simplex algorithm component were never active at the same time we implemented a simple heap manager that dynamically assigned RAM to these components as needed (TinyOS does not include heap management). The memory footprint of our main system components is shown in Table 5.2. In case the coordinator is integrated in a mobile cell phone, significantly more RAM will be available. The additional memory could be used to to optimize the gait monitoring process by storing the history of the acceleration signal. On the other hand, our current implementation only involves one operating system, which makes the approach easier to implement.

5.5.2   Simplex Algorithm

As explained in Section 5.4.2 for our class of problems the pivot operation of the simplex algorithm can be executed very efficiently using a combination of basic boolean algebra operators. As a consequence our LP problem can typically be solved by a sensor node below 100 ms. Figure 5.6 shows the results of a set of experiments we conducted with one Shimmer2 sensor node attached to a laptop. Over the USB connection we periodically injected random excerpts of the RSSI traces we had obtained in the measurements described in Chapter 4. The node used the data to construct the matrix coefficients, and then executed the simplex algorithm. As can be seen in Figure 5.6 the LP problem was typically solved within 50 ms (80 ms maximum). With the same setup we also validated the correctness of our simplex implementation: while the LP problem was being solved on the Shimmer2 node, it was in parallel also solved on the laptop using Solving Constraint Integer Programs (SCIP) [224], a popular non-commercial mixed integer programming solver. A comparison of the results confirmed that our simplex implementation always computed an optimum solution.

Figure 5.6: Execution time of our simplex implementation solving instances of our LP scheduling problem ($N = 7$, $T = 30..40$, $W_n = 1..5$) on a Shimmer2 node.

5.6 Evaluation

We evaluated our reference design COPS by comparison with the IEEE 802.15.4 MAC protocol. We considered the three metrics defined in Chapter 3: *PRR* defined as the ratio of correctly received to transmitted DATA packets; *latency* (also referred to as "queueing delay" below), which is the time between the application passing the packet to the scheduler and the time the application on the same node is signalled the success or failure of the transmission; and *energy consumption* due to longer ACK packets and signal processing.

We implemented an application that generates on average one packet per second on each device.[5] The actual packet generation time is random, uniformly distributed over a one-second window. Since our application does not include a network layer, an outgoing DATA packet is passed immediately to our scheduler component. For each packet the application records the packet generation time as well as the time the scheduler signals the (un)successful completion of the transmission to calculate the queueing delay. On the coordinator the application records the number of successfully received packets and on the devices the application records whether an ACK was received. Note that we disabled MAC retransmissions, i.e. each packet was transmitted at most once. Packets were transmitted with two alternating output power levels: 0 dBm and -10 dBm, which allowed us to explore the influence of transmission power.[6]

An experiment consisted of a subject continuously walking for approximately 5 minutes outdoors on the sidewalk of an urban street. The experiments were conducted in the same environment as described in Section 4.2 by a total of 7 subjects aged between 23 and 64 years: the 5 male subjects who had performed the experiments described in Chapter 4 and two additional female subjects. We used the same number of nodes and node placement as described in Chapter 4. Each experiment lasted 5 minutes: during the first half of the experiment our scheduler was disabled and packets were transmitted using (only) the IEEE 802.15.4 MAC protocol. During the seconds half of the experiment our scheduler was enabled and controlled the packet transmission

---

[5]Note that a low data rate minimizes the collision probability for the standard IEEE 802.15.4 MAC; our results are therefore primarily a result of our opportunistic packet scheduler and not due to the replacement of a CSMA MAC by a TDMA approach.

[6]We made an exception during the RSSI collection phase, when the output power was always set to 0 dBm.
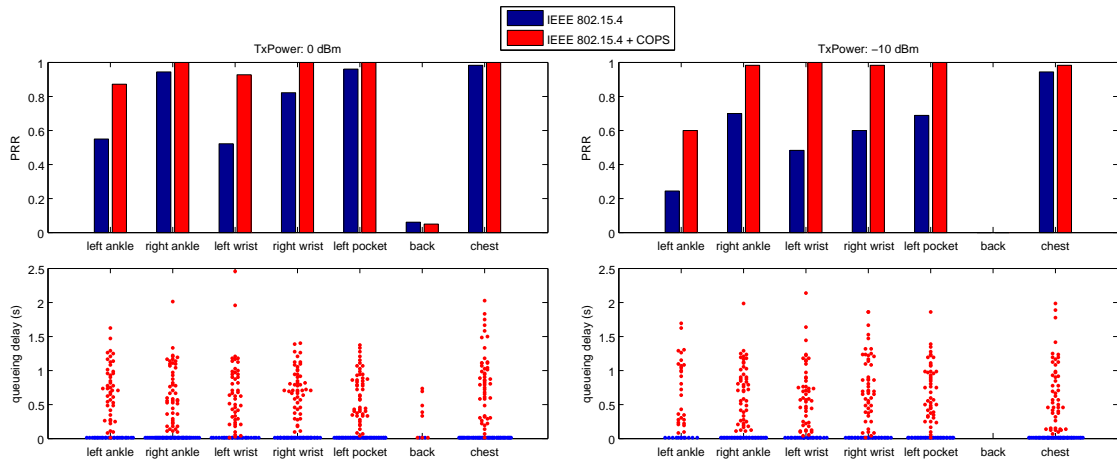
Figure 5.7: PRR and queueing delay for one sample outdoor experiment.

times. In order to have a fair comparison the subjects were asked not to change their walking style during the 5 minute experiment. To check that RF interference did not distort our results (a) during all experiments the devices measured external RF noise in between transmissions and we found that the number of values above the radio sensitivity threshold was negligible; and (b) for each subsequent experiment subjects changed the direction in which a subject walked the path (i.e. a static interferer would not always affect the same time interval of an experiment). We compared PRR, queueing delay and energy overhead during the first 100 seconds when only the IEEE 802.15.4 MAC was active to the 100 second interval once the global schedule had been computed by the coordinator. We conducted a total of 56 outdoor experiments, 8 experiments by 7 different subjects (note that we also conducted some indoor experiments which we will describe in Section 5.6.4).

For illustration Figure 5.7 shows the PRR and queueing delay for one sample experiment. Except for the node positioned on the back of the subject COPS achieves significant PRR increase up to almost 40 percentage points at a transmission power $0\,\mathrm{dBm}$ (Figure 5.7 top left). At $-10\,\mathrm{dBm}$ transmission power the PRR difference is even more pronounced with improvements up to 50 percentage points (Figure 5.7 right). The node positioned on the back of the subject showed no PRR increase, but even a slight PRR decrease at transmission power $0\,\mathrm{dBm}$. This is in line with the preliminary evaluation reported in Section 4.4, which indicated that this position will likely not yield any improvements, because signal strength is often below the radio sensitivity threshold. The queueing delay for the IEEE 802.15.4 MAC is negligible, since packets are transmitted immediately. COPS, on the other hand, introduces an average delay of $0.6\,\mathrm{s}$, with peaks up to $2.5\,\mathrm{s}$. The average delay of $0.6\,\mathrm{s}$ is due to packets being queued for half a stride period on average. A delay of two stride periods may occur when two packets are generated immediately after another at a time when the current OTW has just elapsed.[7]

---

[7]Typically two packets fit into an OTW (which was approximately $1200\,\mathrm{ms}/40{=}30\,\mathrm{ms}$ wide); however, our implementation occasionally did not manage to fit two packets in one OTW, e.g. when the ACK processing was delayed (TinyOS is not a real-time OS).
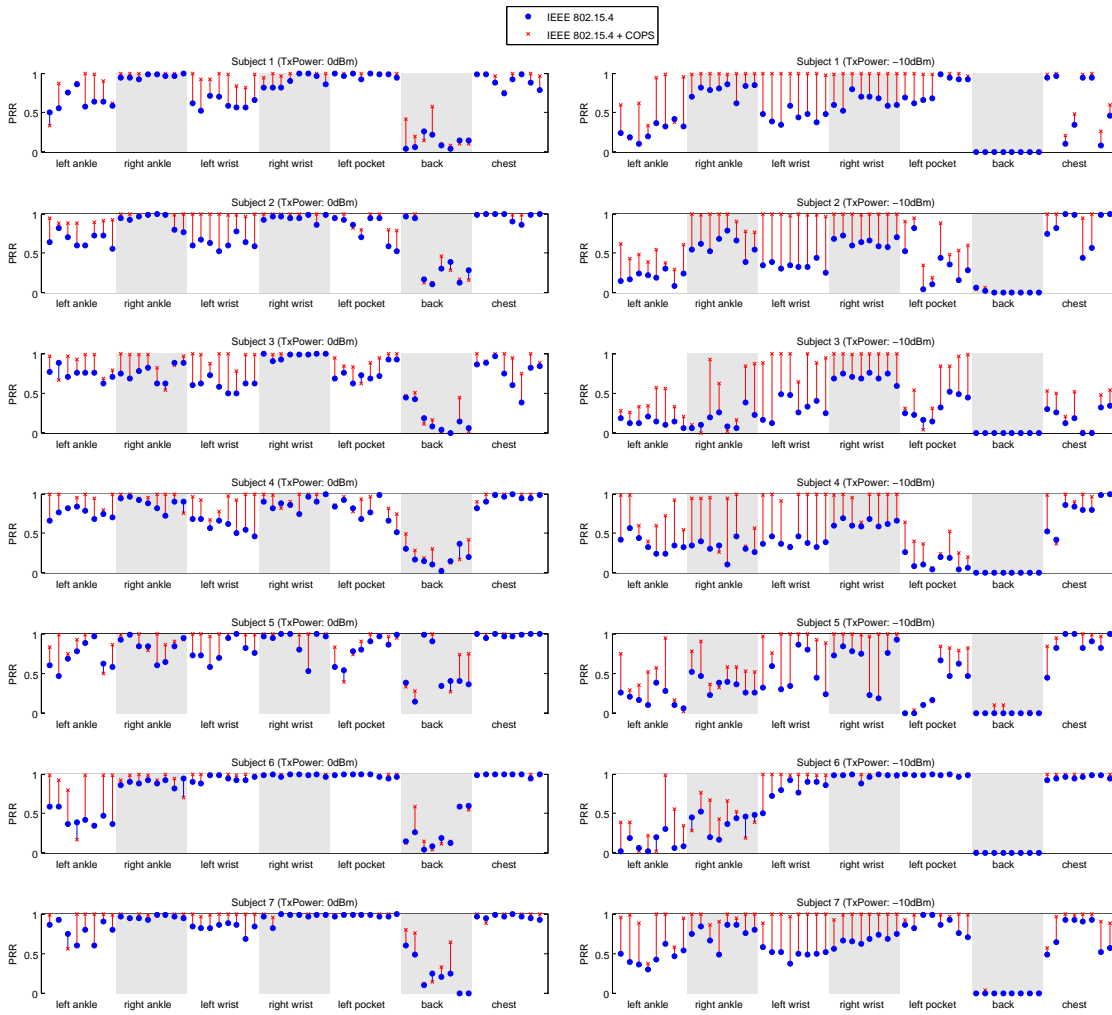
Figure 5.8: PRR results for all 56 outdoor experiments.

5.6.1 Reliability

The PRR results for all 56 outdoor experiments are shown in Figure 5.8. It can be seen that COPS achieves a significant PRR increase compared to the standard IEEE 802.15.4 MAC: typically intermediate links with a PRR between 50 % and 90 % can be turned into reliable links with PRR values well above 90 %. This is most pronounced at a transmission power of -10 dBm where RSSI approaches the sensitivity threshold and small RSSI differences translate into large PRR variations [219]. Links with a PRR below 50 % show smaller improvements, and occasionally a small decrease in PRR. This may be due to insufficient RSSI fluctuations or inaccurate approximation of the RSSI signal when too little RSSI samples are available. In summary, at 0 dBm output power the average PRR increase per node was 10 percentage points and at -10 dBm it was 23 percentage points as indicated by the boxplots in Figure 5.9.

Note that based on the packet transmission timestamps stored by the devices we found that

91
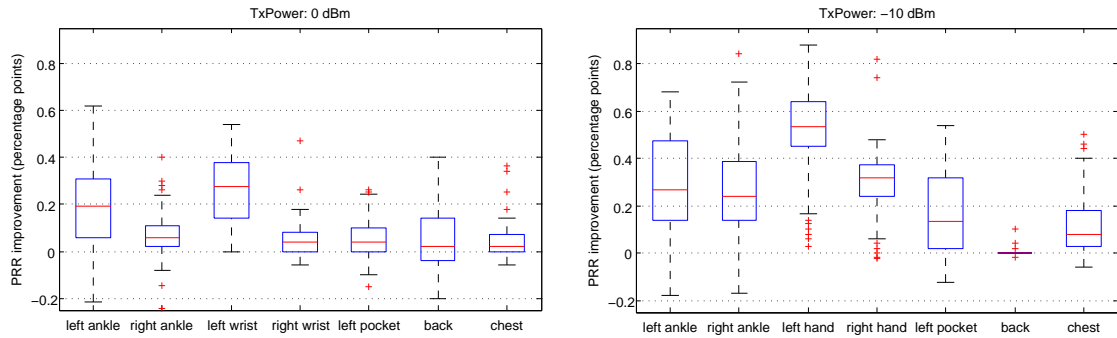
Figure 5.9: PRR gain per node position for all 56 outdoor experiments.

on average 3% (max 5%) of the packets transmitted with the standard IEEE 802.15.4 MAC overlapped in time. This indicates that the reported PRR gains are indeed mostly due to exploiting the signal strength fluctuations rather than due to replacing a CSMA by a TDMA approach.

5.6.2 Queueing Delay

One effect of using the COPS scheduler is an increase in queueing delay which translates into higher packet delivery latency. The results for all 56 outdoor experiments are shown in Figure 5.10. The mean queueing delay with COTS was $715\,\mathrm{ms}$ compared to a mean queueing delay of $5\,\mathrm{ms}$ when using the standard IEEE 802.15.4 MAC. Therefore the delay introduced by COPS is on average about one half stride period larger than the delay introduced by the standard IEEE 802.15.4 MAC but it rarely extends beyond the time of a full stride period. For certain WBAN for epilepsy patients, the latency introduced by COPS may not be acceptable. However, several WBAN applications can accept latencies up to a few seconds, including medical applications that involve low sampling rates such as blood glucose monitoring (for diabetes patients) or body temperature, sports and personal fitness applications, or applications involving monitoring of the environment such as air quality sensing.

5.6.3 Energy Consumption

COPS introduces control overhead, because the scheduler on the coordinator adds a payload portion to ACK packets. To store the current stride period and propagate the beginning of the gait cycle to the devices COPS introduces an ACK payload of 4 byte. Assuming that the radio transmit and receive operation consume the same amount of energy, the radio turnaround time is 12 IEEE 802.15.4 symbols and the IEEE 802.15.4 MAC header is 9 byte, the relative increase in radio energy consumption caused by COPS can be calculated based on the DATA packet's payload (MSDU) size. Figure 5.11 shows the relative increase in energy consumption caused by COPS for different MSDU sizes. The energy overhead ranges between 12% for an empty DATA packet to 3% for a maximum sized IEEE 802.15.4 DATA packet. Note that this overhead occurs only while COPS is actively scheduling packets, i.e. when the packet scheduler is inactive (because the person is not moving) standard IEEE 802.15.4 ACK frames are used and COPS consumes no energy on the devices.

Figure 5.10: Queueing delay for all 56 outdoor experiments.

In addition, the coordinator spends energy for sampling the acceleration signal, which on our platform consumes about $1.5\,\text{mW}$ when active and $9\,\mu\text{W}$ in sleep state [70]. By proper duty cycling the accelerometer during a $30\,\%$ window of a $1.2\,\text{s}$ stride period, the energy spent for acceleration signal sampling corresponds transmitting about 2 maximum-sized DATA packets (without ACKs). Extracting sufficient samples to start the AMDF consumes 11 times more energy.[8] The energy consumption on the coordinator may therefore be significantly higher than on the devices. The overhead will, however, often be acceptable if we assume that the coordinator

---

[8]Note that our MCU may remain in very low-power state ($3\,\mu\text{A}$ in LPM3) while the signal is sampled, i.e. we consider only the energy spent by the accelerometer.
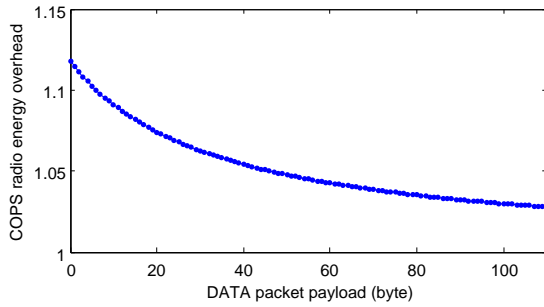
Figure 5.11: Radio energy overhead introduced by COPS due to larger ACK packets.
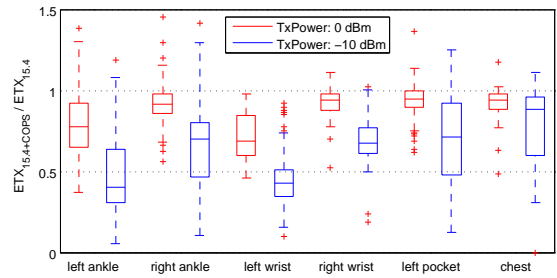
Figure 5.12: ETX ratio for all 56 outdoor experiments.

has more energy resources than the devices, because it is integrated in the user's cell phone.

To better understand the impact on the overall energy budget of a WBAN we used the Expected Transmission Count (ETX) metric [225] to calculate the estimated number of transmissions required to send a packet over a WBAN link. ETX is defined as the inverse of the product of forward delivery ratio (PRR) and reverse delivery ratio (i.e. the probability that the ACK packet was successfully received). We computed the average ETX per node position for our approach ($ETX_{15.4+COPS}$) and the standard IEEE 802.15.4 ($ETX_{15.4}$) for each experiment. In an application whose energy consumption is dominated by the radio chip the ratio $\frac{ETX_{15.4+COPS}}{ETX_{15.4}}$ represents an estimate of how much energy COPS could save. The results in Figure 5.12 show that at an output power of 0 dBm the ETX ratio ranges between 7 % and 29 % (mean 13%); at an output power of -10 dBm it spans from 20 % to 53 % (mean 30 %).[9] If the energy consumption of hardware components other than the radio is negligible COPS thus significantly extends the lifetime of the devices while a person: for example, at transmission power -10 dBm the lifetime of the nodes at left wrist or left ankle could be doubled.

5.6.4 Indoor Measurements

So far we have only considered an outdoor scenario. Indoors the propagation characteristics are typically different, for example, multipath effects due to reflections from near objects may have a strong impact on the signal strength pattern. To understand the implications for our approach we conducted a small measurement study in an indoor office environment. An experiment consisted of a subject walking 20 times through a corridor, which was approximately 2.5 m wide and 20 m long. We used the same application as described in the beginning of this section. During 10 iterations our scheduler was disabled and packets were transmitted using (only) the IEEE 802.15.4 MAC protocol. During the 10 other iterations our scheduler was enabled and controlled the packet transmission times.

For these measurements we modified COPS in two ways: we computed the optimum packet schedule for each node position and each person in advance based on a reference measurement. This was necessary, because with an application data rate of 1 packet per second a walking distance of 20 m was too short to perform the RSSI collection phase, compute a schedule and

---

[9]Note that in Figure 5.12 we omitted the node positioned on the back, because its PRR was very low.

Figure 5.13: PRR results for 12 indoor experiments.

afterwards schedule a significant amount of packets to calculate the PRR. Second, we bootstrapped the gait monitoring process with a reference accelerometer signal, which we had also measured in advance, to reduce the time to start our scheduler.

We conducted a total of 12 indoor experiments, 4 experiments by 3 different subjects. The PRR results are summarized in Figure 5.13. It can be seen that with a transmission power of 0 dBm only small gains are possible, because the PRR values are already very high in the first place. It is interesting to see that — in contrast to the outdoor measurement results — indoors even the node positioned on the back often had a PRR over 90 %. Consequently COPS can only achieve marginal improvements (a mean PRR increase of 2 percentage points). However, at a transmission power level of -10 dBm a significant improvement with a mean PRR increase of 10 percentage points was achieved as can be seen in Figure 5.13 (right). In contrast to the outdoor measurements, the PRR on many links could, however, not be raised to a full 100 %, which we attribute to a harsher multipath environment.

5.7 Summary

In this chapter we have analyzed how periodic fluctuations in received signal strength can be exploited in the WBAN protocol stack through *context-based prediction of the channel state*. As reported in Chapter 4 signal strength fluctuations often occur during regular activities such as walking and are mainly due to oscillating large-scale fading effects (path-loss and shadowing). In this chapter we examined how to monitor the human movement pattern with the help of a tri-axial accelerometer to estimate the current position within the gait cycle while a person is walking. Once periodic mobility has been detected we first extract a set of RSSI samples from existing application traffic (15 samples are typically sufficient) to estimate the RSSI peak within the gait cycle. In parallel we utilize the accelerometer to continuously monitor the stride period and the current offset within the gait cycle. This information is used by our algorithm to estimate future peaks. To avoid packet collisions we have formulated the task of finding a non-overlapping packet schedule based on the collected RSSI samples as a linear programming problem. We have presented an efficient way of solving our class of problems with the simplex method, which allows a sensor node to compute the optimum solution within 100 ms. Finally, we presented our implementation and the results of an extensive experimental evaluation, which show that

- **outdoors**, intermediate links with a PRR between 50 % and 90 % can often be turned into

reliable links with PRR values well above 90 %; in comparison to a system that uses no packet scheduling the average PRR increased by 13 percentage points at 0 dBm transmission power and 26 percentage points at -10 dBm transmission power;

- **indoors**, we have measured a PRR increase of on average 10 percentage points at -10 dBm transmission power; at 0 dBm many links could hardly be improved because PRR was already close to 100 %.

The reported increase in communication reliability is, however, only achievable while a person is walking. In a static scenario our approach is not applicable (although Hall et al. have reported that even when a person is standing still the breathing process often creates noticeable signal strength changes [125]). One limitation is therefore that scenarios in which people perform mostly static tasks, such as office work, are not covered.

On the other hand, in our system the energy overhead for the devices is only marginal, but the coordinator may spend more energy due to signal processing tasks (recall that *only* the coordinator node uses an accelerometer, but not the devices). This is acceptable, because we assume that the coordinator has access to the energy resources of a more powerful device such as the user's cell phone. The cost of our approach is an increase in latency: our scheduler increases packet delivery latency by on average one half stride period. In summary, the presented approach thus realizes an effective mechanism to trade-off latency and communication reliability in WBANs.

Chapter 6

The Impact of External RF Interference on On-Body Communication

WBANs often operate in unlicensed frequency bands and are therefore exposed to RF interference from other wireless technologies. The goal of this chapter is to study the impact of realistic RF interference on the reliability of wireless on-body communication. Our focus on the 2.4 GHz ISM band, because for the foreseeable future this band is expected to keep playing a major role for WBAN communication (Sect. 2.2.1)

The 2.4 GHz band is shared with many other wireless devices such as WLAN stations, microwave ovens or Bluetooth devices. Due to their good spatial coverage and relatively high transmission power WLANs are currently the dominant technology. Figure 6.1 shows the European channel allocation for 802.15.4 and WLAN (IEEE 802.11) in the 2.4 GHz ISM band. It can be seen that all available 802.15.4 channel overlap with one or more WLAN channels. It is well established that — despite interference mitigation mechanisms like DSSS and "listen-before-send" incorporated in both standards — WSN communication can be severely limited by WLANs. While previous studies have treated the problem both from analytical [226–228] and experimental [21, 229, 230] side they have mostly focused on static WSN deployments. However, the impact of RF interference on intra-WBAN communication in realistic environments is not well-understood. Due to human mobility the impact of RF interference may indeed be quite different than in static networks. For example, on the one hand, a WBAN may experience gradual increase in RF interference as the person approaches a static interference source, like a WLAN access point, and a gradual decrease as the person withdraws from the interferer. On the other, it may happen that due to body shadowing the WBAN may (temporarily) be "shielded" from an interferer.

In this chapter, after discussing related work, we present the results from three measurement studies: first we study the impact of WLAN activity on intra-WBAN communication in a controlled interference environment (Section 6.2). These results show, among other things, that there may be large variation in performance among different WBAN hardware platforms. Then, we first present the results from an extensive RF noise measurement campaign conducted in 12
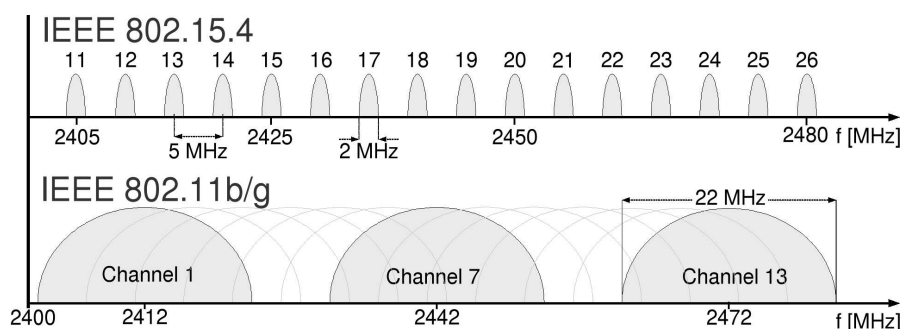
Figure 6.1: IEEE 802.15.4 and 802.11 channel allocation in the 2.4 GHz ISM band.

different urban environments (Section 6.3). Afterwards, we report on a set of link measurements performed in an urban shopping area (Section 6.4). Note that the two measurement studies are based on two different hardware platforms (Section 6.3: Shimmer2, and Section 6.4: Tmote Sky). At the end of this chapter, we put the results from our 2.4 GHz spectrum measurement campaign in perspective with the results reported in Chapter 4: in particular, we estimate what fraction of the packet loss in a WBAN may be caused by RF interference rather than a weak signal below sensitivity threshold (fading).

6.1 Related Work

The problem of coexistence between 802.11 and 802.15.4 networks has received significant interest in the WSN research community. Most early work concentrated on developing probabilistic models that capture the dependence of interference-related packet loss in a 802.15.4 network based on frequency overlap and duty cycle, transmit power and distance of an 802.11 interferer [226–228]. As 802.15.4 technology became more widely available, the focus turned towards experimental approaches to validate the conclusions obtained from the established analytical models. Several studies showed significant interference-related packet loss under moderate to high utilization of the 802.11 channel [231, 232]. Others have analyzed the reverse problem, i.e. the impact of 802.15.4 networks on 802.11 devices, reporting that 802.15.4 devices may also cause loss in an 802.11 network [233].

Acknowledging the problem, several 802.15.4 radio chip manufacturers published guidelines on how to mitigate interference effects between the two technologies [234–236], for example, through minimal frequency offset of 20 MHz or by using ARQ to translate losses into latency [235].

While many coexistence studies have been conducted in controlled laboratory environments, only some have reported on urban 2.4 GHz RF spectrum utilization. For example, an urban power spectrum measurement campaign in the San Fransisco Bay Area revealed "significant levels of the man-made signals regardless of time and location proving the proliferation of unlicensed devices in ISM 2.4 GHz" [237]. More recently, Hanna and Sydor reported on the results of a spectrum monitoring campaign at a (single) urban outdoor location in downtown Ottawa, Canada. Their goal was to assess to what extent 2.4 GHz white space can be used by Cognitive

Radio (CR) WLANs [238]. The authors studied channel occupancy (with the help of a CR platform) and the interference behaviour by WLAN packet type (with a WLAN sniffer). They discovered that 5-15% of the 2.4 GHz is occupied, but the "number of dominant interferers [. . .] is very small" [238]. Others have stated that in urban residential areas 2.4 GHz RF activity is often complex, (temporal) patterns are not obvious and there exists non-negligible noise from non-802.11 devices [239]. On the other hand, Valenta et al. argue that high utilization of the 2.4 GHz band is often contrasted by only sporadically utilized licensed frequency bands above or below the 2.4 GHz ISM band, which might be utilized via cognitive radio concepts [240].

Previous studies have, however, rarely taken the specific characteristics of WBANs into account. One exception is an experimental study on the impact of RF interference among multiple WBAN in an indoor office environment [241]. The authors' main conclusion is that distance is not a good indicator for the received interference power, because the latter is dominated by random variations. Still, only little studies have targeted the quality and quantity of 2.4 GHz RF activity in realistic urban environments as perceived by WBANs. In the rest of this chapter we target this under-explored area.

6.2 Controlled Interference Experiments

The goal of this section is to study how controlled 802.11 interference affects the reliability of a wireless on-body communication link. In addition, we are interested in comparing the results for two different WBAN platforms, the Shimmer2r platform [8] and the Tmote Sky platform [66]. Recall that the two platforms use the same radio 802.15.4 transceiver, but have different antennas, form factor and — when attached to the body — antenna-to-body surface distance (Section 3.2). We are thus able to study the impact of these factors.

6.2.1 Setup and Scenario

We conducted channel measurements outdoors in a park, an environment of negligible RF interference which we verified with the help of a portable spectrum analyzer. We configured two laptops to form an 802.11g ad-hoc network and started a large file transfer from one to the other. One laptop was placed on top of the other on the ground and constantly generated heavy traffic on 802.11 channel 7 at 54 Mbit/s with a constant transmission power of 15 dBm.

We used two WBAN nodes, one attached to the right wrist (sender) and one placed in the right pocket (receiver) of a person. The sender was transmitting 1 packet every 20 ms. Each packet had an MPDU size of 36 byte and packets were transmitted by round-robin iterating over the entire range of 802.15.4 channels (Figure 6.1): 11, 12, . . . , 26, 11, 12, etc. We call an iteration over all 16 channels a *sweep*. After every sweep we modified the transmission power (round-robin) selecting one value from the following set: 0 dBm, -10 dBm and -25 dBm. In order to synchronize the hopping pattern among sender and receiver we introduced cables between the digital I/O pins of the two nodes' MCUs as explained in Section 6.4.1. During an experiment the nodes stored statistics of the exchanged packets on their SD card (Shimmer2r) or transmitted them over USB to a laptop that was carried in a backpack of the person, where the data was stored in a file (Tmote Sky).

Figure 6.2: Environmental RF noise during the controlled RF interference experiment.

The experiment was simple: the person carrying the WBAN first initially stood about $40\,\mathrm{m}$ away from the 802.11 network, at time $t = 0\,\mathrm{s}$ started walking on a straight line towards it, passed $1\,\mathrm{m}$ by the two laptops (at about $t = 30\,\mathrm{s}$) and continued walking on the (virtual) straight line. During this time the receiver recorded the changes environmental RF noise ($P_{Interference}$, Section 3.2.5), the number of failed packet receptions and per-packet RSSI. After the experiment we sorted the measurement results by the transmission power level and analyzed RF noise, PRR and RSSI in the time and frequency domain.

6.2.2 Results

The RF noise results are shown in Figure 6.2. In general, the results for the two platforms are quite similar. As expected, only the 802.15.4 channels that overlap with the WLAN channel 7 (Figure 6.1) are affected: 802.15.4 channels 17 to 20 exhibit significant levels of RF noise, which increases as the distance to the 802.11g network gets smaller. However, even at very close distance (around $t = 80\,\mathrm{s}$), there are a range of different noise floor values, some as low as $-99\,\mathrm{dBm}$. A possible explanation is that the 802.11 stations are not permanently transmitting. Since transmissions are acknowledged the WLAN radios are frequently changing from transmit to receive mode and there may be some idle listening before receiving a frame due to the execution of the CSMA algorithm. Furthermore, the CC2420 averages a single noise floor reading over $128\,\mu\mathrm{s}$ (Section 3.2.5), and it can thus happen that a sample is taken while the channel is (partially) idle. This indicates that a single noise floor value is not reliable for determining presence of an interferer even in close proximity.

The results for the corresponding packet receptions are shown in Figure 6.3a for Tmote Sky and in Figure 6.3b for Shimmer2r. The graphs show a matrix, in which the element entries indicate either successful (white) or failed (black) reception of a packet, respectively. Each graph shows the result for a different 802.15.4 transmission power level. As expected, for a given platform there are more losses when transmission power is set to a lower value. However, what is remarkable is that the results between the two platforms are quite different: the Tmote

Figure 6.3: Packet reception matrix during the controlled RF interference experiment for the right wrist → right pocket link.

Sky platform consistently reported significantly less packet loss than the Shimmer2r platform. Also, while packet losses on the Tmote Sky platform occur only on channels that overlap with the WLAN interferer, at -10 dBm and -25 dBm transmission power the Shimmer2r platform also suffers losses on non-overlapping channels. It seems that received signal strength on Shimmer2r

(a) TxPower: 0 dBm   (b) TxPower: -15 dBm   (c) TxPower: -25 dBm
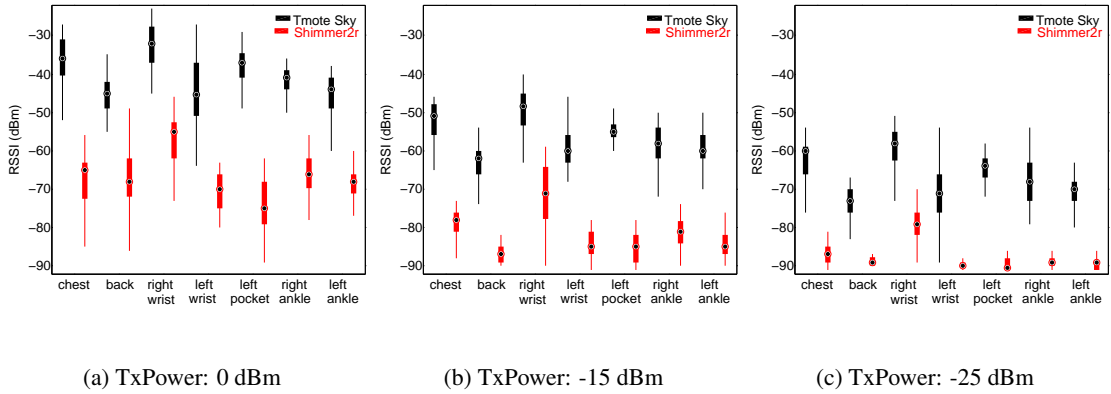
Figure 6.4: RSSI distributions for Tmote Sky and Shimmer2r experiments (channel 11).

is lower than on the Tmote Sky platform. This would also explain the periodic pattern in the bottom graph of 6.3b. As shown in Section 4.3 walking translates into periodic changes in received signal strength (c.f. Figure 4.1 and Figure 4.2): due to the movements of the limbs the signal may (temporarily) be too weak to be decoded by the receiver.

To investigate this hypothesis we analyzed the received signal strength of all successfully received packets on a non-overlapping channel (channel 11). We also repeated the experiment for all node positions shown in Figure 3.3a. The results are shown in Figure 6.4a (0 dBm transmission power), Figure 6.4b (-10 dBm transmission power) and Figure 6.4c (-25 dBm transmission power). The graphs, which summarize the results of 14 experiments (one per receiver node and platform) as a boxplot, confirm our hypothesis: received signal strength on Shimmer2r was consistently about 25 dB lower than on Tmote Sky.

6.2.3 Discussion

There are several possible explanations for the difference in received signal strength: Tmote Sky uses an inverted-F microstrip antenna while Shimmer2r uses an SMD antenna. Shimmer2r has a smaller form factor and when the node is attached to the body the antenna-to-body surface distance is only about 1 cm. The Tmote Sky is larger, it has a bigger container for the batteries and because it has no native enclosure we placed it in a small plastic enclosure. This resulted in a antenna-to-body separation of more than 3 cm. As shown by Roelens et al. the distance between antenna and body surface has a significant impact on received signal strength [110]. Finally, in our experiments only the Tmote Sky nodes were powered over USB, because the data was transmitted over a USB cable to a laptop carried by the WBAN subject in a backpack (in contrast to the Shimmer2r plaform the Tmote Sky has little internal memory). The Shimmer2r nodes were powered by battery. We believe that the Shimmer2r results are more representative, because the platform is specifically designed for WBAN applications, while the Tmote Sky is more suitable for static deployments. Note that the measurement study described in the following Section 6.3 is based on the Shimmer2 platform, while the study described thereafter in Section

6.4 was conducted with Tmote Sky nodes (due to unavailability of the Shimmer2 platform at the time of conducting this particular study).

6.3 Urban RF Noise Measurements

The previous section showed that interference from co-located WLANs may significantly decrease the reliability of a 2.4 GHz wireless on-body communication link. However, the experiments were performed under artificial conditions, in a controlled interference environment. To understand the scope of the problem in realistic environments in this section we present the results from an urban RF noise measurement campaign. In these measurements we collected about half a billion RF noise samples in various urban environments with COTS WBAN hardware. Our setup passively captured RF noise in the entire 2.4 GHz band on five different body positions simultaneously. The goal of this section is to analyze the collected dataset in the time, space and frequency domain. Note that in a later Section 6.5 we will revisit the results and put them in perspective with the results reported in Chapter 4 to estimate what fraction of packet loss is typically caused by RF interference rather than by fading. Note that the measurement study described in this section is based (only) on results obtained with the Shimmer2 hardware platform.

6.3.1 Setup and Scenarios

This section introduces our measurement setup and the scenarios in which we performed the measurements.

Measurement Setup

Our measurement setup consists of six Shimmer2 [8] nodes attached to the human body at the positions marked in Figure3.3a except for the positions at the left ankle and back (we omitted these positions to make the setup more comfortable to wear over longer time). Our setup is completely passive: the main task of the nodes is to continuously perform RF noise measurements by *sweeping* over the entire 2.4 GHz ISM band (16 IEEE 802.15.4 channels). Each *sweep* consists of taking one RSSI (Received Signal Strength Indicator) sample on each of the 16 channels, where each sample represents the ambient RF noise power averaged over a window of $128\,\mu s$.

Our setup introduces an auxiliary wired channel to synchronize the sampling processes of the individual nodes in time. This synchronization is performed over digital I/O signals via pins exposed on the Shimmer2 expansion connector: with a set of dedicated shielded cables we interconnect each of the nodes with a central master node (carried in the right chest pocket of the subject). The master node is responsible for signalling the beginning of a new sweep to the five nodes attached to it via cables. The master node does not perform RF noise measurements itself, because this would interfere with maintaining precise synchronization over the digital I/O connections. Our setup thus effectively five nodes that continuously and simultaneously sample ambient RF noise power in the 2.4 GHz band.[1] Each node stores the data on its SD card and

---

[1]Note that our hopping pattern, i.e. the order in which the 16 channels are covered during a sweep, is randomized,

Figure 6.5: Example 2.4 GHz RF noise trace captured on the left hand and right hand simultaneously.

after an experiment we collect the SD cards and extract the data. Note that we do not calibrate the RSSI data among the different nodes before we use it in our evaluation, because the noise floors of our particular nodes are very similar[2] and previous work has shown that differences in CC2420 RSSI readings are due to a linear offset [116], which can be determined through a noise floor measurement.

In our setup the duration of a sweep over the set of 16 channels is 7 ms; this includes the time for storing the 16 RSSI samples on the SD card. We thus achieve a sweep frequency of 142 Hz resulting in an RSSI sampling frequency of 2.2 kHz. This performance can easily compete with commercial low-cost spectrum analyzers, such as the Wi-Spy 2.4x [242] USB spectrum analyzer. However, in contrast to a spectrum analyzer, our mobile setup allows to perform distributed spectrum measurements on multiple body positions in parallel with authentic WBAN radios.

Scenarios

We had one subject who performed 12 experiments in different urban environments. We selected three categories to classify the experiments: environment (park, campus, residential area, shopping street, urban transportation system); degree of subject mobility (static, mobile or mixed); and indoor vs. outdoor. Table 6.1 describes and classifies all experiments accordingly. Note that some of the experiments covered a certain path multiple times, e.g. a series of walks around the same blocks in a certain neighbourhood. For these experiments Table 6.1 shows the number of iterations. On average an experiment lasted for slightly more than one hour, resulting in a total time of 13 hrs, which resulted in a huge dataset of about half a billion RSSI samples.

6.3.2 Results

In this section we analyze the data collected in the previously described experiments. For our analysis we use two general thresholds: -94 dBm, which is the sensitivity threshold of the radio; and -85 dBm, which can be considered the RSSI of an intermediate link quality ("gray area"). Figure 6.5 gives a first impression of the capabilities of our setup: it shows a one second snapshot

---

in order to not accidentally correlate with other frequency hopping systems. All nodes follow the same pattern as they use the same seed to initialize the pseudo-random number generator.

[2]Max. 2 dB difference in interference-free environment over all channels.

Expt. no.	Environment	Description	Subject mobility	Duration
1	campus (indoor)	sitting at an office desk, working in front of a laptop	static	40 min
2	campus (indoor)	walking through campus buildings (5 iterations, 3.5 min each)	mobile	17.5 min
3	park (outdoor)	walking in an urban recreational park (10 iterations, 6 min each)	mobile	1 hr
4	residential (outdoor)	walking around a suburban residential block (15 iterations, 5 min each)	mobile	75 min
5	residential (indoor)	inhouse activities (kitchen work, watching TV, cleaning, . . .)	static/mobile	66 min
6	residential (indoor)	like expt5, but in a different apartment	static/mobile	70 min
7	residential (indoor)	like expt5 (same apartment)	static/mobile	73 min
8	shopping street (outdoor)	walking an urban shopping street repeatedly (15 iterations, 6 min each)	mobile	90 min
9	shopping street (outdoor)	sitting outdoor in a suburban coffee bar (+few minutes walking)	(mostly) static	90 min
10	shopping street (outdoor)	walking from apartment to supermarket (few minutes inside) + return	mobile	48 min
11	shopping street (indoor)	sitting indoor in a suburban coffee bar (+few minutes walking)	(mostly) static	77 min
12	transportation system (indoor)	using the public transport system (urban railway, underground, tram)	(mostly) static	75 min

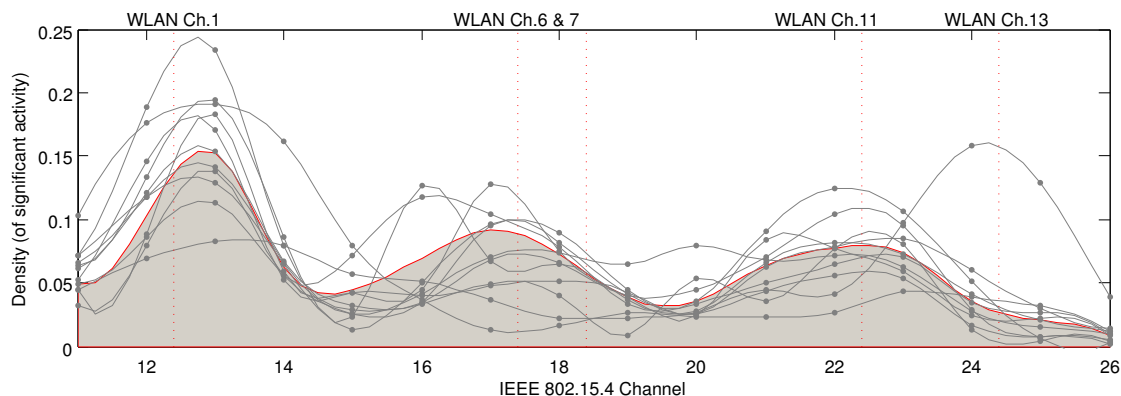Table 6.1: List of RF noise collection experiments.

Figure 6.6: Kernel density plot of spectrum activity $\geq -94\,\mathrm{dBm}$ for all experiments (shaded area) and individual experiments (gray graphs).

of the intensity of 2.4 GHz RF noise measured simultaneously by nodes located on the left and right wrist of the subject. It can be seen that for the majority of time there is little activity above $-94\,\mathrm{dBm}$, but visually there is correlation in frequency and time. Furthermore, there is already strong visual evidence that the data captured on the two different body positions are correlated. In the rest of this section we will explore the data along the dimensions frequency, time and space.

Frequency

To get a better understanding of the distribution of RF noise in the frequency domain, we created kernel density plots showing the density of spectrum activity $\geq -94\,\mathrm{dBm}$ over the 802.15.4 channels. Figure 6.6 shows the results: the shaded area represents the entire dataset; it clearly shows peaks at IEEE 802.15.4 channel 13, 17 and 22. These channels happen to be very close to the popular WLAN channel 1, 6 and 11 center frequencies.[3] Furthermore, the shape of the peaks resembles the Power Spectral Density (PSD) of 802.11 DSSS signals, which is strong evidence that WLAN is indeed the major source of 2.4 GHz activity. The figure also suggests that 802.15.4 channels that are far from WLAN center frequencies, in particular channel 26, should (on average) be least affected by urban RF noise.

We were interested in the correlation of activity on (neighbouring) channels. Such information is a first indicator for the designer of a frequency hopping algorithm, about "where (not) to hop". Figure 6.7 shows the empirical conditional probability for spectrum activity $\geq -94\,\mathrm{dBm}$ on a channel, given that at least one of the two neighbouring channels showed spectrum activity $\geq -94\,\mathrm{dBm}$ activity during the same sweep.[4] With a threshold of $-94\,\mathrm{dBm}$ the average probability is around 20 %, which is roughly a factor 4 increase compared to the average 5 % probability for detecting spectrum activity $\geq -94\,\mathrm{dBm}$ (Figure 6.8). A threshold value of

---

[3]Although the measurements were performed in Europe, the recommended European configuration (1,7 and 13) does not seem to be common. This might be due to a default U.S. configuration of WLAN AP equipment.

[4]In this analysis we omitted channels 11 and 26 as they only have a single neighbouring channel.

Figure 6.7: Conditional probability for RSSI samples being larger or equal than a threshold THR (-94 dBm or -85 dBm), given that the RSSI of a neighbouring channel (± 1 channel) was larger or equal than that threshold during the same sweep.

Figure 6.8: Channel occupancy relative to two thresholds: -94 dBm and -85 dBm.

-85 dBm results in an increase of up to an order of magnitude. These results indicate that there is a significant correlation of activity on neighbouring channels.

Time

To get a first insight in the intensity of 2.4 GHz environmental RF noise we analyzed channel occupancy, i.e. the amount of samples exceeding either of two thresholds, -94 dBm or -85 dBm, per experiment. Or stated differently, the empirical probability that an RSSI sample was larger or equal than a given threshold. The results are shown in Figure 6.8. In several experiments spectrum activity occurred at or above -94 dBm for about 5 % of the time, but there is quite some fluctuation: one experiment has a value well below 1 % (expt3), others show a channel occupancy of 10 % (expt4+5) and up to 20 % (expt8). When applying a -85 dBm threshold

107

Figure 6.9: Empirical probability for RSSI samples being larger or equal than -94 dBm for experiments which included multiple iterations.

channel occupancy decreases considerably: none of the experiments showed more than 5 % of activity. This indicates that the majority of detected environmental RF noise has values below -85 dBm. On the other hand, the IEEE 802.15.4 standard states that typically an SINR of 5 dB to 6 dB is required for good reception [18, Section E.3.2], i.e. even values below -94 dBm can have an impact on the reception rate (c.f. Section 6.5 and Figure 6.20).

Based on Figure 6.8, no obvious trend with respect to our scenario classification is noticeable: the experiment with highest RF noise activity (expt8) was carried out in a shopping street; yet, other experiments in different shopping streets rendered much less activity. To examine the "stability" of RF noise on a larger time scale we analyzed the experiments that included multiple iterations (such as expt8): as shown in Table 6.1 some of the experiments consisted of the subject following a certain path multiple times, e.g. a series of walks around the same blocks in a certain neighbourhood. For such experiments (in total 4) we compared the individual iterations, e.g. the results of two subsequent walks around the same block. Because these iterations were carried out immediately after another, such a comparison can give some insight in the temporal stability of urban 2.4 GHz RF noise. To this end, we compared the relative amount of samples exceeding -94 dBm for the individual iterations of experiments 2,3,4 and 8, which all included an iterative path.

The results are shown in Figure 6.9. Experiment 8, which had the highest RF noise activity (Figure 6.8) shows very stable iterations, all close to the average 20 % value. The situation for experiment 4 is similar. For these two experiments the RF noise conditions thus did not change significantly over several tens of minutes (possibly it was a longer-lasting WLAN file downloads from a device that our subject passed repeatedly). The results for the park environment (expt3) show more variance, but are less representative due to the limited amount of spectrum activity ≥ -94 dBm encountered during this experiment. One conclusion from these results is that it might be promising to take past RF noise data for a given location into account when designing interference mitigation techniques.

To better understand the temporal correlation ("burstiness") of RF noise we examined our dataset on a smaller time scale. We call a sudden increase in the amount of RF noise samples ≥ -94 dBm an RF noise *burst*. More formally, we define a burst as a situation when during a window of size Δ_B seconds 50% or more of the RF noise samples have a value greater or equal than -94 dBm. We choose a threshold of 50%, because this value is about a factor 10 higher than the average RF noise level we have measured for all experiments. Therefore, if RF noise was uniformly distributed over time, then RF noise *burst* would be quite rare. From our traces

we computed the number of distinct[5] bursts and the average duration of bursts per channel for different Δ_B over all experiments. We also computed the same values for a synthetic (simulated) scenario, in which an identical amount of RF noise samples $\geq -94\,\text{dBm}$ is uniformly distributed over the same amount of time. Note that the duration of an RF noise burst is the period of time over which the average number (over Δ_B seconds) of RF noise samples $\geq -94\,\text{dBm}$ exceeds 50%. Therefore, the duration of an RF noise burst may be shorter than Δ_B.

The results are shown in Figure 6.10a and Figure 6.10b. It can be seen that there is a large variance in the number of bursts over different channels. As expected the 802.15.4 channels that overlap with WLAN channels 1, 6 and 11 show one or two orders of magnitude more bursts than the remaining 802.15.4 channels. Also, the number of bursts generally decreases as the interval Δ_B increases, because longer bursts are less frequent than shorter ones. The average duration of bursts is on the order of several hundred milliseconds and also varies per Δ_B, but it is independent of the particular channel (Figure 6.10b). The result for channel 20 in Figure 6.10b can be considered an outlier: we found that the result is caused by a 2-minute uninterrupted RF noise burst on channel 20 during experiment number 9. The most interesting observation, however, is that the empirical results generally differ from the synthetic ones: our experiments exhibit one order of magnitude more burst then we would expect if RF noise was random, uniformly distributed over the same amount of time. The same is also true for the average burst duration. From these results we can conclude that urban RF noise is generally bursty and therefore even if average spectrum utilization is low, it may be necessary to apply RF interference mitigation techniques, such as channel hopping, to avoid temporary blockage of the wireless communication.

Space

The example snapshot shown in Figure 6.5 indicated a strong visual correlation between the RF noise samples captured simultaneously on different body positions. To get a better insight we computed the difference between two RSSI samples taken simultaneously at two different node positions, whenever any of the two nodes had detected spectrum activity $\geq -94\,\text{dBm}$. Figure 6.11 shows a Cumulative Distribution Function (CDF) that summarizes the results for all node pair combinations computed over our entire dataset (in this figure we use a short notation for the body positions: RW = right wrist, LW = left wrist, RA = right ankle, LA = left ankle, CH = chest, RP = right pocket). On average every second sample taken by two nodes simultaneously had a difference of 5 dB or less. In about 80 % of the cases the difference was no more than 10 dB. Figure 6.12 shows a different perspective: it represents the empirical probability of a node detecting RF activity (above -94 dBm, or -85 dBm), given that another node detects the activity in parallel. The probability is more than 50 %, regardless of the particular node pair combination.

This indicates that for the practical task of *binary* interference detection, i.e. the detection if interference is present or not, the results may be similar for different body positions. However, the difference in interference power measured at two different body positions is not negligible,

---

[5]Two bursts are distinct if there exists an interval of size Δ_B seconds between with an average number of RF noise samples $\geq -94\,\text{dBm}$ below 50%.

(a) Number of RF noise bursts.
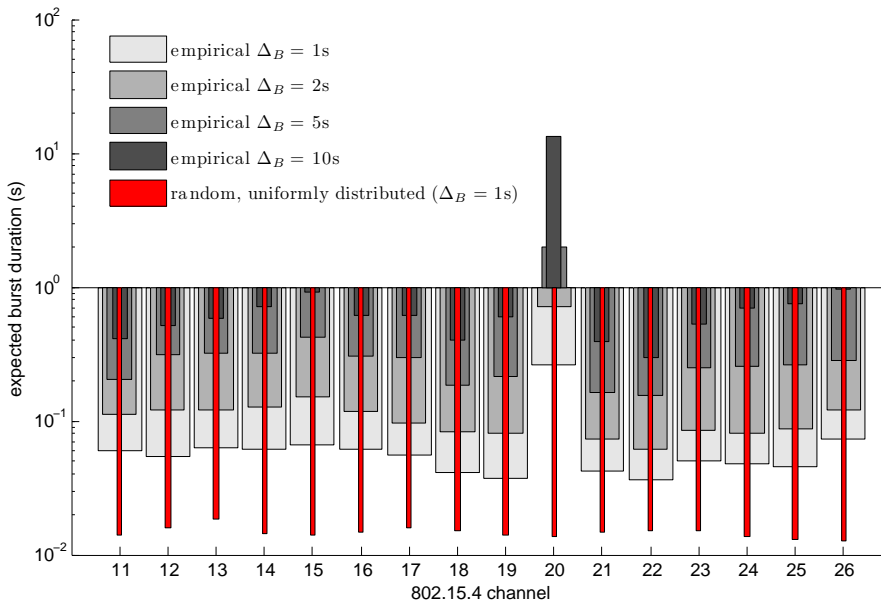
(b) RF noise burst duration.

Figure 6.10: Characteristics of RF noise bursts.

it may vary up to $10\,\mathrm{dB}$ or more.

110
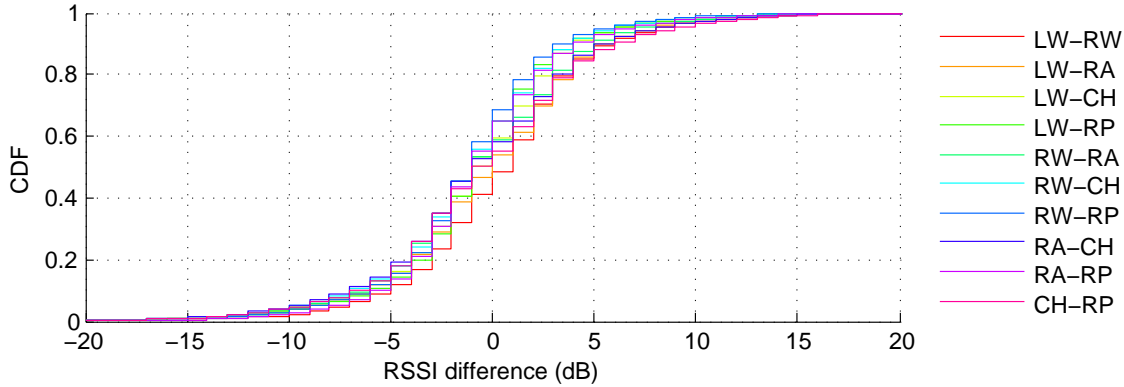
Figure 6.11: CDF of the instantaneous RSSI difference between different node position pairs.

Figure 6.12: Conditional probability for an RSSI sample being larger or equal than a threshold THR (-94 dBm or -85 dBm), given that the simultaneously sampled RSSI on another node was larger or equal than that threshold.

6.3.3   Discussion

We have analyzed the results from an RF noise measurement campaign conducted with a set of five WBAN nodes that synchronously collected in total half a billion 2.4 GHz RF noise samples in various urban environments. The key observations are:

- There is strong evidence that WLAN is often the major source of urban 2.4 GHz spectrum activity. Other measurement campaigns have come to the same conclusion [239, 243].

- In many experiments spectrum activity $\geq -94\,\text{dBm}$ occurred for about 5 % of the time, but there is a large variation among the scenarios. Similar results of 5-15% channel occupancy in the 2.4 GHz band were reported in [238].

- An 802.15.4 intra-WBAN link that achieves an RSSI of more than -90 dBm will often be quite robust against urban RF interference, because the vast majority of samples had values below -94 dBm (an SINR of 5 dB to 6 dB is typically enough for good reception with a 2.4 GHz IEEE 802.15.4 radio [18, Section E.3.2]).

- Urban RF noise is about one order of magnitude more "bursty" than random (uniformly distributed over time) RF noise.

- There is a significant correlation of the simultaneous activity on neighbouring 802.15.4 channels, therefore channel hopping should avoid "short" hops below 10 MHz.

- The difference in interference power measured at two different body positions is not negligible but may vary up to 10 dB or more.

Given that the vast majority of RF noise samples had values below -94 dBm we expect the overall impact of urban RF noise on intra-WBAN communication to be less severe than fading, which had caused significant losses on some links as shown in Chapter 5. Also, it is expected that the selection of am IEEE 802.15.4 channel that does not overlap with the default WLAN center frequencies will likely be a good remedy against the effects of RF interference. On the other hand, since urban RF noise is bursty, even if average spectrum occupation is low it may be necessary to apply RF interference mitigation techniques, such as channel hopping, to avoid temporary blockage of the wireless communication.

From the analysis it is, however, not immediately clear how RF interference affects the reliability of an on-body communication link. In the next Section 6.4 we present the results from a set of link measurements. While the results give valuable insight in the correlation between WBAN losses and WLAN activity they are unfortunately not directly applicable for the Shimmer2 platform (they were conducted with the Tmote Sky platform). Therefore, in Section 6.5 we will revisit the results presented above and analyze them in light of the results reported in Chapter 4. This will allow us to estimate what fraction of packet loss would typically caused by RF interference rather than by fading (a weak signal below sensitivity threshold). This distinction is important, because in practice the countermeasure will vary (e.g. change channel in case of interference or wait until fading conditions have improved in case of a weak signal).

6.4 Urban Link Measurements

The previous section presented the results from an urban measurement campaign, in which a mobile WBAN passively collected a large number of RF noise samples. The results in Section 6.3.2 indicated that WLAN is a major cause of packet loss, but the conclusion was based on an evaluation of spectral footprints rather than detection of real WLAN activity. The goals of this section is to investigate empirical correlation between 802.15.4 packet delivery performance and "real-life" WLAN activity. This is helpful for designing frequency diversity schemes (e.g. deciding how "far to hop" in the frequency domain). To this end we also explore 802.15.4 cross-channel quality correlation and trends that may be used as a potential trigger for channel hopping. In the following we will first describe our experimental setup and then report on a representative sample of our dataset from an urban environment. Afterwards we analyze the empirical traces for cross-channel quality correlation and trends in the noise floor. Note that the measurement study described in this section is based (only) on results obtained with the Tmote Sky hardware platform.

6.4.1 Setup and Scenarios

Our measurements were performed with Tmote Sky [244] sensor nodes (Section 2.1.3). The choice of platform was due to unavailability of the Shimmer2 platform at the time of conducting this particular study. Our setup consisted of two nodes, each node is placed in a thin plastic enclosure and strapped to a person: one to the left upper arm, the other on the right ankle, resulting in a relative distance of about 1.5 m. When the person stood still both nodes had the same alignment and both surface areas were facing in the same horizontal direction. However, in all experiments, the test person was walking at an even speed of about 1.2 m/s (common walking speed). Our setup introduces two auxiliary wired channels: one to synchronize the transmissions on the IEEE 802.15.4 channel; and a second for streaming measurement results to a laptop. This is schematically shown in Figure 6.13a and explained in the following.

Our measurement software accesses the CC2420 radio directly, there is no MAC layer involved and all packets are sent immediately (without clear channel assessment, CCA). Like the passive setup using Shimmer2r nodes described in Section 6.3.1 the two nodes continuously iterate over the 16 channels, however, in addition to measuring RF noise activity the nodes are exchanging one `DATA` and acknowledgement (`ACK`) packet per channel. We call an iteration over all 16 channels – involving 32 packets – a *sweep*. During a sweep the roles of the nodes are fixed: one node sends the `DATA` packets, the other node sends the `ACK` packets; after every sweep the roles are swapped. We use acknowledgements, because this is common to many IEEE 802.15.4 networks, and we let the nodes swap roles, because this allows us to (better) evaluate link (a)symmetry.

The CC2420 supports different transmission power levels, the datasheet documents 8 discrete levels ranging from -25 dBm to 0 dBm [69]. The relevant `TXCTRL` register of the radio, however, accepts 32 different values. The radio manufacturer confirmed to us that all 32 values are valid, however, "the relation between the register setting and the output power is not linear". We experimentally determined an output power of roughly -42 dBm when the `TXCTRL.PA_LEVEL` is set to the value of 2.[6] In our study we then used three different output power levels of -10 dBm, -25 dBm and -42 dBm, alternating every sweep. The Message Sequence Chart (MSC) in Figure 6.13b shows the sequence of operations performed during a sweep.

In our setup whenever a packet is transmitted by one node, the other node's radio is in receive mode, both nodes have tuned to the same channel, and they use the same level of transmission power for the `DATA` and `ACK` packet. This requires very careful synchronization. One option is to perform synchronization through the arrival times of `DATA` and `ACK` packets, but since our experiments are conducted in an environment of possibly strong RF interference we rejected such "in-band" synchronization, because it would require unknown (conservative) guard times to counter the clock drift in case of successive packet losses. Instead we perform synchronization over digital I/O signals through pins exposed on the Tmote Sky expansion header: with a

---

[6]We chose 38 sender/receiver combinations from a batch of 10 Tmote Sky nodes and measured RSSI for different documented as well as the undocumented register settings in a low-interference environment based on a fixed node placement. Assuming a linear relationship between output power and RSSI, we used the documented values in combination with the measured average RSSI as anchors for a linear regression to determine the unknown transmission power level based on the corresponding measured RSSI value.

(a) Schematic measurement setup. The laptop is carried in a backpack.

(b) MSC describing a sweep. Note that only DATA and ACK packets are exchanged over the wireless channel.
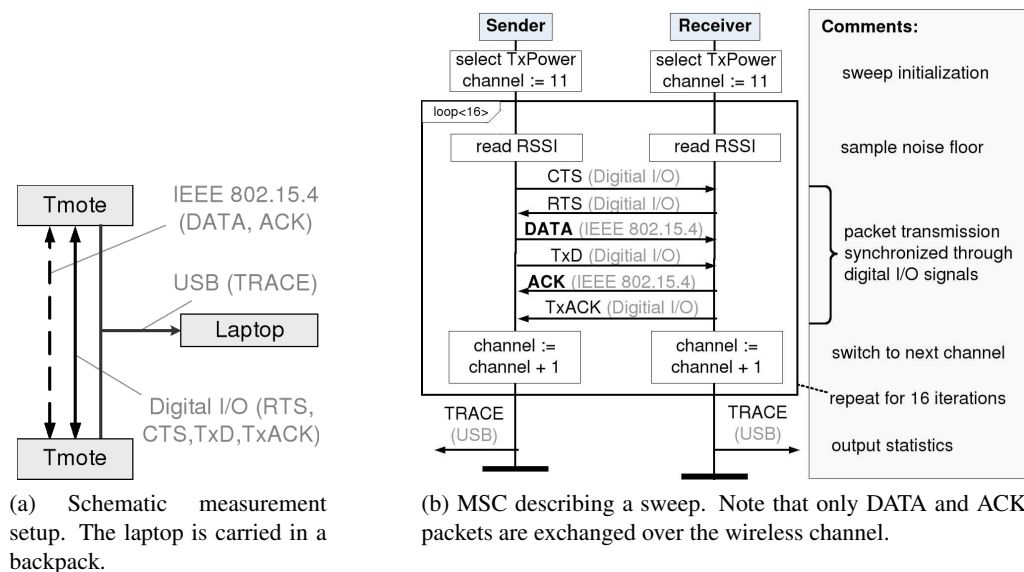
Figure 6.13: Measurement setup and MSC for a single sweep.

shielded cable we interconnect the two MCUs via four digital I/O ports. To increase robustness we use differential signalling (complementary signals sent on two separate wires) and thus allow four different signals, which is sufficient to synchronize the transmission of DATA and ACK packets. The sequence of exchanged digital I/O signals is shown in Figure 6.13b, the vocabulary of messages and signals is listed in Table 6.2.

During our study the sensor nodes collect large amounts of statistics that, due to limited memory, cannot be stored on the nodes themselves. Instead the nodes are connected to and continuously output the results to a laptop through the USB interface of the Tmote Sky, which also serves as the power supply. The laptop is carried in a backpack by the same person that wears the two sensor nodes and is also used to monitor 802.11b/g traffic on selected WLAN channels during the experiments.

In our setup the time required for a sweep over 16 channels is about 87 ms, which results in around 12 sweeps (384 packets) per second including streaming the measurement results over USB. This is achievable because we increase the Tmote Sky CPU frequency to the maximum of 8 MHz and because synchronization over digital I/O is very fast. We took particular care to minimize the impact of streaming the statistics over the USB on the actual measurement and its periodic workflow: most operations related to the transmission of 802.15.4 packets and synchronization occur in interrupt context, while sending serial packets over USB is divided in many small blocks of code that are executed in non-interrupt context. All outgoing/incoming packets are timestamped. After the measurement we use the hardware generated timestamps of successful transmissions as anchors, perform a linear regression to cancel out the clock drift, and obtain precise timing information to verify that in our setup 802.15.4 data packets are indeed transmitted periodically with an average interarrival time of around 5.5 ms. For example, in the experiment described in Sect. 6.4.2 we determine an average interarrival time of 5.43 ms

114

Message/ Signal	Channel	Description
DATA	802.15.4	DATA packet (MPDU of 36 byte)
ACK	802.15.4	ACK packet (MPDU of 5 byte)
RTS	Digital I/O	Sender requests to send a DATA packet
CTS	Digital I/O	Receiver is ready to receive DATA packet
TxD	Digital I/O	Sender has sent a DATA packet
TxACK	Digital I/O	Receiver has sent an ACK packet
TRACE	USB	Measurement results for a sweep over 16 channels

Table 6.2: Messages and signals exchanged during the measurements.

(maximum: $6.45\,\mathrm{ms}$, minimum: $4.39\,\mathrm{ms}$).

We made measurements in three different urban environments: at a shopping street, in a central residential area and in an office area. During all measurements the test person was walking outdoors on the sidewalk of an urban streets. A single measurement typically lasted around 30 minutes.

6.4.2 Results

In the following we first report on one measurement in detail and afterwards present an evaluation of the data from all measurements.

Correlation with 802.11b/g Traffic

The measurement described in the following was made at a central urban shopping street. Buildings were located on either side of the roughly $30\,\mathrm{m}$ wide street, which was moderately frequented by cars and other pedestrians on a weekday evening at 7 p.m. The person carrying the WBAN took a 30 minute walk outdoors along the pavement passing by shops, coffee bars and offices as well as other pedestrians. The walk was one-way and close to straight-line at even walking speed (stopping only at red traffic lights).

Figure 6.14 shows failed 802.15.4 transmissions averaged over 100 transmissions per channel (about $26\,\mathrm{s}$) sorted by transmission power. The losses for $-10\,\mathrm{dBm}$ transmission power (top) are negligible, even at $-25\,\mathrm{dBm}$ (middle) we never see more than $60\,\%$ loss within a window of $26\,\mathrm{s}$ on any channel. At $-42\,\mathrm{dBm}$ (bottom) transmissions failed more frequently and some channels were temporarily completely blocked. The figure suggests that losses were not completely random. Instead they showed some correlation in time and frequency, sometimes lasting for a few tens of seconds up to multiple minutes and spanning over multiple consecutive 802.15.4 channels.

Figure 6.15 shows a 2 minute excerpt of the bottom Figure 6.14 at around $19\!:\!03\,\mathrm{h}$. In this figure transmissions are averaged over a window of 10 transmissions ($2.6\,\mathrm{s}$). When comparing these two figures we find that in Figure 6.15 the pattern at around $19\!:\!03\,\mathrm{h}$ to $19\!:\!04\,\mathrm{h}$ on channels
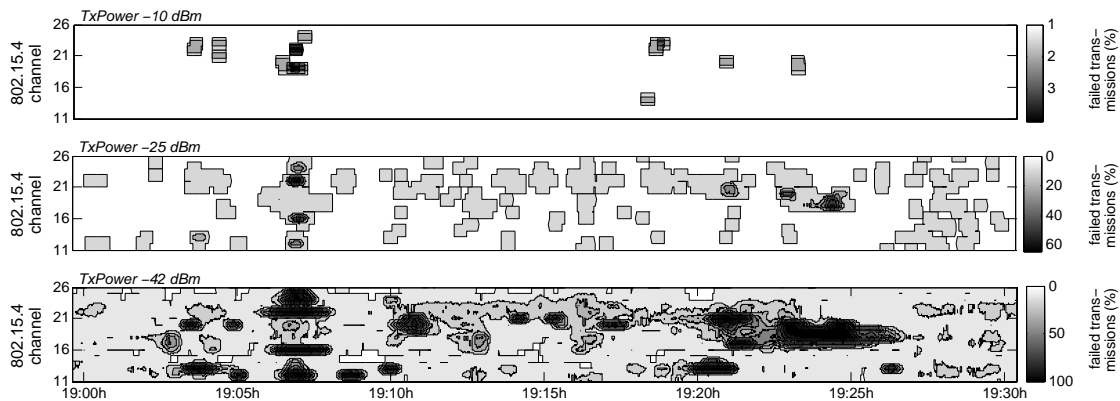
Figure 6.14: Failed 802.15.4 transmissions while walking along an urban shopping street.

Figure 6.15: Two minute excerpt from bottom Figure 6.14, averaged over 10 transmissions.

13 closely resembles the 802.11b "footprint" in Figure 6.3a, which suggests that these losses might have been caused by 802.11 traffic.

During the experiments the 802.11b/g card of the laptop that collected the statistics from the nodes was set to passive monitoring mode so that all 802.11 traffic on a given channel was captured. We tuned the card to 802.11 channel 7, because it is one of the most commonly used. In this way we measured two things in parallel: 802.15.4 failures on all channels and 802.11b/g traffic on 802.11 channel 7.

In Figure 6.16 one can see both, failures on 802.15.4 channel 18 using transmission power $-42\,\mathrm{dBm}$ and the number of received 802.11 packets on channel 7, averaged over 100 transmissions in the 802.15.4 network ($26\,\mathrm{s}$). As shown in Figure 6.1, 802.11 channel 7 completely overlaps with 802.15.4 channel 18. The results appear (negatively) correlated both, visually and statistically. The empirical correlation coefficient is r = -0.89, which when squared is 0.79 and describes the proportion of variance in common between the number of 802.15.4 failures and received 802.11 packets when averaged over a window of $26\,\mathrm{s}$. This suggests that in this experiment 802.11 was indeed the cause for considerable packet loss, at least on channel 18.

When we repeated the measurement at other locations we observed less correlation (correlation coefficients around r = -0.4). It must be noted, however, that the correlation coefficients indicate only linear dependency and that the number of received 802.11 packets is a rather coarse metric because it does not take, for example, packet size or signal strength into consideration.

Figure 6.16: 802.15.4 transmission failures on channel 18 using transmission power $-42\,\mathrm{dBm}$ (left Y-Axis) and 802.11b/g traffic on WLAN channel 7 (right Y-Axis) averaged over a window of $26\,\mathrm{s}$.

Figure 6.17: Successful transmissions per sweep for the measurements described in Sect. 6.4.2.

Cross-channel Quality Correlation

The measurement results from three different urban environments provided us with a large dataset to explore. In the following we analyze this dataset to determine the interrelation of the transmission quality on different channels at a given point in time. Afterwards, we examine whether trends in the noise floor can indicate a (future) decrease of transmission quality. The first helps understanding to what extent channel hopping could improve transmission quality, the second gives insight in a potential (dynamic) trigger for channel switching. Our evaluation is an important starting point, but also has its limitations: it is trace-driven and therefore the results are applicable only for the particular environments, the specific mobility pattern and node placement.

The first topic we address is the correlation of transmission failures over different channels at (roughly) the same time. Little correlation would mean that they share little variance and thus there is a greater probability that a "good" channel is available at a given point in time. In the following we first examine the fraction of good channels over time, then report on the empirical correlation between the channels and finally evaluate post factum how many hops an ideal channel hopping scheme would have required to achieve $0\,\%$ transmission failures.

Figure 6.17 shows the number of successful transmissions per sweep over time for the measurement described in Sect. 6.4.2 (it was the environment where we observed most transmission failures). At $-10\,\mathrm{dBm}$ and $-25\,\mathrm{dBm}$ transmission power for the majority of time the transmissions succeeded on almost all channels. Even for $-42\,\mathrm{dBm}$ the number of good channels rarely dropped below 10, only at around 19:07 h temporarily 13 out of the 16 available channels were

117

blocked. This means that in principle there were enough good channels at any point in time.

From the figure one cannot conclude how similar the different channels behave in time. We calculated the empirical correlation coefficients for the number of failed transmissions over time between all channels and found that they are often very low (typically around zero). For example, for $-42\,\mathrm{dBm}$ transmission power the maximum correlation coefficients for any two channels in the three measurements were $r^2 = 0.1844$, $r^2 = 0.1052$, and $r^2 = 0.0336$, respectively.

We are interested in the minimum number of channel switches that would have been required to achieve $0\,\%$ transmission failures. In other words, the minimum number of hops that an ideal channel hopping scheme with precise knowledge of the future channel conditions would have taken to guarantee that all transmissions succeeded. We implemented a simple greedy algorithm that replays the empirical traces, starting from the first transmission and proceeding in time. The algorithm examines all transmissions and chooses the channel that provides $0\,\%$ transmission failures for the maximum time ahead. It proceeds in time on this channel until a transmission failure occurs and switches the channel during this very sweep, again by choosing the channel that provides $0\,\%$ transmission failures for the maximum time ahead (from the previous evaluation we know that there is always a "good" channel). It repeats this step until it reaches the end of the empirical trace and then outputs the overall (minimum) number of channel switches (hops).

When the algorithm replayed the traces from the three measurements the total number of channel switches for $-10\,\mathrm{dBm}$ transmission power was zero for all three empirical traces, which means there was at least one channel on which no failures occurred, respectively. At $-25\,\mathrm{dBm}$ transmission power two or three hops were required. And at $-42\,\mathrm{dBm}$ the measurement described in Sect. 6.4.2 required 50 hops (note that Figure 6.14 shows moving averages over 100 transmissions), resulting in an average hopping frequency of about $2\,\mathrm{hops/min}$. The empirical results from the residential and office area corresponded to 27 and 38 channel switches, respectively.

Prediction of the Link Quality Degradation

A brief examination of the measurement results revealed that a substantial increase in noise floor is typically accompanied by heavy packet loss, in particular at $-42\,\mathrm{dBm}$ transmission power. For example, Figure 6.18 shows failed transmissions over time (top) together with a contour plot of the noise floor (bottom) averaged over 100 transmissions; this measurement was made in an urban residential area on a Friday afternoon.

We are interested in the dynamics of RF noise around the point in time when the quality of a communication link experiences significant degradation. For $-42\,\mathrm{dBm}$ transmission power we define the threshold as $\geq 90\,\%$ transmission failures within 100 transmissions on a given channel, that means $\geq 90\,\%$ failures within $26\,\mathrm{s}$. An analysis of the empirical traces reveals 8 (shopping street), 11 (residential area) and 2 (office area) = a total of 21 occurrences of significant link quality degradation. For every occurrence we extract RF noise samples on the given channel within a window of ± 2 minutes (roughly 1000 noise floor samples and corresponding to a RF noise sampling frequency of $4\,\mathrm{Hz}$). Note that in our analysis we count only the first occurrence per channel (otherwise there was a total of 26 occurrences) and we ignore channels
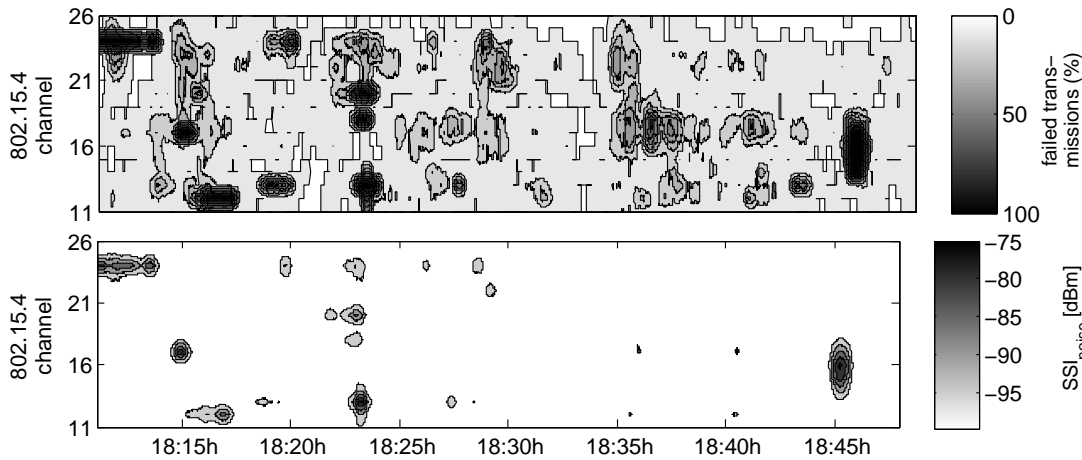
Figure 6.18: Failed 802.15.4 transmissions at $-42\,\mathrm{dBm}$ transmission power while walking through a central urban residential area (top). The bottom figure shows the noise floor measured during the same experiment.

that already had significant losses at the beginning of the measurement, because for them the previous 2 minutes are not available.

Figure 6.19a shows the results for $-42\,\mathrm{dBm}$ transmission power. Since we are interested in general trends independent of the environment or a particular channel it includes the results from all measurements, one graph for every occurrence of link quality degradation. Each graph is aligned relative to the *beginning* of significant link quality degradation, which is represented by the vertical dashed line at time $t = 0\,\mathrm{s}$. This means, starting from $t = 0\,\mathrm{s}$ the next 90 or more out of 100 transmissions failed, respectively. RF noise is arithmetically averaged over the *past* $26\,\mathrm{s}$ (prior moving average). A single point on one of the 21 graphs thus includes the current as well as the history of 99 previous noise floor samples.

At $-25\,\mathrm{dBm}$ output power transmission failures were less frequent and we therefore reduced the threshold to $\geq 50\,\%$ transmission failures within a window of $26\,\mathrm{s}$. This resulted in a total of 5 occurrences of link quality degradation as shown in Figure 6.19b.

The graphs show similar trends: at time $t = -120\,\mathrm{s}$ the average (past) noise floor is typically below $-95\,\mathrm{dBm}$, respectively. Between $t = -120\,\mathrm{s}$ and $t = 0\,\mathrm{s}$ it increases slightly and the most substantial increases are observable around and especially short after $t = 0\,\mathrm{s}$, which is also the *start* of significant transmission failures. Right-shifting the dashed line by $13\,\mathrm{s}$ as well as left-shifting the noise floor graphs by $13\,\mathrm{s}$, respectively, establishes temporal alignment between the moving averages. It is then clearly observable that significant link quality degradation corresponds with a *simultaneous* increase in average noise floor, which strongly suggests that external RF interference is indeed causing the packet loss. The graphs are typically bell-shaped, which is likely a result from walking past a stationary interferer (for example, a WLAN access point).

The results suggest also that (the history of) noise floor observations may be valuable input to a link estimator. Especially at lower transmission frequency trends in the noise floor may be observable before a communication link experiences significant degradation. Increasing the

119

(a) $-42\,\mathrm{dBm}$ transmission power: between $0\,\mathrm{s}$ and $26\,\mathrm{s}$ $\geq 90\,\%$ transmissions failed.

(b) $-25\,\mathrm{dBm}$ transmission power: between $0\,\mathrm{s}$ and $26\,\mathrm{s}$ $\geq 50\,\%$ transmissions failed.

Figure 6.19: Noise floor within a window of ± 2 minutes around heavy transmission failures (starting at 0s).

noise floor sampling frequency and using more elaborated statistical techniques than a simple moving average are likely to have better predictive quality. We consider these topics part of our future work.

6.4.3 Discussion

The results reported in this section were obtained with the Tmote Sky platform; as discussed in Section 6.2.3 we expect packet loss results to be different (more severe) for the Shimmer2 platform. Yet, the effects that we observed in the isolated baseline measurements described in Section 6.2 could also be identified in urban environments: transmission failures sometimes span over multiple consecutive 802.15.4 channels, they are often correlated in time and substantial losses are typically accompanied by an increase in the noise floor. We could also show that WLAN is indeed the cause for considerable WBAN packet loss, which confirms the assumption from Section 6.3.2. Significant losses occurred only when using low transmission power: already at $-10\,\mathrm{dBm}$ transmit power the amount of losses was negligible. This confirms the assumption that the overall effect of RF noise on intra-WBAN communication reliability may be less severe than the effects of fading — we will revisit this question in the following section.

6.5 Analyzing the Causes of Packet Loss

From the findings in Section 6.3 it was not immediately clear what impact urban RF noise has for the reliability of an intra-WBAN communication link. From the results presented in Section 6.4 it is also not possible to quantify the amount of losses *only* due to RF interference. Unfortunately, in practice it is indeed hard to differentiate between the amount of packets lost due to interference vs. fading: typically for lost packets we neither have information about RSSI nor can we measure the amount of simultaneous RF noise. Understanding the cause of packet loss is, however, quite important, because it helps to apply a suitable countermeasure (e.g. to change channel in case of interference or to wait until fading conditions have improved in case

Figure 6.20: The CC2420 SINR/PLR relationship from TOSSIM.

of a weak signal). Based on our RF noise measurement results reported in Section 6.3 in the following we would like to answer the question: *what is the ratio of packet loss caused by RF interference to the losses caused by a too weak (below sensitivity threshold) signal?*

To answer this question we performed a statistical analysis in which we used the data collected in the experiments described in Section 5.6 ("fading") and Section 6.3.2 ("interference"). Note that for the following analysis we assume that fading and RF noise are statistically independent. We also assume that the fading characteristics reported in Section 5.6 are applicable in the scenarios in which we performed RF noise measurements (although this assumption may not hold for static indoor scenarios listed in Table 6.1).

Our analysis is based on the relationship between SINR and Packet Loss Rate (PLR) for the CC2420 [69] radio. PLR is the ratio of the number of packets that were not successfully received to the number of transmitted packets, i.e. $PLR = 1 - PRR$. We use the empirical SINR/PLR relationship based on the measurements reported in [29], which is also part of the radio model of the TinyOS simulator TOSSIM [179]. The corresponding graph is shown in Figure 6.20 (recall that the SINR/PRR relationship has been shown in Figure). We abbreviate the formula that expresses the relationship shown in Figure 6.20 by $PLR(SINR)$ and estimate the ratio of packets that are lost **only** due to external RF interference as follows:

$$PLR_{Interference} = \sum_{x=P_{thr}}^{P_{max}} \left(P(P_{Signal} = x) \sum_{y=P_{min}}^{P_{max}} \left(P(P_{Interference} = y) \, PLR(x-y) \right) \right)$$

(6.1)

where P_{thr} is the sensitivity threshold of the radio (-94 dBm for the CC2420),[7] P_{min} is the minimum detectable P_{Signal} value (we use a value of -100 dBm) and P_{max} is the maximum detectable P_{Signal} value (we use a value of 0 dBm). $PLR(x-y)$ is the PLR for a given SINR $x-y$ (x and y are in dBm) based on the relationship shown in in Figure 6.20. $P(P_{Interference} = y)$ is the empirical probability that RF noise has a value of x and $P(P_{Signal} = x)$ is the empirical probability that the received signal strength has a value of x. We compute $P(P_{Interference} = y)$ from the RF noise measurements described in Section 6.3.2 for each of the 12 scenario listed

---

[7]In the CC2420 datasheet the sensitivity threshold is defined as the signal strength resulting in a 1% PLR.
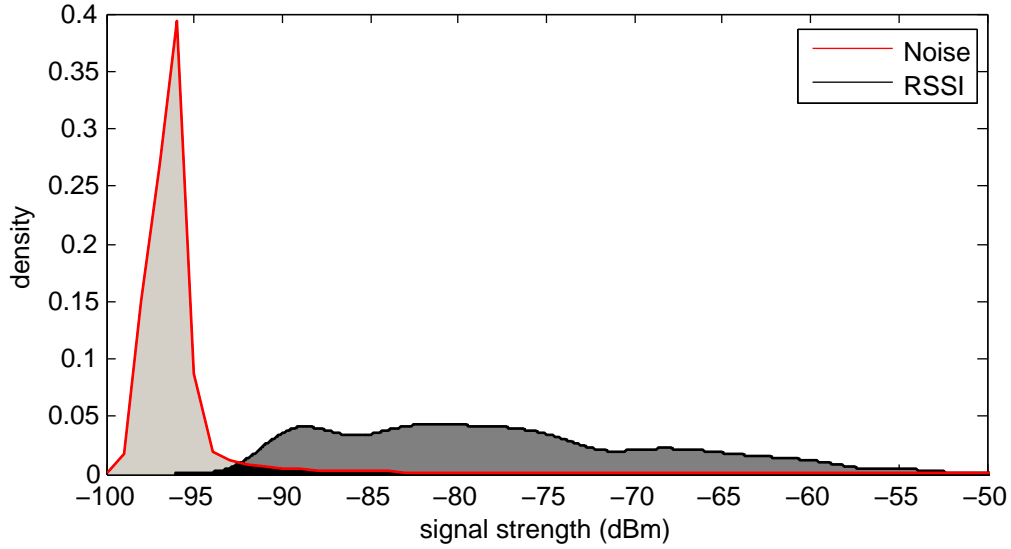
Figure 6.21: Empirical noise (RF noise or noise floor) and RSSI distributions.

in Table 6.1. We compute $P(P_{Signal} = x)$ from the experiment results reported in Section 5.6: recall that in these experiments half of all packets were sent at random points in time based on the standard 802.15.4 MAC. We extract the RSSI of these packets for all subjects, experiments and node positions and use the data to compute $P(P_{Signal} = x)$. The RF noise and RSSI distributions that have been used to compute the probabilities are shown in Figure 6.21. Note that the results are based (only) on the Shimmer2 hardware platform. Note also that $PLR_{Interference}$ in Equation 6.1 considers only the packets above the sensitivity threshold. $PLR_{Interference}$ therefore reflects the expected fraction of packet loss only due to RF interference.

The results are summarized in Figure 6.22. Considering all scenarios described in Section 6.3.1 on average we expect only 3% of packets to be corrupted only due of RF interference. In contrast, we expect about 25% of losses due to fading (the average PLR for the standard 802.15.4 MAC in Figure 5.8). Figure 6.22 also shows the results per scenario and the best, worst and average channel, respectively. It can be seen that the average amount of losses varies over the scenarios and the (static) selection of the best channel can reduce packet loss by up to 10 percentage points (expt8), respectively. While the results indicate that the overall impact of RF interference on WBAN link reliability may be low, it is important to keep in mind that RF interference occurs in bursts (Section 6.3.2). Therefore, for applications with latency bounds it may still be necessary to mitigate the effects of RF interference — we will explore this topic in Chapter 7.

6.6 Summary

In this Chapter we have studied how urban RF interference affects the reliability of wireless on-body communication. To this end we collected about half a billion RF noise samples in various urban environments with COTS WBAN hardware. Our analysis showed that that there is strong

Figure 6.22: Expected packet loss due to RF interference.

evidence that WLAN is the major source of 2.4 GHz activity and there is a significant correlation of simultaneous activity on neighbouring 802.15.4 channels. One important conclusion was drawn by combining the collected RF noise data with the measurement results that captured the signal strength variation on an intra-WBAN communication link (Chapter 4): we estimate that RF interference will typically be responsible for 3% of the losses in a WBAN, while fading effects (a signal below sensitivity threshold) will cause about 25% of the losses. The results indicate that the overall impact of RF interference on WBAN link reliability is low, however, our study also showed that urban RF noise is about one order of magnitude more "bursty" (correlated in time) than random, i.e. uniformly distributed over time, RF noise. Therefore, even if average spectrum occupation is low, it may be necessary to apply RF interference mitigation techniques, in particular in applications with low latency bounds. In the next section we will investigate two mechanisms that are able to mitigate the effects of RF interference in WBANs.

123

Chapter 7

RF Interference Mitigation

The results from the previous chapter suggest that the overall impact of urban RF interference on WBAN communication reliability is rather small. However, the analysis also showed that RF interference is *bursty*, i.e. correlated in time. When considering average reliability and latency over larger time periods such as hours RF interference may not have a significant impact. But on the order of seconds both, reliability and latency, may temporarily be severely impaired. Depending on the particular application this may be acceptable or not. For example, medical monitoring applications have rather diverse latency requirements: the acceptable time between collection and processing/visualization of data may span from 50 ms for capsule endoscopes [59] over 500 ms for ECG sampling [60] to several minutes for blood glucose monitoring [61]. For vital signs (temperature, blood pressure, pulse or respiratory rate) several tens of seconds up to minutes may suffice [46]. However, in general the quality of the medical care delivery is expected to improve as vital signs can be monitored in real-time [62, 63].

In this chapter, after studying related work, we first analyze to what extent *channel hopping* can mitigate packet errors, and in particular error burstiness, caused by RF interference. Channel hopping switches the carrier frequency in between packet transmissions and thus leverages frequency diversity to avoid interferers that occupy the same portion of the RF spectrum for a longer time. This mechanisms has been integrated in several low-power wireless communication standards (Section 2.3.2), but the implications for the WBAN domain are still not well-investigated. In the second part of this chapter we focus on another techniques that can mitigate the effects of RF interference: packet combining. Packet combining is an error-correction scheme that allows to recover an original packet from multiple (corrupted) copies. Our packet combining technique is an extension to a standard ARQ approach that is used in many existing MAC protocols.

7.1 Related Work

Communication standards for unlicensed frequency bands often include mechanisms to improve coexistence with other wireless systems. For example, the 802.15.4-2006 standard [18] lists in its Annex E the following set of mechanisms: first of all, on the 2.4 GHz physical layer the use of *DSSS* maps each symbol to one of 16 nearly orthogonal pseudo-random sequences, which improves the general robustness against interference. Furthermore, an *energy detection*

primitive allows the next higher layer to determine the amount of energy on a set of channels before network formation; the next higher can then decide to start the network on a channel that is least occupied. In the latest revision of the standard *channel hopping* has been integrated in the MAC [157]: by continuously changing the transmission channel the robustness against potential interferers can be increased. In addition, during the CSMA/CA algorithm a *CCA* may be used to detect if the channel is idle before transmitting. This decreases the probability collisions with other 802.15.4 packets as well as transmissions from other wireless systems. Finally, *low duty cycle*, *low transmit power* and *channel alignment* (avoiding frequencies that overlap with other well-known wireless technologies such as WLAN) are suggested as mechanisms to further enhance coexistence with other wireless devices. The research community has analyzed the suitability of these mechanisms and proposed further ones. For example, Huang et al. [245] have pointed out that the 802.15.4 CSMA algorithm is not well-suited to enable coexistence with WLANs. Others have suggested to dynamically change the 802.15.4 CCA threshold [246] or transmission power [247]. Furthermore, Martelli et al. have investigated dynamic selection of the modulation scheme in WBANs [248]. In the following, however, we focus on work related to the two mechanisms that we investigate later in this chapter: channel hopping and packet combining. For other 802.15.4 interference mitigation mechanisms we refer to survey papers [249, 250].

Channel Hopping Over the past years channel hopping has become an active field of research and several low-power wireless standards have adopted this mechanism in some variant (c.f. Section 2.3.2 and [249]). Channel hopping exploits frequency diversity by switching among different carrier frequencies when transmitting packets. This results in two main benefits: it mitigates the effects of interference and it can reduce adverse propagation effects such as frequency selective fading. The majority of research studies have focused on the first effect. For example, an early experimental study investigated ways to mitigate the impact of WLAN interference on 802.15.4 networks in an office setting by adapting the 802.15.4 communication channel [21]. Targeting stationary networks, their measurement setup was optimized for stable interference configurations. Their results confirm the correlation between 802.15.4 packet loss and 802.11 activity as well as the suitability of noise-based predictors of WLAN interference. Watteyne et al. have conducted a similar study and concluded that channel hopping is generally a suitable means to increase connectivity along communication links [24]. In contrast, the results from a study conducted by Ortiz et al. indicated that channel hopping is typically not necessary in the presence of multihop routing [23]. Stabellini et al. conducted a comparison study of different frequency hopping mechanisms in a small 802.15.4 network and concluded that adaptive schemes are better than using pseudo-random hopping patterns [251]. More recently, authors have analyzed the 802.15.4e channel hopping mechanism TSCH (Time Synchronized Channel Hopping) [20] and proposed improvements [22].

The above studies have addressed channel hopping in static WSN deployments. Although the 802.15.4e, BLE and 802.15.6 standards already include optional channel hopping (c.f. Section 2.3.2), not many studies have focused on channel hopping in the WBAN context. One exception is MBStar [252], a WBAN MAC protocol based on a TDMA approach with very short (2.5 ms) time slots. Channel hopping is based on a fixed hopping pattern and channel blacklisting is also integrated in MBStar. The authors report significant reliability improvements

based on a set of small experiments in controlled interference environments. Yao et al. [253] have addressed channel hopping in WBANs in the context of security: they exploit symmetric channel fading to improve secret key entropy. Closest to our work, however, is a recent paper entitled "What can be gained by channel adaptation?" [254]. To answer this question the authors analyzed the potentials of different channel hopping schemes in WBANs under WLAN interference. The investigated schemes are: static (always the same channel), adaptive (channels are switched periodically/pseudo-randomly or based on RF noise measurements) and ideal (assuming perfect knowledge of channels). The evaluation is performed in the Castalia simulator [255], where WLAN traffic is simulated based on a simple interference model. WLAN access points are placed at random locations and the WBAN passes by on a straight line. The results show that under heavy WLAN traffic in particular the adaptive (sensing-based) scheme can greatly increase reliability and it comes very close to the performance of the ideal scheme. The main difference to our work is that rather than modelling WLAN traffic in the simulator we feed our RF noise traces into the simulator and thus analyze channel hopping under real-world interference characteristics.

Packet Combining In packet combining schemes a receiver stores erroneous packets to combine them with parts from later trials [256]. A key issue is to identify the incorrect parts of a packet. In coded transmission systems it is possible that the decoder delivers not only decided bits, but also additional information specifying the confidence in its decisions on a per-bit basis (soft-information). In [257] soft-information is used in a selection-decode-and-forward relaying scheme. In the absence of true soft-information other methods are needed to infer the positions of the correct and incorrect parts of a packet. One method is to insert additional checksums into the packet [258].

The RSSI profiling technique discussed in Section 7.3 has not yet been proposed for usage in ARQ schemes, but RSSI has been used in other "non-traditional" contexts: Demirbas et al. [259] proposed to use the RSSI power-sum of acknowledgement frames in the process of estimating the number of neighbors for which a certain predicate is true. B-MAC [260] samples RSSI with high frequency to achieve a more accurate clear channel assessment; and in [261] per-bit RSSI information has been used (in conjunction with other per-bit/per-symbol information) to gather information about the cause of packet losses in Wifi systems and to adapt the MAC/PHY parameters accordingly: when a collision is inferred, the backoff process of the MAC is triggered whereas a weak signal results in an execution of the rate- or power-adaptation algorithm.

7.2 Channel Hopping

Channel hopping is a mechanism that switches carrier frequencies between or during transmissions. It thus spreads the risk of encountering bad conditions on one particular channel over the available frequency region. The main advantages lie in mitigating the effects of interference and weakening adverse propagation effects on certain frequencies. In this thesis, however, we focus only on the first effect, i.e. we analyze the potential of channel hopping to mitigate the effects of RF interference in the context of WBANs.

Channel hopping can be implemented in different ways. The most common variant switches communication channels based on a pseudo-random sequence of channels known to both trans-

mitter and receiver. The switching of the carrier frequency can happen during transmission (after a certain number of transmitted symbols) or after the transmission of entire packet(s). Rapid changes of the carrier frequency is also referred to as FHSS and has been applied in the military domain, because the signal is hard to intercept and jam. But FHSS is also used in civilian technology such as Bluetooth, which switches channels every 625 μs.[1]

Due to their relative simplicity pseudo-random hopping patterns are integrated in low-power standard such as Bluetooth, 802.15.4e and 802.15.6 (Section 2.3.2). In addition, over the past years dynamic channel hopping techniques have also been considered: rather than following a pre-computed sequence of channels, dynamic mechanisms can update the hopping sequence during network operation based on present channel conditions. To this end the channel quality is monitored and the hopping sequence is modified when applicable. This technique is, for example, often used in cognitive radio based wireless networks [262]. Another related technique is to exclude certain channels from the set of available channels, which is also referred to as *blacklisting*. For example, with AFH Bluetooth avoids crowded frequencies in its hopping sequence. Blacklisting can be performed statically (before network formation) or as part of dynamic approaches.

In this thesis we focus on one specific channel hopping mechanism that is part of the latest 802.15.4e standard [157], in particular, we analyze the TSCH mechanism. We believe TSCH is a representative candidate for the WBAN domain, because like 802.15.6 and BLE it uses a pseudo-random hopping sequence and it changes carrier frequency only after packets have been transmitted (rather than during transmissions like "classical" Bluetooth). Furthermore, TSCH also supports blacklisting of channels. The central goal of this section is to analyze to what extent the TSCH hopping pattern can mitigate the effects of RF interference in comparison to single-channel approaches. In particular, we focus on the short-term effects on latency and reliability, because the previous chapter indicated that over large time periods reliability is not degraded significantly by urban RF interference. To have an upper bound, our investigation also includes an "ideal channel hopping" approach, which is explained below.

7.2.1   Analyzed Variants

We consider three different types of channel hopping: (a) the TSCH hopping pattern with and without blacklisting, (b) a static approach that always remains on the same channel, and (c) an "ideal" hopping pattern, which always selects the channel with the least amount of RF noise. In the following we provide a brief overview of these variants.

TSCH

TSCH is a channel hopping variant introduced in the latest revision of the 802.15.4 standard [157]. It originates from the Time Synchronized Mesh Protocol (TSMP) protocol [263], which also founds the basis of WirelessHART [155]. TSCH is a TDMA approach that divides time into discrete timeslots, which should be large enough so that one data frame and one acknowledgement frame can be exchanged between a pair of devices. A timeslot is either idle (unused)

---

[1]Note that BLE switches channels much slower than classical Bluetooth as explained in Section 2.3.2.

	Slotframe 2	Slotframe 3	Slotframe 4	Slotframe 5	Slotframe 6	Slotframe 7	...

| ASN=39	ASN=40	ASN=41	ASN=42	ASN=43	ASN=44 ... 51	ASN=52	ASN=53	ASN=54	ASN=55	ASN=56	ASN=57 ... 64
Timeslot 0	Timeslot 1	Timeslot 2	Timeslot 3	Timeslot 4		Timeslot 0	Timeslot 1	Timeslot 2	Timeslot 3	Timeslot 4	
Left Wrist	Right Wrist	Left Ankle	Right Ankle	Chest	Timeslot 5 ... 12	Left Wrist	Right Wrist	Left Ankle	Right Ankle	Chest	Timeslot 5 ... 12
→ Right Hip	→ Right Hip	→ Right Hip	→ Right Hip	→ Right Hip		→ Right Hip	→ Right Hip	→ Right Hip	→ Right Hip	→ Right Hip	
CH=14	CH=18	CH=13	CH=22	CH=21	Idle	CH=16	CH=12	CH=17	CH=14	CH=18	Idle

Figure 7.1: An example TSCH configuration.

or it represents a directed/shared link between two/more devices on a specific channel. When a device has data to send to another device, it thus has to wait until a suitable slot has started. The assignment of slots and channel to devices is, however, beyond the scope of the standard and has to be performed by the next higher layer.

Multiple timeslots are contained in one *slotframe*, and a slotframe continuously repeats in time. The channel associated with a timeslot depends on the Absolute Slot Number (ASN), which represents the total number of timeslots that has elapsed since the start of the network. The channel switches are based on a pseudo-random pattern defined by the following formula [157]:

$$CH = hoppingSequenceList[(ASN + channelOffset) \% hoppingSequenceListLen]$$
(7.1)

where *hoppingSequenceList* is a list of pseudo-randomly shuffled channels, *channelOffset* is an optional offset for *hoppingSequenceList* and *hoppingSequenceListLen* is the length of *hoppingSequenceList*. The next higher layer has to assign values to these parameters. An example configuration in which *hoppingSequenceList* includes all 16 channels and the slotframe has 13 timeslots out of which only the first five are reserved for different on-body communication links is shown in Figure 7.1. By excluding certain channels from *hoppingSequenceList*, a (static) blacklisting approach can be realized.

Static Channel Approaches

The analysis in Section 6.3 has shown that some channels are consistently less occupied than others. Therefore, in our analysis we also consider static approaches, i.e. approaches which never switch the channel. For better comparison we implement these approaches as a special case of TSCH by setting *hoppingSequenceList* to one particular channel, i.e. we adopt all other TSCH functions such as the TDMA approach. We compute the results for all 16 channels, respectively, but for clarity of presentation sometimes report only on the best, worst and average (median) channel.

Ideal Approach

Finally, we consider an "ideal" channel hopping approach, which for each packet transmission selects the best channel. The best channel is the channel for which the receiver senses the least

amount of RF noise at the time of transmission. With our trace-driven simulation method introduced below we have this information available, but in practice this approach is not applicable, because it assumes that the transmitter has direct access to the receiver antenna. Yet it is valuable, because it provides an upper bound for the performance gains of channel hopping. Again, we use the TSCH approach as basis for our implementation, but rather than applying Equation 7.1 to compute the channel, the approach simply selects the best channel.

7.2.2 Simulation Setup

One option to evaluate channel hopping approaches is to perform experiments in uncontrolled urban RF interference environments. This method, however, implies a lack of reproducibility of the experiment conditions. In Chapter 6 we reported on experiments that had been conducted in uncontrolled urban interference environments, which had ensured a high degree of realism, but also showed a large variance among the different experiment scenarios. A fair comparison among channel hopping techniques, however, requires comparable RF interference conditions, which cannot be ensured in uncontrolled environments. Therefore in this section we depart from our previous evaluation method: instead of performing experiments, we analyze channel hopping with the help of the Castalia simulator [255], which we have adapted to make use of our RF interference traces.

Castalia

Castalia [136] is a WSN and WBAN simulator based on the OMNeT++ network simulator. In contrast to other network simulators it models temporal changes with the help of conditional probabilities derived from a large number of WBAN experiments [121]. As explained in Section 2.2.4 the model determines current signal strength based on the last observed value and a set of pre-computed conditional distributions. While this allows to model the effects of fading, Castalia does not include a model for external RF interference: at the heart of Castalia's radio module the SINR model is applied to compute whether a packet has been received correctly, but the only possible source of interference are intra-WBAN packet collisions. Therefore we have integrated the large RF noise dataset reported in Section 6.3 into the Castalia simulator. Rather than creating a new interference model, which is beyond the scope of this thesis, in our simulations we "replay" the RF interference characteristics captured in the different experiment scenarios listed in Table 6.1. Since our RF noise traces include data for all channels with a resolution of 142 Hz we can thus precisely recreate a particular RF interference environment in the simulator. This *trace-driven* simulation technique naturally captures the temporal RF interference characteristics, which are pivotal to our analysis. A single simulation run thus corresponds to a specific experiment trace and the total simulated time is bounded by the duration of the traces. The disadvantage is that we have a limited total time of 13 hrs simulated time. On the other hand, we are able to compare different approaches under identical, realistic RF interference conditions.

Property	Value
available channels (except for TSCH-BL)	11..26
available channels (TSCH-BL)	11,20,26
number of timeslots per slotframe	13
timeslot duration	15 ms
scheduled slots	0..4
idle slots	5..12
channelOffset	0
channel access	slottedALOHA
number of MAC retransmissions	0
transmission power	-25 dBm

Table 7.1: Simulation parameters.

Configuration

Castalia does not include the 802.15.4e MAC protocol. Therefore, we have implemented a reduced version of the protocol, which includes the TSCH TDMA channel access and channel hopping method. Since the standard does not provide default parameter values, in our setup we adopt the "Minimal 6TSCH Configuration" parameter set defined in a working document of the IETF [264]. However, we make a few minor modifications: rather than using a slotframe length of 101 timeslots, our slotframe has a length of 13 timeslots. The resulting data rate of roughly 5 packets/s per device ensures that we can analyze reliability and latency over short time intervals. Furthermore, we set the number of retransmissions to zero, because we are interested in how channel hopping manages the immediate channel effects without any additional reliability enhancements. Finally, we only use uplink traffic (devices → coordinator). Note that like suggested in [264] we use slottedALOHA for channel access. In addition to the standard hopping pattern over all 16 available channels we consider a second variant that uses only three channels (11, 20 and 26) which are expected to show little interference (c.f. Figure 6.6). We refer to the second variant as TSCH-BL ("BL" for blacklisting) and to the standard variant that uses all 16 channels as *TSCH*.

Since Castalia's fading model is only available for 6 nodes we use a star topology with 5 devices and 1 coordinator. The coordinator is located at the right hip, which is close to the "right trouser pocket" location we have always selected for the coordinator node. The devices are located at the same positions shown in Figure 3.3a, except the nodes at the back and left pocket are not present, because Castalia does not support these positions. The SINR/PRR relationship in Castalia is slightly different than in TOSSIM. Since we believe that the latter has been validated more thoroughly (and to be consistent with our previous evaluation in Section 6.5) we integrate the SINR/PRR relationship from TOSSIM in Castalia. Also, like the study reported in [254] in our simulations we set the transmission power to -25 dBm. Each device transmits a packet every 195 ms to the coordinator. To this end each device has exclusive access to one 15 ms timeslots in each slotframe. A slotframe consists of 13 timeslots of which the last 8 timeslots are unused as shown in the example in Figure 7.1. The channel associated with a timeslot is calculated based on Equation 7.1, which results in a pseudo-random hopping pattern over all 16 channels,
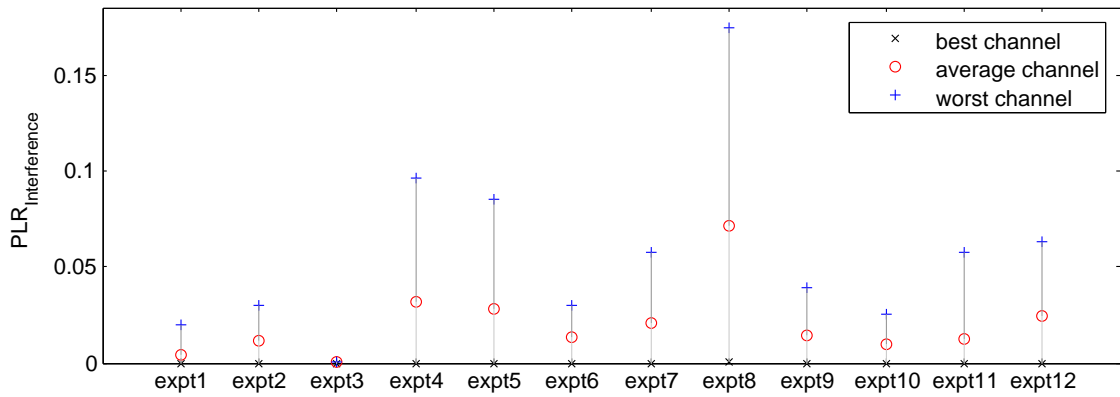
Figure 7.2: Packet loss caused by RF interference per channel.
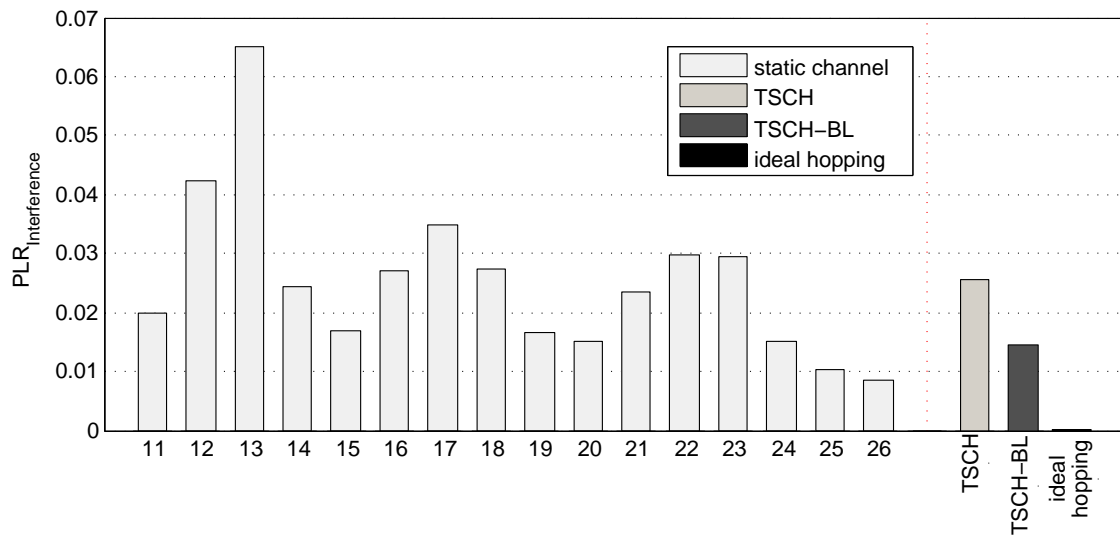
Figure 7.3: Packet loss caused by RF interference per channel hopping variant.

each channel being equally likely (TSCH); or over a reduced set of 3 channels (TSCH-BL). A summary of the simulation parameters used in our setup is shown in Table 7.1.

7.2.3   Results

In this section we report on the performance results of the different channel hopping variants described in Section 7.2.1. Recall that all variants use the parameters listed 7.1 and realize a TDMA as approach shown in Figure 7.1. The variants only differ with respect to what channel is selected at a particular point in time: channel selection is either defined by Equation 7.1 (TSCH and TSCH-BL), the channel is never changed (static approaches) or for each packet from a "god's view" perspective the channel with the least amount of RF noise is selected ("ideal" approach).

We first analyze the amount of packet loss per scenario for the static channel approaches. This allows us to compare the simulator outcome with the analytical results presented in Section 6.5. Figure 7.2 shows the results from the simulation. It can be seen that they match the analytical results (in Figure 6.22) quite well: the results for the individual experiments show a similar trend and on average there are 2.5% of losses caused by only interference (vs. 17% losses caused by a too weak signal). This is only slightly less than estimated in Section 6.5, which used a different fading model.[2]

Next, we analyzed the packet loss rate for each of the channel hopping approaches summarized over all scenarios, respectively. The results are shown in Figure 7.3. Similarly as reported in Section 6.3.2 there is quite some variance among the static channel scenarios and the channels that show the least PLR match the ones far from the popular WLAN center frequencies 1,6 and 11 as in Figure 6.6. TSCH achieves the average PLR over all static channel approaches. This is expected since TSCH spreads the risk of encountering bad conditions over the entire frequency band. Correspondingly, TSCH-BL achieves the average PLR over channel 11, 20 and 26. With 0.003% the "ideal" hopping approach shows a negligible number of packet losses, which indicates that virtually always there is at least one interference-free channel.
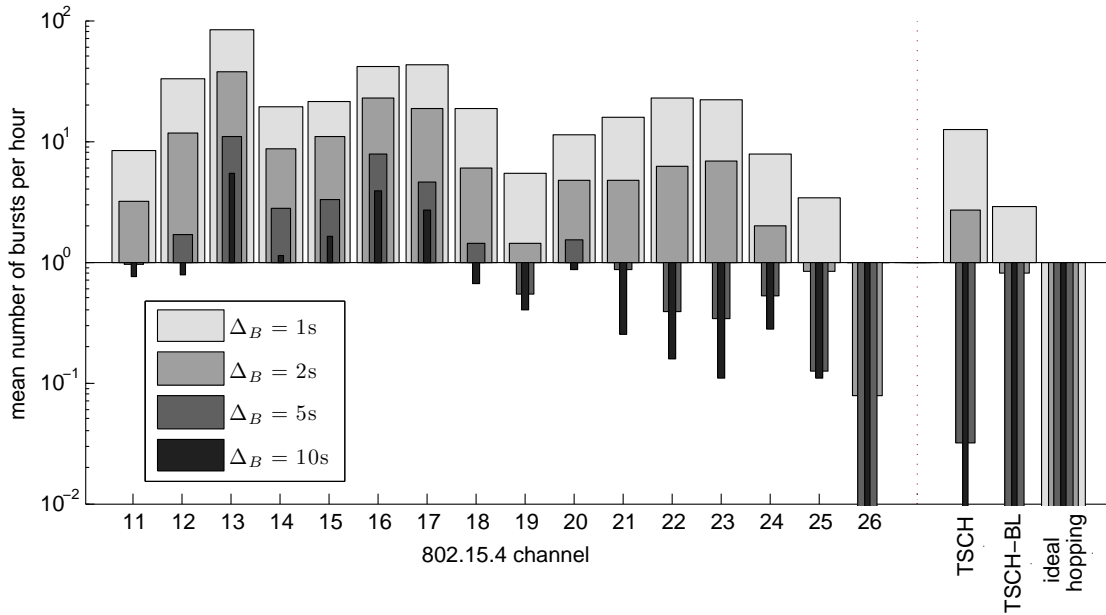
Figure 7.4: Number of packet error bursts caused by RF interference.

Finally, we analyzed the temporal correlation of packet losses due to RF interference. We use a similar notion of "burstiness" as defined in Section 6.3.2: a burst error is a situation when over a window of size Δ_B seconds the average number packets corrupted by RF interference

---

[2]Another reason for this difference may be that in Section 6.5 we also considered the node located at the back of the subject when computing $P(P_{signal} = x)$ for Equation 6.1, while the simulation did not include this node. The node located at the back consistently produced low RSSI readings, i.e. a for this node large number of its packets are expected to be corrupted by RF interference.

Figure 7.5: Duration of packet error bursts caused by RF interference.

exceeds a threshold of 50%. Figure 7.4 shows the results for the mean number of packet error bursts. As expected the static channel approaches show peaks at channels that overlap with the "default" WLAN channels (IEEE 802.15.4 channel 13, 17 and 22). This trend is identical to the analytical results reported in Section 6.3.2 (Fig. 6.10a), but the absolute number of bursts are much lower. This is expected, since an RF noise burst does not necessarily imply a packet error burst. Interestingly, for short time intervals ($\Delta_B = 1s$ and $\Delta_B = 2s$) the results for TSCH are only slightly better than the average results over all static approaches (we provide a possible explanation below). TSCH-BL, on the other hand, outperforms TSCH by almost an order of magnitude, which indicates that even simple channel blacklisting can improve performance significantly. However, the static approach for channel 26 exceeds both, TSCH and TSCH-BL. This is likely due to the fact that channel 26 is generally very quiet, because it does not overlap with any of the default WLAN channels used in the U.S. (which seem to be used in Europe as well). Finally, the "ideal" channel hopping pattern is able to achieve zero packet error bursts for any Δ_B. Based on our data we can conclude that error bursts due to interference can therefore theoretically always be avoided.

Figure 7.5 shows the duration of packet error bursts caused by RF interference. Note that analogous to Section 6.3.2 the duration of an error burst is the period of time over which the average number (over Δ_B seconds) of packet errors exceeds 50%. Therefore, the duration of a burst may be shorter than Δ_B. The results among the static channel approaches are very similar, except the approach for channel 26 shows a significantly lower burst duration. TSCH shows comparable results to the latter and TSCH-BL outperforms both. This means that regarding burst duration TSCH-BL is the best practical choice. The results for the "ideal" channel hopping pattern are not available (set to zero), because as stated above this approach was able to achieve zero packet error bursts for any Δ_B.

7.2.4 Discussion

The PLR results for the static channel approaches (Figure 7.2 and Figure 7.3) confirm the analytical results from Section 6.5: with on average 2.5% PLR the overall impact of RF interference on intra-WBAN communication reliability is rather small. As expected channel hopping can level out differences between channels, which is important, because the PLR among channels can differ by more than a factor three (Figure 7.3). Interestingly, the number of short-term packet error bursts cannot be reduced significantly with TSCH, which indicates that interference is correlated over a larger frequency range.

The main conclusion of our analysis is, however, that currently the best strategy is to simply select channel 26: not only are the number of packet losses significantly lower than on other channels or when applying channel hopping approaches, but even the number and duration of packet error bursts due to interference can hardly be exceeded by other variants. On the other hand, it is expected that channel 26 will eventually exhibit much more inter-WBAN interference, which (presumably) has not played a significant role in our study.

7.3 Packet Combining

In the previous section we have analyzed how frequency diversity may help in mitigating the effects of RF interference. However, even on channel 26 a WBAN will not be able to avoid occasional packet corruption through RF interference. Whenever this happens the 802.15.4 MAC protocol (like many other MAC protocols) makes use of the standard Send-and-Wait ARQ scheme: through the absence of an acknowledgement the transmitter notices a failure in the transmission and after some time retransmits the packet. Often, however, only a few bits have been corrupted in the original packet and therefore the retransmitted packet carries a large amount of redundant information. In this section we study an approach that can reduce the transmission of redundant bits in case a packet has been corrupted by RF interference. By transmitting less bits our approach achieves a reduction in energy consumption, which may, for example, be leveraged to improve communication reliability by increasing the maximum number of retransmissions.

To understand the approach it is helpful to recall that on a wired transmission medium the transmitter may be able to detect a collision at the receiver by monitoring the medium during the transmission, e.g. for an abnormal change in voltage. RF transceivers are usually half-duplex and the SINR conditions at transmitter and receiver can differ greatly. Therefore, transmitter-side collision detection schemes are generally not applicable in wireless communication. However, on the receiver-side RF transceivers can often measure the power of the received radio signal and provide a corresponding RSSI. Previous work in low-power wireless networking has shown that a packet collision (RF interference) distorts the received signal and typically manifests as an additive increase in RSSI [265]. Also, RSSI has proven a relevant parameter when identifying RF interference as the cause of packet loss [261, 266].

In this section we explore to what extent monitoring RSSI with high sampling rate *during packet reception* can be useful not only for receiver-side interference detection, but also for estimating the bit error positions inside corrupted frames. We start with a set of baseline

Figure 7.6: RSSI sampled on a Tmote Sky with 62.5 kHz while receiving a maximum-sized IEEE 802.15.4 frame (133-byte PPDU) with an airtime of 4.256 ms.

measurements to investigate the dynamics of RSSI during controlled collisions and in several environments of realistic, uncontrolled RF interference. Afterwards we introduce and evaluate a simple, threshold-based algorithm that estimates bit error positions with the help of RSSI profiles. Finally, we present an analytical model and empirical results that show the performance improvements of an ARQ scheme when it is coupled with the algorithm. Note that we perform experiments in static testbeds as well as realistic outdoor WBAN scenarios. Testbeds are a convenient and powerful tool for performing long-term measurements with a large number of nodes. Since we are not primarily interested in the WBAN channel conditions but in the performance of our mechanism in realistic WLAN interference environments the majority of experiments is conducted in testbeds.

7.3.1 RSSI Profiling

Like many other radios the CC2420 [69] automatically provides for every received packet an RSSI value, which represents the average signal strength during packet reception. In conformance with the IEEE 802.15.4 standard the CC2420 also allows to obtain the current RSSI by software. Within the radio RSSI is continuously updated and averaged over the last 8 symbol periods (128 μs). This allows to measure ambient RF noise, for example during an IEEE 802.15.4 energy detection scan, but also to obtain an RSSI sample *during frame reception*. When one samples the RSSI while receiving a frame with high frequency — we always use 62.5 kHz, the reciprocal of the time of a symbol — one obtains a time series as depicted in Fig. 7.6. We call such a time series the *RSSI profile* of the incoming frame: a moving average over a window of 128 μs over the received signal strength during packet reception. In practice, the sampling of the RSSI profile of an incoming frame is triggered by the reception of its start-of-frame delimiter (Start Frame Delimiter (SFD)), which results in an interrupt on the MCU. The end of an RSSI profile is marked by the reception of its last byte. Note that in contrast to the illustrative results shown in Fig.7.6 and 7.7, RSSI profiles therefore do not include samples of the noise floor.
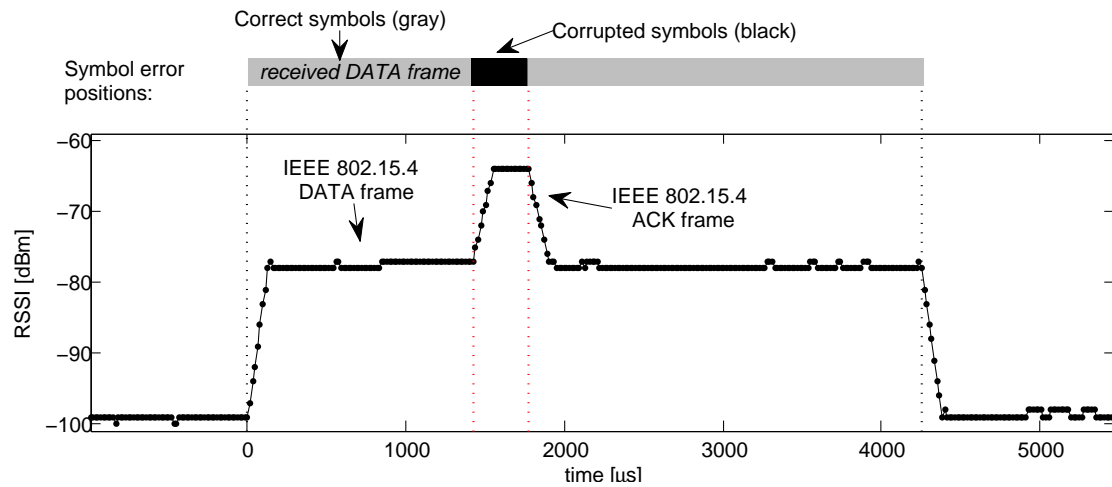
(a) Collision with an 802.15.4 ACK (11 byte PPDU).
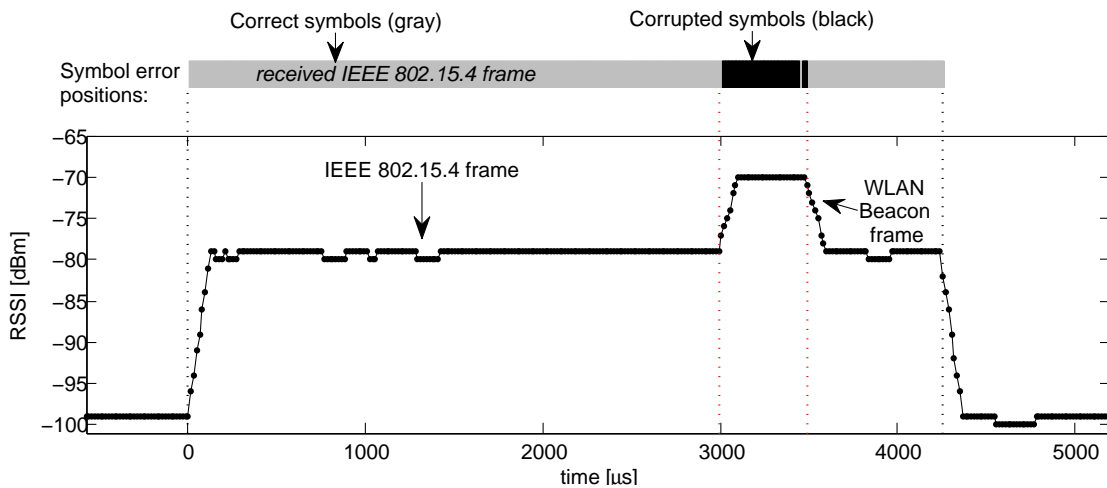
(b) Collision with a WLAN beacon transmitted at 2 Mbps.

Figure 7.7: Bit error positions and RSSI profiles of an IEEE 802.15.4 frame (133-byte PPDU) colliding with an IEEE 802.15.4 ACK (top) and IEEE 802.11 beacon (bottom).

Controlled Collisions

In order to understand the effect of a packet collision for the RSSI profile we created controlled collisions between different types of WLAN and 802.15.4 frames. These experiments were conducted in an environment of negligible external RF interference, which we verified with the help of a portable *Wi-Spy 2.4x* USB spectrum analyzer. The basis for our measurement were two Tmote Sky sensor nodes: one periodically transmitted maximum-sized 802.15.4 frames (133-byte PPDU with pre-defined content), the other listened for incoming frames while it continuously measured RSSI with a rate of 62.5 kHz. Whenever a frame (with a correct or incorrect Cyclic Redundancy Check (CRC)) was received, it output the frame content and its RSSI profile

137

over USB to a laptop. This allowed us to analyze the dynamics in RSSI during reception and compare it with the bit error positions in a corrupted frame.

In a first experiment we used a third Tmote Sky node to act as interferer: it periodically transmited small 802.15.4 DATA or ACK frames without clear channel assessment (CCA) and thus provoked collisions with the frames exchanged between the two other nodes. In a second experiment we replaced the Tmote Sky interferer with a laptop which used its WLAN PC card, based on an Atheros AR5212 chipset, to inject different types of WLAN frames (Beacon, ACK, Request To Send (RTS), Clear To Send (CTS) or DATA). WLAN frames were also sent without CCA and the channels were chosen such that they have maximum overlap. In all experiments the distances between the nodes/WLAN interferer were small ($<5\,\text{m}$) and the interferer was located closer to the receiver than the transmitter. Two representative results can be seen in Fig. 7.7.

The results show that (1) collisions are clearly visible by elevations in the RSSI profile, (2) the duration of an elevation matches the airtime of the colliding frame and, most importantly, (3) the positions of the bit errors in the received frame are temporally correlated with the elevation. Note that whenever there was no frame error (collision) the RSSI profile was typically stable (± 1 dBm) as shown in Fig. 7.6.

Uncontrolled RF Interference

The previous results were obtained in a controlled interference environment. Before analyzing the approach in realistic interference environments we analyze three issues that might exclude the applicability of RSSI profiling in practice. They are related to the following questions:

1. How often does the radio hardware discard corrupted frames before they can be processed and potentially recovered by software?

2. When a corrupted frame is received, how many bits are erroneous?

3. Is there a substantial difference between the RSSI profile of a corrupted and a correctly received frame?

If only a small fraction of the corrupted frames is decoded and forwarded by the radio hardware; if in a corrupted frame typically most bits are erroneous; or if RSSI profiles are very similar regardless of a frame being corrupted or correct, then an error recovery scheme based on RSSI profiling is likely to be of little practical use. In the following we deal with these questions by performing a set of experiments in three different environments of uncontrolled, realistic RF interference. Two of them in a static testbed and one in a WBAN scenario.

The first two measurements were conducted in two publicly accessible indoor sensor network testbeds: TKN Wireless Indoor Sensor Network Testbed (TWIST) [25] and MoteLab [26]. Both testbeds are situated on a university campus and have several WLANs co-located. Each testbed contains a large number of Tmote Sky nodes and we program one of them, located close to the geographic center of the testbed, to broadcast a total of 100,000 frames on channel 21[3],

---

[3]Channel 21 overlaps with (the popular) WLAN channel 10.

Figure 7.8: Each point represents an 802.15.4 link over which 100,000 (TWIST/MoteLab) or 10,000 (WBAN) frames were transmitted. The y-axis represents what fraction of those frames that were not received correctly was still accessible by the receiver MAC protocol.

one frame every 125 ms. All other nodes listen for incoming frames and output every received frame and its RSSI profile over USB to be stored in a trace file for our later examination. Frames have an MPDU size of 64 byte and are transmitted with CCA: whenever the sender determines a busy channel the transmission is deferred for a small random time interval. Since we use only one sender a busy channel can only be caused by other users of the spectrum such as co-located WLAN. Each measurement lasts for about 4 hours and was, in both cases, performed on a weekday in the daytime.

For the WBAN measurements we used a similar setup as described in Section 6.4.1: our setup consisted of two nodes, each node was placed in a thin plastic enclosure and strapped to a person. One node was attached to the left upper arm, the other on the right ankle, resulting in a relative distance of about 1.5 m. The node on the ankle was transmitting a total of 10,000 frames on channel 21 with a transmission power of -25 dBm, the other node was listening and forwarding received frames and their RSSI profiles over USB to a laptop carried in a backpack. Again, the MPDU size was 64 byte and frames were transmitted with CCA. In all six experiments the person was walking outdoors on urban streets for about 30 minutes: a central urban shopping street and streets in a central residential area.

Per setup we calculated for every link, i.e. sender-receiver pair, PRR as the ratio of correctly received to transmitted packets, and Corrupt Packet Reception Rate (CPRR) as the ratio of received corrupt packets (CRC incorrect) to transmitted packets. Thus $CPRR/(1 - PRR)$ describes what fraction of those frames that were not received correctly was still accessible by software. The remaining frames were discarded by the receiver radio supposedly due to a weak signal and/or errors in the PHY header. In our evaluation we only consider links that had a PRR between 0.1 and 0.995 and we only count as corrupt packets those that have the expected length.

Fig. 7.8 shows per-link $CPRR/(1 - PRR)$ for all setups. On bad links sometimes only 20 % of those frames that were not received correctly were accessible by software, on good links typically at least 50 %. One possible explanation is that on bad links signal strength is often below the radio sensitivity threshold and frames are lost due to an insufficiently strong signal.
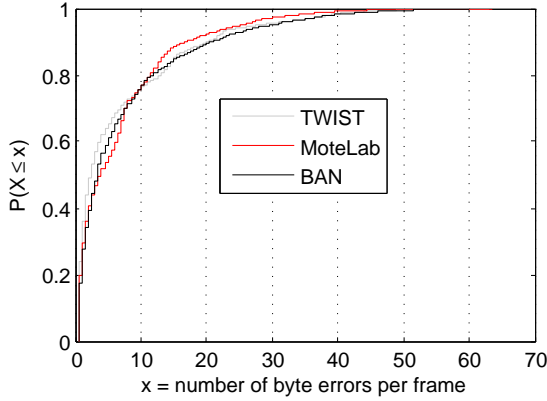
Figure 7.9: Empirical CDFs for the number of byte errors per corrupted frame.
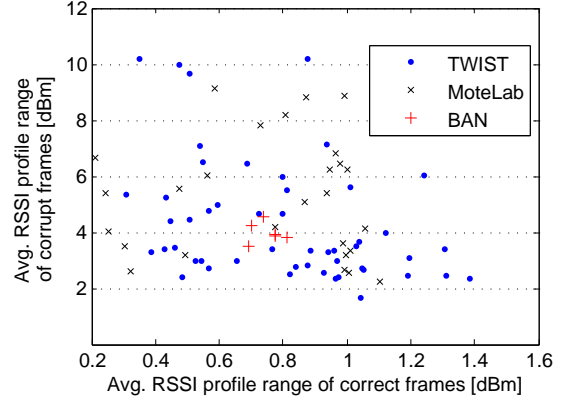
Figure 7.10: Average range of the RSSI profile per link for correct vs. corrupt frames.

The number of byte errors per corrupted frame is shown in the empirical CDFs in Fig. 7.9: on average no more than 5 bytes were incorrect, i.e. less than $10\,\%$ of the 64 byte MPDU. Finally, for every received packet we calculated the range over the 128 samples (2 samples correspond to one byte) contained in its RSSI profile, i.e. the interval between the minimum and maximum RSSI value. We then compared the average range for all correctly received vs. corrupted packets per link. The results are shown in Fig. 7.10 and suggest that typically the RSSI profile of a frame that had at least one bit error had a considerably higher range than the RSSI profile of a frame that was received correctly.

In summary, the results indicate that (1) often a large portion of the corrupted frames is accessible and thus potentially recoverable, (2) the average number of erroneous bytes compared to a medium-sized frame is small and (3) RSSI profiles might contain enough information to estimate the bit error positions correctly.

7.3.2 Bit Error Position Estimation

In the following we examine how to estimate the bit error positions in a corrupted frame based on its RSSI profile. We call an algorithm that performs this task a RSSI-based Bit Error Position Estimation (REPE) algorithm. A REPE algorithm is invoked on the receiver upon reception of a frame with an incorrect checksum (CRC). It takes the RSSI profile of the corrupted frame as input and tries to output an estimate of the bit error positions. Note that in some cases the algorithm will not be able to output a decision, for example when the RSSI profile does not contain enough variance. A good REPE algorithm will maximize the number of decisions, while minimizing the number of false positives (bits classified as incorrect, although they are correct) and false negatives (bits classified as correct, although they are incorrect). It should also impose little computational and memory overhead. In the rest of this subsection we introduce a simple REPE algorithm and evaluate its performance based on the traces we collected in the previous measurements.

The REPE algorithm we propose simply marks all symbols that were received while the

---

Algorithm: REPE-THRESHOLD($RSSI_{Profile}, \Delta_{threshold}, \Delta_{front}, \Delta_{rear}, sensitivity$)

$i \leftarrow 0$

$RSSI_{Base} \leftarrow \mathbf{min}(RSSI_{Profile})$

$symbolErrorMask[length(RSSI_{Profile})] \leftarrow \{0\}$

if $RSSI_{Base} \geq sensitivity$

then $\begin{cases} \textbf{while } i < length(RSSI_{Profile}) \\ \textbf{do } \begin{cases} \textbf{if } RSSI_{Profile}[i] \geq RSSI_{Base} + \Delta_{threshold} \\ \textbf{then } \begin{cases} first \leftarrow \mathbf{max}(i - \Delta_{front}, 0) \\ \textbf{while } (i < length(RSSI_{Profile}) \textbf{ and} \\ \quad RSSI_{Profile}[i] \geq RSSI_{Base} + \Delta_{threshold}) \\ \quad \textbf{do } i \leftarrow i + 1 \\ last \leftarrow \mathbf{min}(i + \Delta_{rear}, length(RSSI_{Profile}) - 1) \\ symbolErrorMask[first..last] \leftarrow 1 \end{cases} \\ \textbf{else } i \leftarrow i + 1 \end{cases} \end{cases}$

return $(symbolErrorMask)$

comment: an empty *symbolErrorMask* (all zeros) means "no decision".

Figure 7.11: The REPE algorithm in pseudocode and an illustration of the terminology.

RSSI was above a certain threshold as incorrect. The threshold is defined relative to the RSSI of the incoming 802.15.4 frame, which we denote as $RSSI_{Base}$. In our scheme $RSSI_{Base}$ is set to the minimum value of the frame's RSSI profile, because previous work has shown that interference (collisions) results in an RSSI increase rather than a decrease [265]. To exclude errors caused by an insufficiently strong signal the algorithm does not output a decision if $RSSI_{Base}$ is below the radio's sensitivity threshold. Before the RSSI rises above and after it falls below the threshold an additional (temporal) safety margin is added. All bits between and including
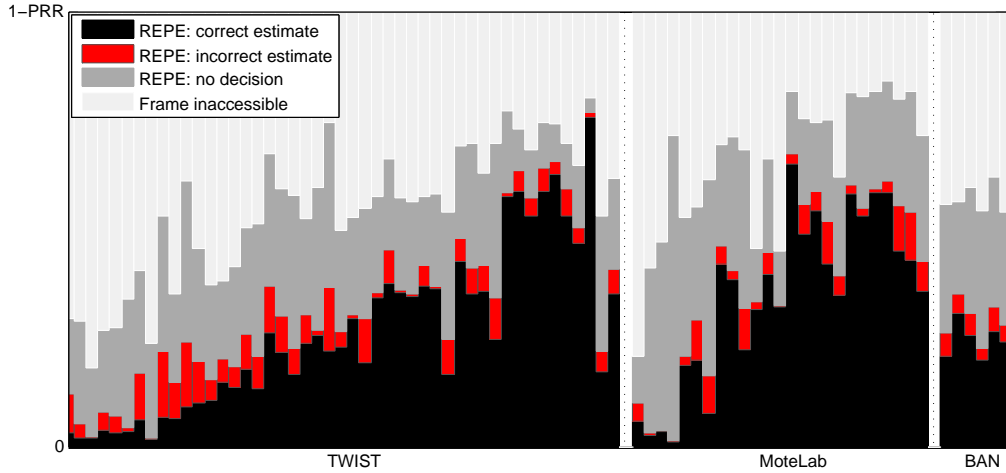
Figure 7.12: Cases for which REPE made a correct, incorrect or no decision.
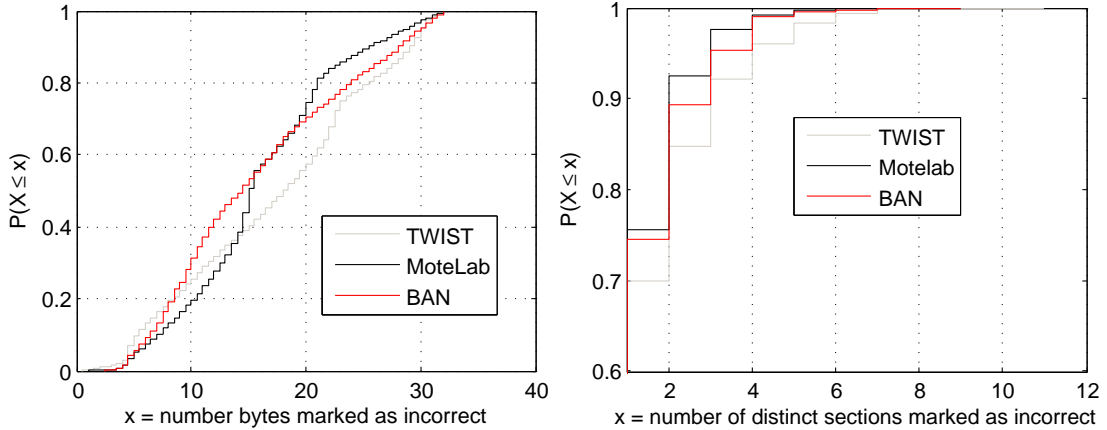
Figure 7.13: Empirical CDFs for the errors estimates.

the safety margins are marked as incorrect and there can be multiple such subsections per frame. Fig.7.11 shows a pseudocode representation of the algorithm and visualizes its notation using an example.

The time complexity of the algorithm is $\mathcal{O}(n)$, where n is the number of samples in the RSSI profile. The memory requirements are n byte to store the RSSI profile — after the REPE algorithm has made a decision the same memory can be used to buffer the corrupted frame.

We evaluated the presented REPE algorithm "offline", with the help of the traces we collected in the measurements described in Sect. 7.3.1. For every received frame our traces contained the RSSI profile as well as the content of the frame. From the latter we could infer the actual bit errors — the "ground truth". The algorithm was instantiated with parameters similar to the ones used in the example shown in Fig. 7.11: $\Delta_{threshold} = 2\,\mathrm{dBm}$, $\Delta_{front} = 6\,\mathrm{symbols}$, $\Delta_{rear} = 2\,\mathrm{symbols}$ and $sensitivity = -93\,\mathrm{dBm}$, which were chosen on the basis of a few trials. For every corrupted frame we let the REPE algorithm try to estimate the bit error subsec-

tion(s) and compared the result with the ground truth. The estimate was correct, if it resulted in no false negative, otherwise it was incorrect. In order to prevent trivial estimates (all bits marked as incorrect) an estimate that in total contained more than half of the MPDU size, i.e. more than 32 byte, was treated as "no decision". We evaluated every link separately and again we only considered links that had a PRR between 0.1 and 0.995.

Fig. 7.12 shows how the REPE algorithm performed for each link. Every link is represented by a bar that is subdivided into four parts, reflecting what happened to those frames that were not received correctly: they were either not accessible by software, or the REPE algorithm made either a correct or incorrect estimate, or it could not make a decision at all. The bars are ordered by increasing PRR and each bar is normalized to $1 - PRR$. It can be seen that the testbed results exhibit strong variance, but there is also a clear trend: on bad links (low PRR, i.e. left side, respectively) typically the vast majority of frames that were not received correctly was either not accessible by software or the REPE algorithm could not output a decision. Also, at low PRR the number of incorrect estimates sometimes outnumbered the correct ones. With growing PRR, however, the proportion of REPE decisions increased to 50% and above, and for $PRR > 0.7$ the ratio of correct to incorrect decisions was typically at least 4. Again, a possible explanation for this trend is that on bad links packet loss is often also caused by a weak signal rather than (only) RF interference. The results for the WBAN scenarios have less variance and are very similar to the average testbed results. Given that the average testbed results seem to match the results in the WBAN scenarios quite well in the following we perform any further performance evaluation in the testbeds (since it is significantly easier to perform long-term measurements with a large number of nodes).

Fig. 7.13 shows CDFs for the number of bytes and distinct subsections marked as incorrect by the REPE algorithm. With 15-20 bytes the average number of bytes classified as incorrect is about 4-5 times higher than the actual average number of byte errors (c.f. Fig. 7.9). This suggests that the proposed REPE algorithm might still be improved in terms of false positives, potentially by better parameter tuning or more fine-grained information from the radio hardware (e.g. shorter RSSI time windows or per-symbol Link Quality Indicator (LQI)/correlation values). According to the bottom graph of Fig. 7.13 the estimated errors were often confined to a single subsection and in about 90% of all cases no more than two subsections were identified. These results confirm the general applicability of the algorithm, in the following we examine one practical application more closely.

7.3.3 REPE-ARQ

There are several application areas that could benefit from the coupling with a REPE algorithm, for example Forward Error Correction (FEC) schemes or techniques that identify/exclude a certain type of interferer technology based on the airtime/inter-frame spacings of colliding frames [267]. In this subsection, however, we explore how a REPE algorithm can improve ARQ protocols [268].

In the standard Send-and-Wait ARQ scheme a frame is retransmitted when a bit error has occurred, typically noticed at the transmitter by the absence of an acknowledgement. We extend this scheme as follows: when a corrupted frame (CRC incorrect) is received, the REPE algorithm is invoked on the receiver side. If the algorithm can make an error estimate, the corrupted frame

Figure 7.14: The NACK frame format.

is buffered and a Negative ACK (NACK) frame is transmitted. As depicted in Fig. 7.14 a NACK is similar to an IEEE 802.15.4 ACK frame, except that it contains at least four bytes of payload, denoting the byte-offset and length of the estimated corrupted subsection(s) in the received frame, and a 16-bit REPE CRC. The number of *REPE offset/length* fields is variable to account for multiple error subsections — alternatively a bitmask could be used to specify the error estimates, but it would typically impose more overhead since the number of error subsections is expected to be small (c.f. Fig. 7.13). The REPE CRC is calculated over the *remaining* presumably uncorrupted portion of the received frame and allows the transmitter to determine whether the REPE algorithm on the receiver side was successful: it simply calculates the CRC over the original frame less the subsection(s) specified by the REPE offset and REPE length, and compares the result with the REPE CRC (all contained in the NACK). If they match, the receiver's estimate was correct, i.e. there was no error outside the estimated corrupted subsection[4], otherwise it was incorrect. In the first case only the corrupted portion (plus a 3-byte 802.15.4 MAC header) is retransmitted, indicated by an unused bit-flag in the header, which we call the R-Flag (we use bit 7 of frame control field). Upon correct reception the receiver is then able to assemble the original frame, pass it to the next higher layer and transmit a final ACK. In case this ACK is not received, the transmitter retransmits the reduced-size frame. If the transmitter discovers that the estimate was incorrect the entire frame is retransmitted. A flowchart of the REPE-ARQ scheme is shown in Fig.7.15a (transmitter-side) and Fig. 7.15b (receiver-side).

Analytical Model

We use a time-homogeneous discrete-time Markov chain model [269] to analyze the performance of the REPE-ARQ algorithm with a maximum of k allowable trials. This model is based on the assumption that different trials to transmit a frame are independent of each other.

The model's state transition diagram is shown in Fig. 7.16. From this diagram the underlying state transition matrix **P** can be directly derived. In the TxFull,i states (where i represents the number of trials) the transmitter sends a full-sized packet, whereas in the TxPartial,i states it sends a partial packet, i.e. retransmits only those parts of the payload that were previously corrupted and correctly identified as such by the REPE algorithm. In the SUCCESS state the transmitter has received an acknowledgement, whereas in the FAIL state the transmitter has exhausted all k allowable trials without receiving an acknowledgement. A transition from

---

[4]Like the IEEE 802.15.4 MAC we ignore the residual error rate of the 16-bit CRC.

(a) Transmitter-side   (b) Receiver-side
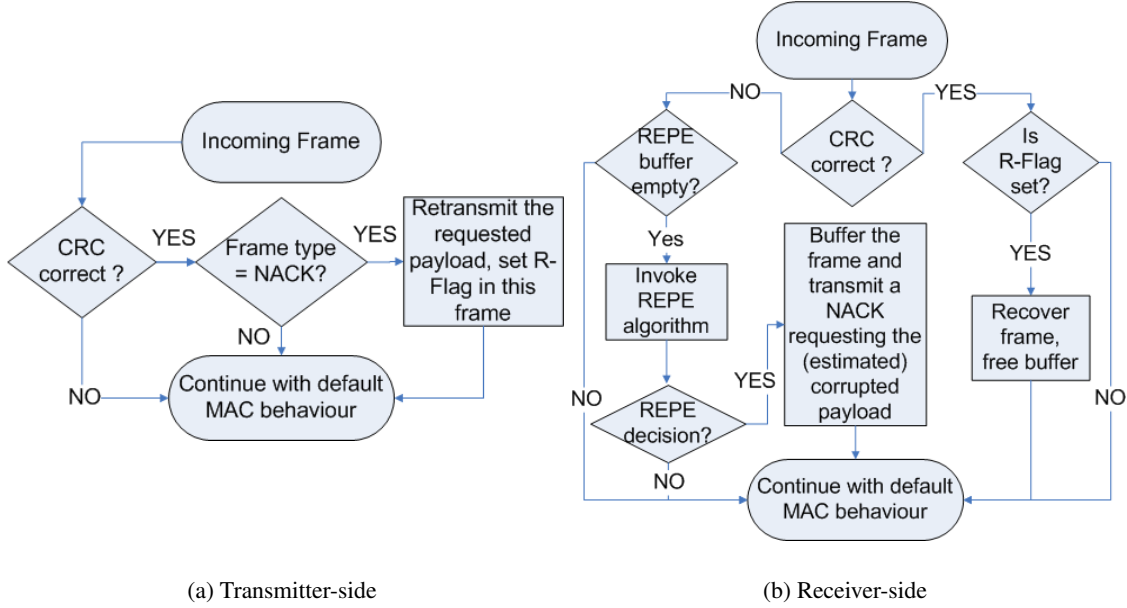
Figure 7.15: The REPE-ARQ scheme.

`TxFull,i` to `SUCCESS` occurs when a full-sized packet and its ACK were received without errors (with probability $p_{F,Succ}$); a transition `TxFull,i` to `TxPartial,i+1` occurs when a full-sized packet with errors was received, the REPE algorithm made a correct estimate and the NACK was received without errors (with probability $p_{F,P}$); and a transition `TxPartial,i` to `SUCCESS` occurs when a partial packet and its ACK were received without errors (with probability $p_{P,Succ}$).

The remaining state transition probabilities can be derived from the three basic probabilities $p_{F,Succ}$, $p_{P,Succ}$ and $p_{F,P}$ in a straightforward way. They are shown in Fig. 7.16. To instantiate the model we obtain the three basic probabilities from measurements (c.f. Sect. 7.3.3).

Our goal is to compute the average total number of bits transmitted within the k possible trials. To achieve this, we utilize the framework of potential theory for Markov chains [269, Sec. 4.2], the relevant definitions and theorems are paraphrased in Appendix B. In our case, all states `TxFull,i` and `TxPartial,i` are inner states, whereas the absorbing states `FAIL` and `SUCCESS` are final states. In the final states no transmission costs are incurred, all states `TxFull,i` have the same cost c_F and all states `TxPartial,i` have the cost c_P.

The average number of transmitted bits in the `TxFull,i` states, c_F, consists of the following components:

- the number of bits in a full-sized data frame, $l_{D,F}$,

- the number of bits in a positive acknowledgement, l_{ACK}, weighted with probability $p_{F,Succ}$, and

- the average number of bits in a negative acknowledgement, l_{NACK}, weighted with prob-
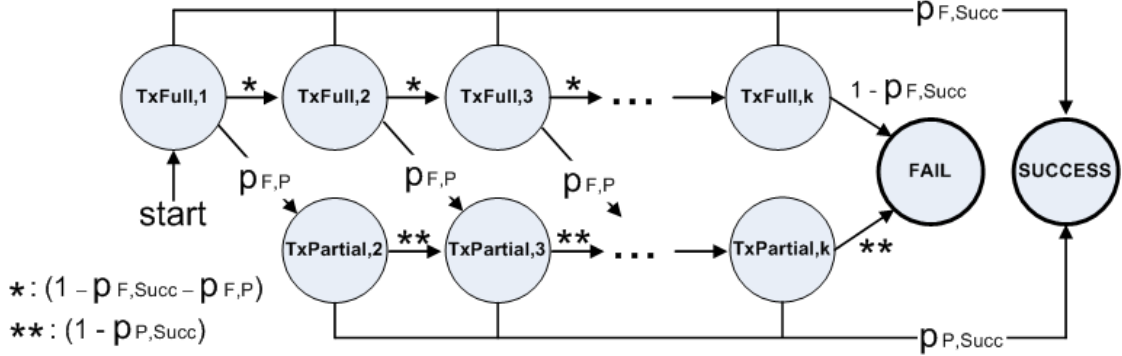
Figure 7.16: The Markov chain model for the REPE-ARQ scheme.

ability $p_{F,P}$.

Summarizing:

$$c_F = l_{D,F} + p_{F,Succ} \cdot l_{ACK} + p_{F,P} \cdot l_{NACK}$$

Similarly:

$$c_P = l_{D,P} + p_{P,Succ} \cdot l_{ACK}$$

where $l_{D,P}$ is the average number of bits for a partial data frame.[5] The average number of transmitted bits is computed by solving $\phi = \mathbf{P} \cdot \phi + \mathbf{c}$ for ϕ and reading off from ϕ the component corresponding to state `TxFull,1`, which amounts to solving a simple linear equation system (see Appendix B). Below we show results for probabilities and average frame sizes we obtained through measurements.

A similar model has been developed for the Send-And-Wait-ARQ protocol with k allowable trials. The major difference to the REPE model is the missing `TxPartial,i` states.

Experimental Setup

We implemented the REPE-ARQ scheme on the Tmote Sky platform in TinyOS 2.1 and integrated it with an existing CSMA MAC protocol. The MAC requires an idle channel before transmission and performs a maximum of three retransmissions in case an acknowledgement is not received correctly. We made measurements in the TWIST and MoteLab testbeds using the following setup: from the set of available nodes we selected one node to make periodic unicast transmissions (with a constant PPDU size of 133 byte) to a subset of about 25 other nodes in a round-robin fashion. In order to evaluate both, the REPE- and the standard ARQ, under identical conditions, the REPE algorithm on the transmitter side was slightly modified: after the reception of a correct NACK the sender never reduced the frame size, but instead it inserted in the payload an additional, intermediate CRC at an offset where the reduced frame would have ended.

---

[5]The model contains an approximation: in the calculation of c_F and c_P it is implicitly assumed that ACKs and NACKs are always received correctly.

146

Figure 7.17: The ratio of bytes transmitted with the REPE-ARQ vs. standard ARQ scheme (REPE Gain) during a 4-hour measurement / according to the analytical model. Each bar represents an 802.15.4 link and the bars are ordered by increasing PRR, ranging from about 0.1 to 0.99 (with a similar distribution as in Fig. 7.8).

In case of a bit error the receiver could thus determine whether the reduced-, the full-sized or both frames were affected[6], although in practice only a single full-sized frame was transmitted. With this approach both variants were running virtually *at the same time*. The sender/receiver could identify whether the transmission based on the REPE- or standard ARQ (or both) were (un)successful and count the number of transmitted bytes separately.

For example, assume that the first frame is completely destroyed; the second frame, i.e. first retransmission, is corrupted, the REPE-ARQ makes a correct estimate and the NACK (15 byte PPDU) is correctly received; the third frame has a single bit error, which is outside the region covered by the intermediate CRC and the corresponding ACK is received correctly; the fourth frame and its ACK are received correctly. Then the number of transmitted bytes accounted for the REPE-ARQ scheme is $133 + (133 + 15) + (40 + 11) + 0 = 332$ byte (assuming the retransmitted frame had a PPDU size of 40 byte and an ACK is 11 byte), for the standard ARQ it is $133 + 133 + 133 + (133 + 11) = 543$ byte.

Analytical and Experimental Results

Fig. 7.17 displays the experimental results from the TWIST and MoteLab testbeds. It shows a comparison of the REPE- with the standard ARQ scheme with respect to the total number of transmitted bytes per link during a 4-hour measurement (red bars). The results take all exchanged packets into consideration, including packets for which the REPE algorithm could not make a decision or packets that were inaccessible by software. From the traces we derived the probabilities $p_{F,Succ}$, $p_{P,Succ}$ and $p_{F,P}$, as well as the average frame sizes l_{NACK} and $l_{D,P}$ required by our analytical model (Sect. 7.3.3). The analytical results are shown as black bars in the same Fig. 7.17.

The experimental results differ for the two testbeds, possibly due to the diverse interference environments: in TWIST the gain is to 0.7 to 5.2 % (on average 2.1 %), while in MoteLab it is -0.5 to 14.7 % (on average 6.0 %). The highest gain is typically achieved on intermediate

---

[6]Only in the first case the receiver would transmit an ACK frame and it would include a special flag to distinguish it from the ACK for a full-sized frame.

links with a PRR of around 0.6. On very good links the potential performance gain is low, because only a small fraction of the frames is corrupted and can be recovered in the first place. The analytical model matches these results quite well: the average difference is 0.0059 and the maximum is 0.029. This suggests that the model is suitable and that with a simple software extension on a typical mote platform the REPE scheme can indeed achieve a considerable performance gain.

7.3.4 Discussion

We have shown that by monitoring RSSI with high sampling rate during packet reception the receiver can often not only identify RF interference as the cause for packet corruption, but also estimate the bit error positions correctly. Our algorithm simply correlates the instantaneous signal strength with the arrival time of the individual symbols and marks all bits received while the signal level was above a certain threshold as incorrect. It is particularly effective for large packets and when the interfering signal has short airtime. Through an analytical and experimental evaluation we showed that the approach can achieve high success rates and that an ARQ scheme generally benefits when coupled with our algorithm. We expect an additional performance gain if the radio hardware could provide more fine-grained signal strength information (shorter RSSI time windows or per-symbol LQI/correlation values).

While many of our results were obtained through testbed measurements, we have also conducted experiments with WBANs in three different urban scenarios. The corresponding results match the average results from the testbed studies quite well. Therefore, based on the performance evaluation (Section. 7.3.3) we estimate practical savings of 5-10 % of energy for intermediate (PRR below 0.9) WBAN links. While this is a moderate increase only, the overhead in terms of protocol complexity for the devices is rather small: in line with our central design goal specified in Section 3.2.6 we push the resource-expensive tasks in terms of computation and memory (RSSI profiling) to the WBAN coordinator. The resulting energy savings by the devices may then be leveraged, for example, to improve communication reliability by increasing the maximum number of retransmissions.

7.4 Summary

In the beginning of this chapter we confirmed the main conclusion drawn from the analysis in Section 6.5: the overall impact of RF interference on intra-WBAN communication reliability is expected to be rather small (about 2-3% RF interference-related packet losses). However, since interference is bursty (correlated in time) it may still be necessary to apply interference mitigation techniques in applications with tight latency bounds. To this end we have analyzed two techniques: channel hopping and packet combining. We have shown that channel hopping not only levels out the differences in the reliability among channels, but it also reduces the number of packet error bursts due to RF interference. However, even when blacklisting notoriously busy channels, channel hopping can hardly beat a very simple strategy: static selection of IEEE 802.15.4 channel 26. Because it is spaced far from the practically used WLAN center frequencies, channel 26 generally shows very little RF interference-related (burst) errors. There-

fore selecting channel 26 is the most recommended RF interference mitigation technique — at least as long as inter-WBAN interference is not an issue. [7] In the second part of this chapter we showed that by monitoring RSSI during frame reception we can estimate bit-error positions within packets that have been corrupted by RF interference. By exploiting this fact in an extended retransmission approach we achieved practical energy savings of 5-10 %. These savings can, for example, be directly translated in reliability improvements through an increase in the number of maximum retransmissions.

---

[7]Note however, that although channel hopping may not provide substantial performance robustness against RF interference (in comparison to IEEE 802.15.4 channel 26), it may still provide significant robustness to fading.

Chapter 8

Conclusions

The focus of this thesis has been on RF communication among human-mounted sensors. The collection of such sensors is envisioned to enable a range of applications from mobile health monitoring over personal fitness to assisted living. Within the communication network sensor data is typically transferred from the sensor devices to a more capable device such as the user's cell phone. Since this communication is performed over a wireless channel it is inherently unreliable: the wireless channel is an unstable medium for communication and it becomes even more dynamic when the communicating entities are mobile. At the same time many envisioned WBAN applications require reliable data transmission.

This thesis has been addressing the two main sources of intra-WBAN communication unreliability on the level of the wireless channel: fading and RF interference. While both effects have been studied well in the context of other wireless networks, their impact on intra-WBAN communication is not well understood. Therefore, one part of this thesis has been dedicated to establishing a better understanding of how the specific WBAN operating conditions, such as node placement on the human body surface and user mobility, impact intra-WBAN communication. This analysis relates to the intended communication among WBAN devices as well as the impact of interfering signals from other co-located devices. In the remaining second part of this thesis we have been analyzing mechanisms that have the potential to improve communication reliability without sacrificing a significant amount of additional energy (since WBANs are battery-powered and their energy resources are scarce). To this end, we have investigated how deferring communication until fading conditions are "good" can increase intra-WBAN communication reliability; and we have studied the potential of two interference mitigation techniques: channel hopping and packet combining. In the following we summarize the main conclusions that have been drawn within the scope of this thesis.

- While a person is walking, the signal strength of an on-body communication link fluctuates strongly, sometimes more than $20\,\text{dB}$ during a single stride period. The fluctuation is caused, among other things, by changes in relative node distance and body shadowing, sometimes in conjunction with an alternation between LOS and NLOS communication.

- Since some of these effects occur periodically, as long as the person continues walking signal strength peaks follow a *regular* pattern: past observations allow to estimate future

151

peaks relative to the user's gait cycle. Outdoors, on average 7.5 dB can be gained, if on-body communication is scheduled at signal strength peaks rather than at arbitrary points in time.

- The previously described effect can be exploited within the WBAN protocol stack. By extracting a set of RSSI samples from existing application traffic (15 samples are typically sufficient) and monitoring the person's gait cycle with the help of an accelerometer, a WBAN software component can effectively predict signal strength peaks as long as the person continues walking.

- When packets are transmitted at predicted signal strength peaks, *outdoors* unreliable (intermediate) links can often be turned into reliable links with PRR values well above 90 %. However, *indoors* an increase in reliability is only noticeable at -10 dBm transmission power, because at 0 dBm many links can hardly be improved. The gains in reliability come at the cost of an increase in latency by on average half a stride period, but the additional energy spent by devices is negligible.

- There is strong evidence that WLAN is the major source of urban 2.4 GHz interference. We expect that interfering signals with a power ≥ -94 dBm — which is a representative sensitivity threshold for the class of radios we consider — are present for on average about 5 % of the time, although there is a large variance over different environments.

- In comparison to the effects of fading we expect that RF interference will cause relatively little packet loss: we estimate that RF interference is responsible for only about 3% of WBAN packet loss, while outdoors fading effects (a signal below sensitivity threshold) cause around 17% to 25% of losses. This is due to the fact that an SINR of 5 dB is often sufficient to successfully decode a IEEE 802.15.4 packet and (a) the vast majority of RF noise samples are below -94 dBm while (b) the majority of time the signal strength of an intra-WBAN communication is outside the critical "corridor" between -94 dBm and about -90 dBm, where RF interference is harmful (we assume that a packet below the sensitivity threshold of -94 dBm cannot be harmed by RF noise, because it is too weak to be received in the first place).

- Urban RF noise is *bursty*, i.e. correlated in time: in many urban scenarios we found that consecutive packet losses due to interference were about ten times more probable than in a synthetic scenario where the same amount of RF noise is randomly (uniformly) distributed over time. Therefore even if average spectrum utilization is low, it may be necessary to apply RF interference mitigation technique if applications have low latency bounds.

- Pseudo-random channel hopping techniques — like the TSCH approach specified in the latest revision of the IEEE 802.15.4 standard — mitigate the effects of RF interference by spreading the risk of encountering bad conditions on one particular channel over the available frequency region. TSCH can thus level out differences between channels, which is important, because the PLR among channels can differ by more than a factor three.

- However, currently the best strategy is to simply select IEEE 802.15.4 channel 26: not only is the amount of packet loss significantly lower than on other channels or when applying channel hopping (even when blacklisting notoriously busy channels), but also the number and duration of packet error bursts due to interference can hardly be outperformed by other variants. In contrast to the U.S., in Europe it is in fact possible to select a WLAN channel that overlaps with IEEE 802.15.4 channel 26. However, our measurements show that this is rarely happening in practice (possibly because interoperability with U.S. WLAN client devices is favored).

- In practice even on channel 26 occasional packet corruption through RF interference cannot be avoided. Whenever this happens the MAC protocol normally retransmits the entire packet, even if only a few bits have been corrupted. By monitoring signal strength during packet reception a software component on the receiver side can often detect the bit error positions. Our experimental evaluation indicates practical savings of 5-10 % of energy for intermediate (PRR below 0.9) WBAN links if only corrupted bits retransmitted.

In summary, this thesis has established a better understanding of how the characteristic fading and RF interference conditions affect intra-WBAN communication reliability. It has also demonstrated how coupling WBAN communication with the movements of the user can help predicting channel conditions; and it has examined the suitability of interference mitigation techniques in the context of WBANs based on a large dataset of realistic, urban interference traces.

There are, however, still several aspects to investigate. First, we have not considered that fading effects are frequency-dependent. Although channel hopping may not provide substantial performance robustness against RF interference in comparison to the static selection of IEEE 802.15.4 channel 26, we do expect that frequency diversity techniques can significantly increase robustness against multipath fading. A thorough investigation in the context of WBANs is, however, still missing. Second, spatial diversity schemes are another means of improving robustness against fading: for example, the 802.15.6 standard [81] allows to introduce relayer nodes to form a "two-hop extended star [topology]". This topic also warrants further investigation from the scientific community. Harmonizing the requirement of reliable communication with minimizing emitted interference via adaptive transmission power schemes represents another underexplored WBAN topic. Furthermore, in addition to providing more realistic simulation models and tools alternative WBAN communication technologies such as UWB or inductive coupling require a closer examination. Finally, WBAN propagation and interference characteristics in the 868 MHz or 60 GHz band justify further exploration.

Appendix A

Truth table for Equation 5.11 and Equation 5.12

Table 8.1 and Table 8.2 show the truth tables for Equation 5.11 and Equation 5.12. Since a pivot operation preserves the matrix' property of being totally unimodular we can ignore any row, where the assignment results in a value other than $+1, -1, 0$ (in column "$x-r$" or "$x+r$"), respectively, because these assignments will never be executed in your system.

x	x_p	x_n	r	r_p	r_n	x-r	$(x_p \wedge \neg r_p) \vee (\neg x_n \wedge r_n)$	$(x_n \wedge \neg r_n) \vee (\neg x_p \wedge r_p)$
0	0	0	0	0	0	0	0	0
0	0	0	1	1	0	-1	0	1
0	0	0	-1	0	1	1	1	0
1	1	0	0	0	0	1	1	0
1	1	0	1	1	0	0	0	0
1	1	0	-1	0	1	2	X	X
-1	0	1	0	0	0	-1	0	1
-1	0	1	1	1	0	-2	X	X
-1	0	1	-1	0	1	0	0	0

Table 8.1: Truth table for Equation 5.11 and Equation 5.12 for the case $c_p = 1$.

x	x_p	x_n	r	r_p	r_n	x+r	$(x_p \vee r_p) \wedge (\neg x_n \wedge \neg r_n)$	$(x_n \vee r_n) \wedge (\neg x_p \wedge \neg r_p)$
0	0	0	0	0	0	0	0	0
0	0	0	1	1	0	1	1	0
0	0	0	-1	0	1	-1	0	1
1	1	0	0	0	0	1	1	0
1	1	0	1	1	0	2	X	X
1	1	0	-1	0	1	0	0	0
-1	0	1	0	0	0	-1	0	1
-1	0	1	1	1	0	0	0	0
-1	0	1	-1	0	1	-2	X	X

Table 8.2: Truth table for Equation 5.11 and Equation 5.12 for the case $c_n = 1$.

Appendix B

Potentials of Markov chains

In the following we paraphrase definitions and theorems from the framework of potential theory for Markov chains [269, Sec. 4.2], which we have used in 7.3.3. Be $(X_n)_{n \geq 0}$ a time-homogeneous Markov chain with discrete (i.e. finite or countably infinite) state space $\mathcal{S}$ and state-transition matrix $\mathbf{P}$. The state space is partitioned into *inner states* D and *boundary states* or *final states* ∂D so that $\mathcal{S} = D \cup \partial D$. Suppose that $\mathbf{c} = (c_i)_{i \in D}$ and $\mathbf{f} = (f_i)_{i \in \partial D}$ are non-negative vectors representing the costs c_i when the chain is in the inner state $i \in D$ and the costs f_i when the chain is in the boundary state $i \in \partial D$. Let the random variable T be the hitting time for the boundary: $T = \inf \{n \geq 0 : X_n \in \partial D\}$. Set

$$\phi_i = \mathrm{E}_i \left[\sum_{n < T} c(X_n) + f(X_T) \cdot \mathbf{1}_{\{T < \infty\}} \right]$$

Then ϕ_i is the expected total costs when the chain starts in state $X_0 = i$ and operates in the inner states D, each time incurring a cost c_i, until it reaches a final state in ∂D, incurring a final cost corresponding to the final state. The final costs are incurred only when the hitting time T is finite. Then the following holds [269, Theorem 4.2.3]:

- The potential $\phi = (\phi_i)_{i \in \mathcal{S}}$ satisfies:

$$\begin{cases} \phi = \mathbf{P} \cdot \phi + \mathbf{c} & : \quad \text{in } D \\ \phi = \mathbf{f} & : \quad \text{in } \partial D \end{cases} \tag{8.1}$$

- If $\mathrm{Pr}_i[T < \infty] = 1$ (i.e. the probability to hit the final states when the starting state is $X_0 = i$) for all i then Equation 8.1 has at most one bounded solution.

In other words, we are looking for a solution of the system of linear equations given in 8.1.

Appendix C

List of the Author's Scientific Publications

The topics in this thesis have partially been covered by the author in the following publications:

Journal Article

- **J.-H. Hauer**, "Leveraging Human Mobility for Communication in Body Area Networks" *ACM Transactions on Sensor Networks (TOSN)*, vol. 10, no. 3, pp. XX–XX, ISSN: X, *ACCEPTED FOR PUBLICATION (2014)*

Bookchapter

- **J.-H. Hauer**, S. Tennina, R. Severino, A. Koubâa, Chapter 3: "Open-source implementations of the IEEE 802.15.4 and ZigBee protocols", in *IEEE 802.15.4 and ZigBee as Enabling Technologies for Low-Power Wireless Systems with Quality-of-Service Constraints*, Editor: Mario Alves, pp. 27–43, ISBN: 978-3-642-37368-8, SpringerBriefs in Electrical and Computer Engineering, Springer, 2013

Conference Proceedings

- **J.-H. Hauer** and D. Willkomm, "An Empirical Study of Urban 2.4 GHz RF Noise from the Perspective of a Body Sensor Network", in *Proc. of the 9th International Conference on Wearable and Implantable Body Sensor Networks (BSN'12)*, pp. 1–8, ISBN: 978-1-4673-1393-3, London, United Kingdom, May 2012.

- **J.-H. Hauer**, R. Daidone, R. Severino, J. Büsch, M. Tiloca and S. Tennina, "Poster Abstract: An Open-Source IEEE 802.15.4 MAC Implementation for TinyOS 2.1" in *Proc. of the 8th European Conference on Wireless Sensor Networks (EWSN'11)*, pp. 1–2, ISBN: 978-3-642-19185-5, Bonn, Germany, February 2011.

- K.S. Prabh and **J.-H. Hauer**, "Opportunistic Packet Scheduling in Body Area Networks" in *Proc. of the 8th European Conference on Wireless Sensor Networks (EWSN'11)*, pp. 114–129, ISBN: 978-3-642-19185-5, Bonn, Germany, February 2011.

- **J.-H. Hauer**, A. Willig and A. Wolisz, "Mitigating the Effects of RF Interference through RSSI-Based Error Recovery" in *Proc. of the 7th European Conference on Wireless Sensor Networks (EWSN'10)*, pp. 224–239, ISBN: 978-3-642-11916-3, Coimbra, Portugal, February 2010.

- **J.-H. Hauer**, V. Handziski and A. Wolisz, "Experimental Study of the Impact of WLAN Interference on IEEE 802.15.4 Body Area Networks" in *Proc. of the 6th European Conference on Wireless Sensor Networks (EWSN'09)*, pp. 17–32, ISBN: 978-3-642-00223-6, Cork, Ireland, February 2009.

- V. Handziski, Joseph Polastre, **J.-H. Hauer**, C. Sharp, A. Wolisz and David Culler, "Flexible Hardware Abstraction for Wireless Sensor Networks" in *Proc. of the 2nd European Workshop on Wireless Sensor Networks (EWSN'05)*, pp. 145 - 157, ISBN: 0-7803-8801-1, Istanbul, Turkey, February 2005.

Technical Reports

- P. Levis, D. Gay, V. Handziski, **J.-H. Hauer**, B. Greenstein, M. Turon, J. Hui, K. Klues, C. Sharp, R. Szewczyk, J. Polastre, P. Buonadonna, L. Nachman, G. Tolle, D. Culler and A. Wolisz, "T2: A Second Generation OS For Embedded Sensor Networks", in *TKN Technical Report Series, TKN-05-007, Telecommunication Networks Group, Technische Universität Berlin*, November 2005.

- **J.-H. Hauer**, "TKN15.4: An IEEE 802.15.4 MAC Implementation for TinyOS 2", in *TKN Technical Report Series, TKN-08-003, Telecommunication Networks Group, Technische Universität Berlin*, March 2009.

The following list represents publications of the author published in addition to those referenced in this thesis:

Journal Article

- A. Willig and N. Karowski and **J.-H. Hauer**, "Passive discovery of IEEE 802.15.4-based body sensor networks", in *Ad Hoc Networks*, pp. 742–754, vol. 8, no.7, ISSN: 1570-8705, September 2010.

Conference Proceedings

- M. Chwalisz, **J.-H. Hauer**, J. Büsch, "Demo Abstract: An Infrastructure for Automated Body Area Network Experimentation" in *Proc. of the 9th European Conference on Wireless Sensor Networks (EWSN'13)*, pp. 1–2, ISBN: 978-3-642-36671-0, Ghent, Belgium, February 2013. BEST DEMO AWARD.

- D. Finn, J. Tallon, L. DaSilva, P. Wesemael, S. Polin, W. Liu, S. Bouckaert, J. Vanhie-Van Gerwen, N. Michailow, **J.-H. Hauer**, D. Willkomm and C. Heller, "Experimental

Assessment of Tradeoffs among Spectrum Sensing Platforms", in *Proc. of the ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation and Characterization (WiNTECH'11)*, pp. 67–74, ISBN: 978-1-4503-0867-0, Las Vegas, USA, September 2011.

- M. Zúñiga, I. Irzynska, **J.-H. Hauer**, T. Voigt, C. Boano and K. Römer, "Link Quality Ranking: Getting the Best out of Unreliable Links" in *Proc. of the 7th Int. Conference on Distributed Computing in Sensor Systems (DCOSS'11)*, pp. 1–8, ISBN: 978-1-4577-0512-0, Barcelona, Spain, June 2011.

- A. Sanchez, I. Moerman, S. Bouckaert, D. Willkomm, **J.-H. Hauer**, N. Michailow, G. Fettweis, L. DaSilva, J. Tallon and S. Polin, "Testbed Federation: An Approach for Experimentation-driven Research in Cognitive Radios and Cognitive Networking", in *Proc. of the Future Network and Mobile Summit*, pp. 1–9, ISBN: 978-1-4577-0928-9, Warsaw, Poland, June 2011.

- D. Kipnis, A. Willig, **J.-H. Hauer** and N. Karowski, "The ANGEL IEEE 802.15.4 Enhancement Layer: Coupling Priority Queueing and Service Differentiation", in *Proc. of the European Wireless Conference (EW'08)*, pp. 1–7, ISBN: 978-3-8007-3102-2, Prague, Czech Republic, June 2008.

- **J.-H. Hauer**, V. Handziski, A. Köpke, A. Willig and A. Wolisz, "A Component Framework for Content-based Publish/Subscribe in Sensor Networks", in *Proc. of 5th European Conference on Wireless Sensor Networks (EWSN'08)*, pp. 369–385, ISBN: 978-3-540-77689-5, Bologna, Italy, January 2008.

- V. Handziski, J. Polastre, **J.-H. Hauer** and C. Sharp, "Flexible hardware abstraction of the TI MSP430 microcontroller in TinyOS" in *Proc. of the 2nd International Conference on Embedded Networked Sensor Systems (SenSys'04)*, pp. 277–278, ISBN:1-58113-879-2, Baltimore, USA, November 2004.

- A. Köpke, V. Handziski, **J.-H. Hauer** and H. Karl, "Structuring the Information Flow in Component-Based Protocol Implementations for Wireless Sensor Nodes", in *Proc. of Work-in-Progress Session of the 1st European Workshop on Wireless Sensor Networks (EWSN'04)*, pp. 41–45 Berlin, Germany, January 2004.

Technical Reports

- N. Karowski, A. Willig and **J.-H. Hauer**, "Passive discovery schemes for opportunistic message relaying schemes based on IEEE 802.15.4", in *TKN Technical Report Series, TKN-08-008, Telecommunication Networks Group, Technische Universität Berlin*, August 2008

- A. Köpke and **J.-H. Hauer**, "IEEE 802.15.4 Symbol Rate Timer for TelosB", in *TKN Technical Report Series, TKN-08-006, Telecommunication Networks Group, Technische Universität Berlin*, May 2008

Bibliography

[1] Benoît Latré, Bart Braem, Ingrid Moerman, Chris Blondia, and Piet Demeester, "A survey on wireless body area networks," *Wireless Networks*, pp. 1–18, 2010.

[2] T. Penzel, B. Kemp, G. Klosch, A. Schlogl, J. Hasan, A. Varri, and I. Korhonen, "Acquisition of biomedical signals databases," *Engineering in Medicine and Biology Magazine, IEEE*, vol. 20, no. 3, pp. 25 –32, may/jun 2001.

[3] Mathias Benedek and Christian Kaernbach, "Decomposition of skin conductance data by means of nonnegative deconvolution," *Psychophysiology*, vol. 47, no. 4, pp. 647–658, 2010.

[4] M.D. Breton, D.P. Shields, and B.P. Kovatchev, "Optimum subcutaneous glucose sampling and fourier analysis of continuous glucose monitors," *Journal of diabetes science and technology (Online)*, vol. 2, no. 3, pp. 495, 2008.

[5] C.V.C. Bouten, K.T.M. Koekkoek, M. Verduin, R. Kodde, and J.D. Janssen, "A triaxial accelerometer and portable data processing unit for the assessment of daily physical activity," *Biomedical Engineering, IEEE Transactions on*, vol. 44, no. 3, pp. 136 –147, march 1997.

[6] Zexi Liu and Chang-Hee Won, "Knee and waist attached gyroscopes for personal navigation: Comparison of knee, waist and foot attached inertial sensors," in *Position Location and Navigation Symposium (PLANS), 2010 IEEE/ION*, may 2010, pp. 375 –381.

[7] D. Hasenfratz, O. Saukh, S. Sturzenegger, and L. Thiele, "Participatory air pollution monitoring using smartphones," *Mobile Sensing*, 2012.

[8] Shimmer Research, "Shimmer2 capabilities overview," http://www.shimmer-research. com/download/documentation, 2010.

[9] OECD, "Organisation for Economic Co-operation and Development: Projecting OECD health and long-term care expenditures: What are the main drivers?," Economics Department Working Paper 477, 2006.

[10] Ratko Magjarevic, "Home care technologies for ambient assisted living," in *11th Mediterranean Conference on Medical and Biomedical Engineering and Computing 2007*, Ratko

161

Magjarevic, Tomaz Jarm, Peter Kramar, and Anze Zupanic, Eds., vol. 16 of *IFMBE Proceedings*, pp. 397–400. Springer Berlin Heidelberg, 2007.

[11] J. Wilson, V. Bhargava, A. Redfern, and P. Wright, "A wireless sensor network and incident command interface for urban firefighting," in *Mobile and Ubiquitous Systems: Networking Services, 2007. MobiQuitous 2007. Fourth Annual International Conference on*, Aug. 2007, pp. 1 –7.

[12] Reed W Hoyt, Jaques Reifman, Trinka S Coster, and Mark J Buller, "Combat medical informatics: present and future," *Proceedings of the AMIA Symposium*, , no. 3, pp. 335–339, 2002.

[13] S. Drude, "Requirements and application scenarios for body area networks," in *Mobile and Wireless Communications Summit, 2007. 16th IST*, July 2007, pp. 1 –5.

[14] Peter Barrie, Andreas Komninos, and Oleksii Mandrychenko, "A pervasive gesture-driven augmented reality prototype using wireless sensor body area networks," in *Proceedings of the 6th International Conference on Mobile Technology, Application and Systems*, New York, NY, USA, 2009, Mobility '09, pp. 61:1–61:4, ACM.

[15] Guang-Zhong Yang, Ed., *Body Sensor Networks*, Springer Verlag, 2006.

[16] Jamil Y . Khan and Mehmet R. Yuce, *Wireless Body Area Networks*, Pan Stanford Publishing, Dec 2011.

[17] Prabal Dutta and David Culler, "Mobility changes everything in low-power wireless sensornets," in *Proceedings of the 12th conference on Hot topics in operating systems*, Berkeley, CA, USA, 2009, HotOS'09, pp. 7–7, USENIX Association.

[18] "IEEE standard for information technology- telecommunications and information exchange between systems- local and metropolitan area networks- specific requirements part 15.4: Wireless medium access control (MAC) and physical layer (PHY) specifications for low-rate wireless personal area networks (WPANs)," *IEEE Std 802.15.4-2006 (Revision of IEEE Std 802.15.4-2003)*, pp. 1–305, 2006.

[19] 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*, New York, NY, USA, 2010, SenSys '10, pp. 309–322, ACM.

[20] Branko Kerkez, "Adaptive time synchronization and frequency channel hopping for wireless sensor networks," M.S. thesis, Electrical Engineering and Computer Sciences University of California at Berkeley , 2012.

[21] Razvan Musaloiu and Andreas Terzis, "Minimising the effect of WiFi interference in 802.15.4 wireless sensor networks," *Int. J. Sen. Netw.*, vol. 3, no. 1, pp. 43–54, Dec. 2008.

[22] Peng Du and G. Roussos, "Adaptive time slotted channel hopping for wireless sensor networks," in *Computer Science and Electronic Engineering Conference (CEEC), 2012 4th*, 2012, pp. 29–34.

[23] Jorge Ortiz and David Culler, "Multichannel reliability assessment in real world wsns," in *Proc. of IPSN'10*, Stockholm, Sweden.

[24] Thomas Watteyne, Ankur Mehta, and Kris Pister, "Reliability through frequency diversity: why channel hopping makes sense," in *Proc. of PE-WASUN'09*, Tenerife, Spain, 2009.

[25] V.Handziski, A.Köpke, A.Willig, and A.Wolisz, "Twist: A scalable and reconfigurable testbed for wireless indoor experiments with sensor network," in *Proc. of the 2nd Intl. Workshop on Multi-hop Ad Hoc Networks: from Theory to Reality, (RealMAN 2006)*, Florence, Italy, May 2006.

[26] Geoffrey Werner-Allen, Patrick Swieskowski, and Matt Welsh, "Motelab: a wireless sensor network testbed," in *IPSN '05: Proceedings of the 4th international symposium on Information processing in sensor networks*, Piscataway, NJ, USA, 2005, p. 68, IEEE Press.

[27] iMinds, Ghent, Belgium, "w-ilab.t testbed," http://www.crew-project.eu/portal/wilabdoc.

[28] Manjunath Doddavenkatappa, MunChoon Chan, and A.L. Ananda, "Indriya: A low-cost, 3d wireless sensor network testbed," in *Testbeds and Research Infrastructure. Development of Networks and Communities*, Thanasis Korakis, Hongbin Li, Phuoc Tran-Gia, and Hong-Shik Park, Eds., vol. 90 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pp. 302–316. Springer Berlin Heidelberg, 2012.

[29] Kannan Srinivasan, Prabal Dutta, Arsalan Tavakoli, and Philip Levis, "An empirical study of low-power wireless," *ACM Trans. Sen. Netw.*, vol. 6, no. 2, pp. 16:1–16:49, Mar. 2010.

[30] P. Misra, N. Ahmed, and Sanjay Jha, "An empirical study of asymmetry in low-power wireless links," *Communications Magazine, IEEE*, vol. 50, no. 7, pp. 137–146, 2012.

[31] Marco Zuniga and Bhaskar Krishnamachari, "Analyzing the transitional region in low power wireless links," in *In First IEEE International Conference on Sensor and Ad hoc Communications and Networks (SECON)*, 2004, pp. 517–526.

[32] Omprakash Gnawali, Rodrigo Fonseca, Kyle Jamieson, David Moss, and Philip Levis, "Collection tree protocol," in *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*, New York, NY, USA, 2009, SenSys '09, pp. 1–14, ACM.

[33] Daniele Puccinelli, Silvia Giordano, Marco Zuniga, and Pedro José Marrón, "Broadcast-free collection protocol," in *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems*, New York, NY, USA, 2012, SenSys '12, pp. 29–42, ACM.

[34] Olaf Landsiedel, Euhanna Ghadimi, Simon Duquennoy, and Mikael Johansson, "Low power, low delay: opportunistic routing meets duty cycling," in *Proceedings of the 11th international conference on Information Processing in Sensor Networks*, New York, NY, USA, 2012, IPSN '12, pp. 185–196, ACM.

[35] Jan-Hinrich Hauer, Roberta Daidone, Ricardo Severino, Jasper Büsch, Marco Tiloca, and Stefano Tennina, "Poster abstract: An open-source ieee 802.15.4 MAC implementation for TinyOS 2.1," in *Proc. of 8th European Conference on Wireless Sensor Networks (EWSN)*, Bonn, Germany, Feb. 2011.

[36] Jan-Hinrich Hauer, "TKN15.4: An IEEE 802.15.4 MAC Implementation for TinyOS 2," TKN Technical Report Series TKN-08-003, Telecommunication Networks Group, Technical University Berlin, 2009.

[37] S. Tennina, A. Koubâa, R. Daidone, M. Alves, R. Severino, P. Jurcik, M. Tiloca, J.-H. Hauer, G. Dini, M. Bouroche, and Tovar E., *Engineering IEEE 802.15.4/ZigBee-based WSN/COBS systems*, Springer, September 2012.

[38] Jan-Hinrich Hauer, "Leveraging human mobility for communication in body area networks," *ACM Trans. Sen. Netw.*, vol. X, no. X, pp. X, X To Appear 2014.

[39] K. Shashi Prabh and Jan-Hinrich Hauer, "Opportunistic packet scheduling in body area networks," in *Proc. of 8th European Conference on Wireless Sensor Networks (EWSN)ISBN 978-3-642-19185-5*, Bonn, Germany, Feb. 2011, pp. 114–129.

[40] J. H. Hauer, V.Handziski, and A.Wolisz, "Experimental study of the impact of WLAN interference on IEEE 802.15.4 body area networks," in *Proc. of 6th European Conference on Wireless Sensor Networks (EWSN)*, Cork, Ireland, Feb. 2009.

[41] Jan-Hinrich Hauer and Daniel Willkomm, "An empirical study of urban 2.4 GHz RF noise from the perspective of a body sensor network," in *Ninth International Conference on Wearable and Implantable Body Sensor Networks (BSN), 2012*, May 2012, vol. 2012.

[42] Jan-Hinrich Hauer, Andreas Willig, and Adam Wolisz, "Mitigating the effects of RF interference through RSSI-based error recovery," in *Proc. of 7th European Conference on Wireless Sensor Networks (EWSN)ISBN 978-3-642-11916-3*, Coimbra, Portugal, Feb. 2010.

[43] M. D. Buist, E. Jarmolowski, P. R. Burton, S. A. Bernard, B. P. Waxman, and J. Anderson, "Recognising clinical instability in hospital patients before cardiac arrest or unplanned admission to intensive care. A pilot study in a tertiary-care hospital.," *Med J Aust*, vol. 171, no. 1, pp. 22–25, July 1999.

[44] Dorothy W Curtis, Esteban J Pino, Jacob M Bailey, Eugene I Shih, Jason Waterman, Staal A Vinterbo, Thomas O Stair, John V Guttag, Robert A Greenes, and Lucila Ohno-Machado, "Smart-an integrated wireless system for monitoring unattended patients," *Journal of the American Medical Informatics Association*, vol. 15, no. 1, pp. 44–53, 2008.

[45] Jeonggil Ko, Jong Hyun Lim, Yin Chen, Rvăzvan Musvaloiu-E, Andreas Terzis, Gerald M. Masson, Tia Gao, Walt Destler, Leo Selavo, and Richard P. Dutton, "MEDiSN: Medical emergency detection in sensor networks," *ACM Trans. Embed. Comput. Syst.*, vol. 10, no. 1, pp. 11:1–11:29, Aug. 2010.

[46] Octav Chipara, Chenyang Lu, Thomas C. Bailey, and Gruia-Catalin Roman, "Reliable clinical monitoring using wireless sensor networks: experiences in a step-down hospital unit," in *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, New York, NY, USA, 2010, SenSys '10, pp. 155–168, ACM.

[47] Hai ying Zhou and Kun mean Hou, "Pervasive cardiac monitoring system for remote continuous heart care," in *Bioinformatics and Biomedical Engineering (iCBBE), 2010 4th International Conference on*, june 2010, pp. 1 –4.

[48] Eugene I. Shih, Ali H. Shoeb, and John V. Guttag, "Sensor selection for energy-efficient ambulatory medical monitoring," in *Proceedings of the 7th international conference on Mobile systems, applications, and services*, New York, NY, USA, 2009, MobiSys '09, pp. 347–358, ACM.

[49] Shyamal Patel, Konrad Lorincz, Richard Hughes, Nancy Huggins, John Growdon, David Standaert, Metin Akay, Jennifer Dy, Matt Welsh, and Paolo Bonato, "Monitoring motor fluctuations in patients with parkinson's disease using wearable sensors," *Trans. Info. Tech. Biomed.*, vol. 13, no. 6, pp. 864–873, Nov. 2009.

[50] L. Atallah, O. Aziz, G.Z. Yang, and B. Lo, "Detecting walking gait impairment with an ear-worn sensor," in *Wearable and Implantable Body Sensor Networks, 2009. BSN 2009. Sixth International Workshop on*. IEEE, 2009, pp. 175–180.

[51] A. Wood, J. Stankovic, G. Virone, L. Selavo, Zhimin He, Qiuhua Cao, Thao Doan, Yafeng Wu, Lei Fang, and R. Stoleru, "Context-aware wireless sensor networks for assisted living and residential monitoring," *Network, IEEE*, vol. 22, no. 4, pp. 26 –33, 2008.

[52] Francisco Fernández-Luque, Juan Zapata, Ramón Ruiz, and Emilio Iborra, "A wireless sensor network for assisted living at home of elderly people," in *Proceedings of the 3rd International Work-Conference on The Interplay Between Natural and Artificial Computation: Part II: Bioinspired Applications in Artificial and Natural Computation*, Berlin, Heidelberg, 2009, IWINAC '09, pp. 65–74, Springer-Verlag.

[53] S.M. George, Wei Zhou, H. Chenji, Myounggyu Won, Yong Oh Lee, A. Pazarloglou, R. Stoleru, and P. Barooah, "Distressnet: a wireless ad hoc and sensor network architecture for situation management in disaster response," *Communications Magazine, IEEE*, vol. 48, no. 3, pp. 128 –136, march 2010.

[54] Susan P. McGrath, Eliot Grigg, Suzanne Wendelken, George Blike, Michael De Rosa, Aaron Fiske, and Robert Gray, "ARTEMIS: A Vision for Remote Triage and Emergency Management Information Integration," 2003.

[55] Elizabeth Bales, Nima Nikzad, Nichole Quick, Celal Ziftci, Kevin Patrick, and William Griswold, "Citisense: Mobile air quality sensing for individuals and communities design and deployment of the citisense mobile air-quality system," in *Pervasive Computing Technologies for Healthcare (PervasiveHealth), 2012 6th International Conference on*, may 2012, pp. 155 –158.

[56] Prabal Dutta, Paul M. Aoki, Neil Kumar, Alan Mainwaring, Chris Myers, Wesley Willett, and Allison Woodruff, "Common sense: participatory urban sensing using a network of handheld air quality monitors," in *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*, New York, NY, USA, 2009, SenSys '09, pp. 349–350, ACM.

[57] S. Abboud and O. Barnea, "Errors due to sampling frequency of the electrocardiogram in spectral analysis of heart rate signals with low variability," in *Computers in Cardiology 1995*, sept. 1995, pp. 461 –463.

[58] R. Merletti and P. Di Torino, "Standards for reporting EMG data," *Journal of Electromyography and Kinesiology*, vol. 9, no. 1, pp. 3–4, 1999.

[59] Kenichi Takizawa, Masatoshi Homan, Yoshiki Takeoka, Takahiro Aoyagi, Hiroaki Hagiwara, and Tetsushi Ikegami, "IEEE P802.15 working group submission: Capsule endoscope using an implant WBAN," Tech. Rep. 15-08-0154-00-0006, IEEE P802.15, March 2008.

[60] N. Chevrollier and N. Golmie, "On the use of wireless network technologies in healthcare environments," in *Proceedings of the fifth IEEE workshop on Applications and Services in Wireless Networks (ASWN 2005)*, Paris, France, 2005, pp. 147–152.

[61] D. Keenan, J. Mastrototaro, and Steil G. Voskanyan, "Delays in minimally invasive continuous glucose monitoring devices: a review of current technology.," *Diabetes Sci Technology*, pp. 1207–14, Sep 2009.

[62] JW Kinn, JC Marek, MF O'Toole, SM Rowley, and Bufalino VJ, "Effectiveness of the electronic medical record in improving the management of hypertension.," *J Clin Hypertens (Greenwich)*, pp. 415–419, Nov 2002.

[63] St. Mary's Health Care, "Real-time vital sign monitoring ensures timely care for at-risk patients, leading to fewer unplanned intensive care unit admissions and out-of-intensive care unit cardiac arrests," http://www.innovations.ahrq.gov/content.aspx?id=1786, 2009.

[64] P. Pillatsch, E.M. Yeatman, and A.S. Holmes, "Piezoelectric rotational energy harvester for body sensors using an oscillating mass," in *Wearable and Implantable Body Sensor Networks (BSN), 2012 Ninth International Conference on*, 2012, pp. 6–10.

[65] "Texas instruments MSP430 16-bit microcontroller platform," http://www.ti.com/lsds/ti/microcontroller/16-bit_msp430/overview.page.

[66] Joseph Polastre, Robert Szewczyk, and David Culler, "Telos: Enabling ultra-low power wireless research," in *Proceedings of the Fourth International Conference on Information Processing in Sensor Networks: Special track on Platform Tools and Design Methods for Network Embedded Sensors (IPSN/SPOTS)*, 2005.

[67] Chris Otto, Aleksandar Milenković, Corey Sanders, and Emil Jovanov, "System architecture of a wireless body area sensor network for ubiquitous health monitoring," *J. Mob. Multimed.*, vol. 1, no. 4, pp. 307–326, Jan. 2005.

[68] "Texas instruments msp430f1611 datasheet," http://focus.ti.com/docs/prod/folders/print/msp430f1611.html.

[69] Texas Instruments, "CC2420 2.4 GHz IEEE 802.15.4 / ZigBee-ready RF Transceiver," http://www.ti.com/lit/gpn/cc2420, Apr. 2002.

[70] Freescale Semiconductor, "±1.5g - 6g three axis low-g micromachined accelerometer," www.freescale.com/files/sensors/doc/data_sheet/MMA7260QT.pdf, 2008.

[71] Benny Lo and Guang-Zhong Yang, "Architecture for body sensor networks," in *Perspectives in Pervasive Computing*, oct. 2005, pp. 23 –28.

[72] Chulsung Park, P.H. Chou, Ying Bai, R. Matthews, and A. Hibbs, "An ultra-wearable, wireless, low power ECG monitoring system," in *Biomedical Circuits and Systems Conference, 2006. BioCAS 2006. IEEE*, 29 2006-dec. 1 2006, pp. 241 –244.

[73] Emmanuel Munguia Tapia, Natalia Marmasse, Stephen S. Intille, and Kent Larson, "Mites: Wireless portable sensors for studying behavior," in *in Proceedings of Extended Abstracts Ubicomp 2004: Ubiquitous Computing*, 2004.

[74] R. Naima and J. Canny, "The berkeley tricorder: Ambulatory health monitoring," in *Wearable and Implantable Body Sensor Networks, 2009. BSN 2009. Sixth International Workshop on*, june 2009, pp. 53 –58.

[75] Prabal Dutta, Jay Taneja, Jaein Jeong, Xiaofan Jiang, and David Culler, "A building block approach to sensornet systems," in *Proceedings of the 6th ACM conference on Embedded network sensor systems*, New York, NY, USA, 2008, SenSys '08, pp. 267–280, ACM.

[76] Welch Allyn, Inc., URL: http://www.welchallyn.com.

[77] Corventis, Inc., URL: http://www.corventis.com.

[78] Polar Electro, URL: http://www.polar.com.

[79] Dynastream Innovations Inc., URL: http://www.thisisant.com.

[80] K. Wac, R. Bults, B. van Beijnum, I. Widya, V. Jones, D. Konstantas, M. Vollenbroek-Hutten, and H. Hermens, "Mobile patient monitoring: The mobihealth system," in *Engineering in Medicine and Biology Society, 2009. EMBC 2009. Annual International Conference of the IEEE*, sept. 2009, pp. 1238 –1241.

[81] "IEEE standard for local and metropolitan area networks - part 15.6: Wireless body area networks," *IEEE Std 802.15.6-2012*, pp. 1 –271, 29 2012.

[82] Javier Espina, Thomas Falc, and Oliver Mülhens, *Body Sensor Networks*, chapter Network Topologies, Communication Protocols, and Standards, pp. 145–182, Springer Verlag, 2006.

[83] Chris Otto, Aleksandar Milenković, Corey Sanders, and Emil Jovanov, "System architecture of a wireless body area sensor network for ubiquitous health monitoring," *J. Mob. Multimed.*, vol. 1, no. 4, pp. 307–326, Jan. 2005.

[84] Kristof Van Laerhoven, Benny PL Lo, Jason WP Ng, Surapa Thiemjarus, Rachel King, Simon Kwan, Hans-Werner Gellersen, Morris Sloman, Oliver Wells, Phil Needham, Nick Peters, Ara Darzi, Chris Toumazou, and Guang-Zhong Yang, "Medical healthcare monitoring with wearable and implantable sensors," in *International Workshop on Ubiquitous Computing for Pervasive Healthcare Applications (UbiHealth)*, September 2004.

[85] Victor Shnayder, Bor rong Chen, Konrad Lorincz, Thaddeus R. F. Fulford-Jones, and Matt Welsh, "Sensor networks for medical care," April 2005.

[86] A. Wood, G. Virone, T. Doan, Q. Cao, L. Selavo, Y. Wu, L. Fang, Z. He, S. Lin, and J. Stankovic, "Alarm-net: Wireless sensor networks for assisted-living and residential monitoring," Tech. Rep., 2006.

[87] Benoit Latre, Bart Braem, Ingrid Moerman, Chris Blondia, Elisabeth Reusens, Wout Joseph, and Piet Demeester, "A low-delay protocol for multihop wireless body area networks," in *Proceedings of the 2007 Fourth Annual International Conference on Mobile and Ubiquitous Systems: Networking&Services (MobiQuitous)*, Washington, DC, USA, 2007, MOBIQUITOUS '07, pp. 1–8, IEEE Computer Society.

[88] ZigBee Standards Organization, "ZigBee Specification," http://www.zigbee.org.

[89] David Culler, "IP is the future of ubiquitous sensor networks — keynote TinyOS USN applications workshop, seoul korea," Tech. Rep., July 2008.

[90] Zach Shelby and Carsten Bormann, *6LoWPAN: The Wireless Embedded Internet*, Wiley Publishing, 2010.

[91] Zach Shelby, Klaus Hartke, Carsten Bormann, and Brian Frank, "Constrained application protocol (CoAP)," Tech. Rep. draft-ietf-core-coap-18.txt, IETF Secretariat, Fremont, CA, USA, June 2013.

[92] R. Jacobsen, F.O. Hansen, J.K. Madsen, H. Karstoft, P.H. Mikkelsen, T.A. Skogberg, E.S. Rasmussen, C. Andersen, M. Alroe, and T.S. Toftegaard, "A modular platform for wireless body area network research and real-life experiments: The ASE-BAN testbed," *International Journal On Advances In Networks And Services*, vol. 4, no. 3-4, 2011.

[93] T. Rappaport, *Wireless Communications: Principles and Practice*, Prentice Hall PTR New Jersey, dec 2001.

[94] J.E. Turner, M.S. Jessup, and Kin-Fai Tong, "A novel technique enabling the realisation of 60 GHz body area networks," in *Wearable and Implantable Body Sensor Networks (BSN), 2012 Ninth International Conference on*, may 2012, pp. 58 –62.

[95] Eckhard Grass, Klaus Tittelbach-Helmrich, Chang-Soon Choi, Frank Winkler, Thilo Ohlemüller, and Rolf Kraemer, "Communication systems operating in the 60 GHz ISM band: Overview," *International Journal of Microwave and Wireless Technologies*, vol. 3, pp. 89–97, 4 2011.

[96] FCC: Medical Body Area Networks First Report and Order, URL: http://www.fcc.gov/document/medical-body-area-networks-first-report-and-order.

[97] S.D. Baker and D.H. Hoglund, "Medical-grade, mission-critical wireless networks [designing an enterprise mobility solution in the healthcare environment]," *Engineering in Medicine and Biology Magazine, IEEE*, vol. 27, no. 2, pp. 86 –95, march-april 2008.

[98] ETSI, "Electromagnetic compatibility and radio spectrum matters (erm); system reference document (srdoc); medical body area network systems (mbanss) in the 1785 MHz to 2500 MHz range," Tech. Rep. ETSI TR 101 557 V1.1.1 (2012-02), February 2012.

[99] DASH7 Alliance, URL: http://www.dash7.org/.

[100] S. Kim, L. Beckmann, M. Pistor, L. Cousin, M. Walter, and S. Leonhardt, "A versatile body sensor network for health care applications," in *Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP), 2009 5th International Conference on*, dec. 2009, pp. 175 –180.

[101] Bundesnetzagentur, "Allgemeinzuteilung von Frequenzen für die Nutzung durch Anwendungen geringer Leistung der Ultra- Wideband (UWB) Technologie," http://www.bundesnetzagentur.de/cae/servlet/contentblob/32374/publicationFile/2528/AllgemeinzuteilungId12424pdf.pdf, 2008.

[102] L. Ahlin and J. Zander, *Principles of Digital Radio Communications*, Chartwell Learning & Development Limited, 1997.

[103] Gordon L. Stübner, *Principles of mobile communication*, Kluwer Academic Publishers, 2001.

[104] B. Sklar, "Rayleigh fading channels in mobile digital communication systems.," *Comm. Mag.*, vol. 35, no. 9, pp. 136–146, Sept. 1997.

[105] Ruijun Fu, Yunxing Ye, Ning Yang, and K. Pahlavan, "Doppler spread analysis of human motions for body area network applications," in *Personal Indoor and Mobile Radio Communications (PIMRC), 2011 IEEE 22nd International Symposium on*, sept. 2011, pp. 2209 –2213.

[106] D.B. Smith, J. Zhang, L.W. Hanlen, D. Miniutti, D. Rodda, and B. Gilbert, "Temporal correlation of dynamic on-body area radio channel," *Electronics Letters*, vol. 45, no. 24, pp. 1212 –1213, 19 2009.

[107] J. Proakis and M. Salehi, *Digital Communications*, McGraw-Hill higher education. McGraw-Hill Companies,Incorporated, 2008.

[108] E. Reusens, W. Joseph, G. Vermeeren, L. Martens, B. Latre, I. Moerman, B. Braem, and C. Blondia, "Path loss models for wireless communication channel along arm and torso: measurements and simulations," in *Antennas and Propagation Society International Symposium, 2007 IEEE*, june 2007, pp. 345 –348.

[109] Takahiro Aoyagi, Jun-ichi Takada, Kenichi Takizawa, Hirokazu Sawada, Norihiko Katayama, Kamya Yazdandoost, Takehiko Kobayashi, Huan-Bang Li, and Ryuji Kohno, "IEEE P802.15 Working Group for Wireless Personal Area Networks (WPANs) - Channel Model for Body Area Network (BAN)," PDF, 2008.

[110] L. Roelens, S. Van den Bulcke, W. Joseph, G. Vermeeren, and L. Martens, "Path loss model for wireless narrowband communication above flat phantom," *Electronics Letters*, vol. 42, no. 1, pp. 10–11, Jan. 2006.

[111] S.L. Cotton, G.A. Conway, and W.G. Scanlon, "A time-domain approach to the analysis and modeling of on-body propagation characteristics using synchronized measurements at 2.45 ghz," *Antennas and Propagation, IEEE Transactions on*, vol. 57, no. 4, pp. 943 –955, april 2009.

[112] J.S. Seybold, *Introduction to RF Propagation*, Wiley, 2005.

[113] J. Ryckaert, P. De Doncker, R. Meys, A. de Le Hoye, and S. Donnay, "Channel model for wireless communication around human body," *Electronics Letters*, vol. 40, no. 9, pp. 543 – 544, april 2004.

[114] David Reiss, "Electromagnetic Surface Wave. The Net Advance of Physics: Special reports, No. 1," 1996.

[115] K.C. Huang and Z. Wang, *Millimeter Wave Communication Systems*, IEEE Series on Digital & Mobile Communication. Wiley, 2011.

[116] Emiliano Miluzzo, Xiao Zheng, Kristóf Fodor, and Andrew T. Campbell, "Radio characterization of 802.15.4 and its impact on the design of mobile sensor networks," in *EWSN'08: Proceedings of the 5th European conference on Wireless sensor networks*, Berlin, Heidelberg, 2008, pp. 171–188, Springer-Verlag.

[117] R. Rosini, R. D'Errico, and R. Verdone, "Body-to-body communications: A measurement-based channel model at 2.45 GHz," in *Personal Indoor and Mobile Radio Communications (PIMRC), 2012 IEEE 23rd International Symposium on*, 2012, pp. 1763–1768.

[118] Y.I. Nechayev, Z.H. Hu, and P.S. Hall, "Short-term and long-term fading of on-body transmission channels at 2.45 GHz," in *Antennas Propagation Conference, 2009. LAPC 2009. Loughborough*, nov. 2009, pp. 657 –660.

[119] A. Fort, C. Desset, P. Wambacq, and L.V. Biesen, "Indoor body-area channel model for narrowband communications," *Microwaves, Antennas Propagation, IET*, vol. 1, no. 6, pp. 1197 –1203, dec. 2007.

[120] Simon L. Cotton and William G. Scanlon, "An experimental investigation into the influence of user state and environment on fading characteristics in wireless body area networks at 2.45 ghz," *Trans. Wireless. Comm.*, vol. 8, no. 1, pp. 6–12, Jan. 2009.

[121] Athanassios Boulis, Yuriy Tselishchev, Lavy Libman, David Smith, and Leif Hanlen, "Impact of wireless channel temporal variation on MAC design for body area networks," *ACM Trans. Embed. Comput. Syst.*, vol. 11, no. S2, pp. 51:1–51:18, Aug. 2012.

[122] Simon L. Cotton and William G. Scanlon, *Channel Modeling of Narrowband Body-Centric Wireless Communication Systems*, chapter 11, Pan Stanford Publishing, 2011.

[123] K Yazdandoost and K Sayrafian-Pour, "IEEE P802.15 working group for wireless personal area networks (WPANs) - channel model for body area network (BAN)," https://mentor.ieee.org/802.15/dcn/08/15-08-0780-12-0006-tg6-channel-model.pdf, 2010.

[124] M. Gallo, P.S. Hall, Y.I. Nechayev, and M. Bozzetti, "Use of animation software in simulation of on-body communications channels at 2.45 ghz," *Antennas and Wireless Propagation Letters, IEEE*, vol. 7, pp. 321 –324, 2008.

[125] P.S. Hall, M. Ricci, and T.M. Hee, "Measurements of on-body propagation characteristics," in *Antennas and Propagation Society International Symposium, 2002. IEEE*, 2002, vol. 2, pp. 310 – 313 vol.2.

[126] P.S. Hall, Yang Hao, Y.I. Nechayev, A. Alomalny, C.C. Constantinou, C. Parini, M.R. Kamarudin, T.Z. Salim, D.T.M. Hee, R. Dubrovka, A.S. Owadally, Wei Song, A. Serra, P. Nepa, M. Gallo, and M. Bozzetti, "Antennas and propagation for on-body communication systems," *Antennas and Propagation Magazine, IEEE*, vol. 49, no. 3, pp. 41 –58, june 2007.

[127] Dino Miniutti, Leif Hanlen, David Smith, Andrew Zhang, Daniel Lewis, David Rodda, and Ben Gilbert, "IEEE P802.15 working group for wireless personal area networks (WPANs) - narrowband channel characterization for body area networks," Tech. Rep. 15-08-0421-00-0006, IEEE P802.15, Jul 2008.

[128] J. Cai, S. Cheng, and C. Huang, "MAC channel model for WBAN," Tech. Rep. 15-09-0562-00-0006, IEEE P802.15, jul 2009.

[129] N.E. Roberts, Seunghyun Oh, and D.D. Wentzloff, "Exploiting channel periodicity in body sensor networks," *Emerging and Selected Topics in Circuits and Systems, IEEE Journal on*, vol. 2, no. 1, pp. 4 –13, march 2012.

[130] Y. Tselishchev, L. Libman, and A. Boulis, "Reducing transmission losses in body area networks using variable TDMA scheduling," in *World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2011 IEEE International Symposium on a*, june 2011, pp. 1 −10.

[131] E. N. Gilbert, "Capacity of a burst-noise channel," vol. 39, no. 5, pp. 1253–1265, Sept. 1960.

[132] M. M. Zonoozi and P. Dassanayake, "User mobility modeling and characterization of mobility patterns," *IEEE J.Sel. A. Commun.*, vol. 15, no. 7, pp. 1239–1252, Sept. 2006.

[133] Majid Nabi, Marc Geilen, and Twan Basten, "MoBAN: A configurable mobility model for wireless body area networks," in *4th International Conference on Simulation Tools and Techniques*. March 2011, ICST.

[134] Andreas Köpke, Michael Swigulski, Karl Wessel, Daniel Willkomm, P.T. Klein Haneveld, Tom Parker, Otto Visser, Hermann Simon Lichte, and Stefan Valentin, "Simulating wireless and mobile networks in omnet++ – the mixim vision," in *OMNeT++ 2008: Proceedings of the 1st International Workshop on OMNeT++ (hosted by SIMUTools 2008)*, ICST, Brussels, Belgium, Belgium, 2008, ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering).

[135] OMNeT++ Network Simulator, URL: http://www.omnetpp.org.

[136] Athanassios Boulis, "Castalia: revealing pitfalls in designing distributed algorithms in wsn," in *SenSys '07: Proceedings of the 5th international conference on Embedded networked sensor systems*, New York, NY, USA, 2007, pp. 407–408, ACM.

[137] Wei Ye, J. Heidemann, and D. Estrin, "An energy-efficient MAC protocol for wireless sensor networks," in *INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, 2002, vol. 3, pp. 1567 – 1576 vol.3.

[138] Pei Huang, Li Xiao, S. Soltani, M.W. Mutka, and Ning Xi, "The evolution of MAC protocols in wireless sensor networks: A survey," *Communications Surveys Tutorials, IEEE*, vol. 15, no. 1, pp. 101 –120, quarter 2013.

[139] A. Bachir, M. Dohler, T. Watteyne, and K.K. Leung, "MAC essentials for wireless sensor networks," *Communications Surveys Tutorials, IEEE*, vol. 12, no. 2, pp. 222 –248, quarter 2010.

[140] Dongheui Yun, Seong-Eun Yoo, Daeyoung Kim, and Dohyeun Kim, "OD-MAC: An on-demand MAC protocol for body sensor networks based on IEEE 802.15.4," in *Embedded and Real-Time Computing Systems and Applications, 2008. RTCSA '08. 14th IEEE International Conference on*, 2008, pp. 413 –420.

[141] Gengfa Fang and E. Dutkiewicz, "Bodymac: Energy efficient TDMA-based MAC protocol for wireless body area networks," in *Communications and Information Technology, 2009. ISCIT 2009. 9th International Symposium on*, sept. 2009, pp. 1455 –1459.

[142] Yan Zhang and G. Dolmans, "A new priority-guaranteed mac protocol for emerging body area networks," in *Wireless and Mobile Communications, 2009. ICWMC '09. Fifth International Conference on*, 2009, pp. 140 –145.

[143] Huaming Li and Jindong Tan, "Heartbeat-driven medium-access control for body sensor networks," *Information Technology in Biomedicine, IEEE Transactions on*, vol. 14, no. 1, pp. 44 –51, jan 2010.

[144] Hang Su and Xi Zhang, "Battery-dynamics driven tdma mac protocols for wireless body-area monitoring networks in healthcare applications," *Selected Areas in Communications, IEEE Journal on*, vol. 27, no. 4, pp. 424 –434, may 2009.

[145] A. Boulis, D. Smith, D. Miniutti, L. Libman, and Y. Tselishchev, "Challenges in body area networks for healthcare: the MAC," *Communications Magazine, IEEE*, vol. 50, no. 5, pp. 100 –106, may 2012.

[146] "IEEE Standard for Information Technology - Telecommunications and Information Exchange Between Systems - Local and Metropolitan Area Networks - Specific Requirements. - Part 15.1: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Wireless Personal Area Networks (WPANs)," *IEEE Std 802.15.1-2005 (Revision of IEEE Std 802.15.1-2002)*, 2005.

[147] Carles Gomez, Joaquim Oller, and Josep Paradells, "Overview and evaluation of bluetooth low energy: An emerging low-power wireless technology," *Sensors*, vol. 12, no. 9, pp. 11734–11753, 2012.

[148] M. Siekkinen, M. Hiienkari, J.K. Nurminen, and J. Nieminen, "How low energy is bluetooth low energy? Comparative measurements with ZigBee/802.15.4," in *Wireless Communications and Networking Conference Workshops (WCNCW), 2012 IEEE*, april 2012, pp. 232 –237.

[149] E. Mackensen, M. Lai, and T.M. Wendt, "Performance analysis of an Bluetooth low energy sensor system," in *Wireless Systems (IDAACS-SWS), 2012 IEEE 1st International Symposium on*, sept. 2012, pp. 62 –66.

[150] M. Ali, L. Albasha, and H. Al-Nashash, "A Bluetooth low energy implantable glucose monitoring system," in *Radar Conference (EuRAD), 2011 European*, oct. 2011, pp. 377 –380.

[151] Bin Yu, Lisheng Xu, and Yongxu Li, "Bluetooth low energy (BLE) based mobile electrocardiogram monitoring system," in *Information and Automation (ICIA), 2012 International Conference on*, june 2012, pp. 763 –767.

[152] "A 1 v 5 ma multimode IEEE 802.15.6/Bluetooth low-energy WBAN transceiver for biotelemetry applications," *Solid-State Circuits, IEEE Journal of*, vol. 48, no. 1, pp. 186 –198, Jan 2013.

[153] LAN/MAN Standards Committee of the IEEE Computer Society, *IEEE Standard for Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements – Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low Rate Wireless Personal Area Networks (LR-WPANs)*, Oct. 2003.

[154] IETF 6LoWPAN Working Group, "IPv6 over Low power WPAN (6lowpan)," http://www.ietf.org/dyn/wg/charter/6lowpan-charter.html.

[155] HART Communication Foundation, "WirelessHART," http://www.hartcomm.org/protocol/wihart/wireless_technology.html.

[156] ," *IEEE Std 802.15.4a-2007 (Amendment to IEEE Std 802.15.4-2006), title=IEEE Standard for Information Technology - Telecommunications and Information Exchange Between Systems - Local and Metropolitan Area Networks - Specific Requirement Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (WPANs)*, pp. 1 –203, 2007.

[157] ," *IEEE Std 802.15.4e-2012 (Amendment to IEEE Std 802.15.4-2011), title=IEEE Standard for Local and metropolitan area networks–Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs) Amendment 1: MAC sublayer*, pp. 1 –225, 16 2012.

[158] N.F. Timmons and W.G. Scanlon, "Analysis of the performance of IEEE 802.15.4 for medical sensor body area networking," in *Sensor and Ad Hoc Communications and Networks, 2004. IEEE SECON 2004. 2004 First Annual IEEE Communications Society Conference on*, oct. 2004, pp. 16 – 24.

[159] Dharshana Vijayaraghavan, "Performance analysis of IEEE 802.15.4 for body area networks," M.S. thesis, San Diego State University, 2010.

[160] P H. Ghare, A. G. Kothari, and A. G. Keskar, "Evaluation of scalability issue of 802.15.4 MAC for body area networks," *International Journal of Computer and Electrical Engineering*, vol. 4, no. 5, pp. 718–721, 2012.

[161] "TinyOS git repository," February 2013, https://github.com/tinyos/tinyos-main.

[162] Jason Hill, Robert Szewczyk, Alec Woo, Seth Hollar, David Culler, and Kristofer Pister, "System architecture directions for networked sensors," in *Proc. of the ninth international conference on Architectural support for programming languages and operating systems (ASPL 2000)*, 2000.

[163] David Gay, Philip Levis, Robert von Behren, Matt Welsh, Eric Brewer, and David Culler, "The nesC language: A holistic approach to networked embedded systems," in *PLDI '03:*

Proceedings of the ACM SIGPLAN 2003 conference on Programming language design and implementation, New York, NY, USA, 2003, pp. 1–11, ACM Press.

[164] P. Levis, D. Gay, V. Handziski, J.-H.Hauer, B.Greenstein, M.Turon, J.Hui, K.Klues, C.Sharp, R.Szewczyk, J.Polastre, P.Buonadonna, L.Nachman, G.Tolle, D.Culler, and A.Wolisz, "T2: a second generation OS for embedded sensor networks," Tech. Rep. TKN-05-007, Telecommunication Networks Group, Technische Universität Berlin, Nov. 2005.

[165] V.Handziski, J. Polastre, J.-H. Hauer, C. Sharp, A. Wolisz, and D. Culler, "Flexible hardware abstraction for wireless sensor networks," in *Proc. of 2nd European Workshop on Wireless Sensor Networks (EWSN 2005)*, Istanbul, Turkey, Feb. 2005.

[166] Jan Flora and Philippe Bonnet, "Never mind the standard here is the TinyOS 802.15.4 stack," Technical Report 06-10, University of Copenhagen, 2006.

[167] David Gay, Phil Levis, and David Culler, "Software design patterns for TinyOS," in *LCTES '05: Proceedings of the 2005 ACM SIGPLAN/SIGBED conference on Languages, compilers, and tools for embedded systems*, New York, NY, USA, 2005, pp. 40–49, ACM.

[168] MEMSIC Inc., "MICAz wireless measurement system," http://www.memsic.com/, June 2004.

[169] J. Araujo, A. Anta, M. Mazo, J. Faria, A. Hernandez, P. Tabuada, and K.H. Johansson, "Self-triggered control over wireless sensor and actuator networks," in *Distributed Computing in Sensor Systems and Workshops (DCOSS), 2011 International Conference on*, june 2011, pp. 1 –9.

[170] Huasong Cao, Haoming Li, Leo Stocco, and Victor C. M. Leung, "Demonstration of a novel wireless three-pad ecg system for generating conventional 12-lead signals," in *Proceedings of the Fifth International Conference on Body Area Networks*, New York, NY, USA, 2010, BodyNets '10, pp. 29–32, ACM.

[171] N. Karowski, A. Viana, and A. Wolisz, "Optimized asynchronous multi-channel discovery of IEEE 802.15.4-based wireless personal area networks," *Mobile Computing, IEEE Transactions on*, vol. PP, no. 99, pp. 1, 2012.

[172] A. Hernandez, "Wireless process control using IEEE 802.15.4 protocol," M.S. thesis, Royal Institute of Technology (KTH), Sweden, 2010.

[173] R. Daidone, G. Dini, and M. Tiloca, "On experimentally evaluating the impact of security on IEEE 802.15.4 networks," in *Distributed Computing in Sensor Systems and Workshops (DCOSS), 2011 International Conference on*, june 2011, pp. 1 –6.

[174] T. Basmer, H. Schomann, and S. Peter, "Implementation analysis of the IEEE 802.15.4 MAC for wireless sensor networks," in *Mobile and Wireless Networking (iCOST), 2011 International Conference on Selected Topics in*, 2011, pp. 7–12.

175

[175] G. Anastasi, M. Conti, and M. Di Francesco, "A comprehensive analysis of the MAC unreliability problem in IEEE 802.15.4 wireless sensor networks," *Industrial Informatics, IEEE Transactions on*, vol. 7, no. 1, pp. 52 –65, feb. 2011.

[176] E. Toscano and L.L. Bello, "Multichannel superframe scheduling for IEEE 802.15.4 industrial wireless sensor networks," *Industrial Informatics, IEEE Transactions on*, vol. 8, no. 2, pp. 337 –350, may 2012.

[177] HyungJune Lee, A. Cerpa, and P. Levis, "Improving wireless simulation through noise modeling," April 2007, pp. 21–30.

[178] Yin Chen and Andreas Terzis, "On the implications of the log-normal path loss model: an efficient method to deploy and move sensor motes," in *Proceedings of the 9th ACM Conference on Embedded Networked Sensor Systems*, New York, NY, USA, 2011, SenSys '11, pp. 26–39, ACM.

[179] Philip Levis, Nelson Lee, Matt Welsh, and David Culler, "TOSSIM: accurate and scalable simulation of entire tinyos applications," in *Proceedings of the 1st international conference on Embedded networked sensor systems*, New York, NY, USA, 2003, SenSys '03, pp. 126–137, ACM.

[180] K. Pahlavan and P. Krishnamurthy, *Principles of wireless networks: a unified approach*, Prentice Hall communications engineering and emerging technologies series. Prentice Hall PTR, 2002.

[181] P.A. Shaltis, A. Reisner, and H.H. Asada, "Wearable, cuff-less ppg-based blood pressure monitor with novel height sensor," in *Engineering in Medicine and Biology Society, 2006. EMBS '06. 28th Annual International Conference of the IEEE*, 2006, pp. 908–911.

[182] O. Aziz, L. Atallah, B. Lo, M. Elhelw, L. Wang, G. Z. Yang, and A. Darzi, "A pervasive body sensor network for measuring postoperative recovery at home," *Surg Innov*, vol. 14, no. 2, pp. 83–90, June 2007.

[183] Yang Hao and Robert Foster, "Wireless body sensor networks for health-monitoring applications," *Physiological Measurement*, vol. 29, no. 11, pp. R27, 2008.

[184] Ling Bao and StephenS. Intille, "Activity recognition from user-annotated acceleration data," in *Pervasive Computing*, Alois Ferscha and Friedemann Mattern, Eds., vol. 3001 of *Lecture Notes in Computer Science*, pp. 1–17. Springer Berlin Heidelberg, 2004.

[185] R Herren, A Sparti, K Aminian, and Y Schutz, "The prediction of speed and incline in outdoor running in humans using accelerometry," *MEDICINE AND SCIENCE IN SPORTS AND EXERCISE*, vol. 31, no. 7, pp. 1053–1059, JUL 1999.

[186] Cheol-Hong Min, A.H. Tewfik, Youngchun Kim, and R. Menard, "Optimal sensor location for body sensor network to detect self-stimulatory behaviors of children with autism spectrum disorder," in *Engineering in Medicine and Biology Society, 2009. EMBC 2009. Annual International Conference of the IEEE*, 2009, pp. 3489–3492.

[187] L. Brown, B. Grundlehner, J. Van de Molengraft, J. Penders, and B. Gyselinckx, "Body area network for monitoring autonomic nervous system responses," in *Pervasive Computing Technologies for Healthcare, 2009. PervasiveHealth 2009. 3rd International Conference on*, 2009, pp. 1–3.

[188] Thomas Falck, Javier Espina, Jean-Pierre Ebert, and Daniel Dietterle, "Basuma - the sixth sense for chronically ill patients," in *Proceedings of the International Workshop on Wearable and Implantable Body Sensor Networks*, Washington, DC, USA, 2006, BSN '06, pp. 57–60, IEEE Computer Society.

[189] Shan Lin, Jingbin Zhang, Gang Zhou, Lin Gu, John A. Stankovic, and Tian He, "ATPC: adaptive transmission power control for wireless sensor networks," in *Proceedings of the 4th international conference on Embedded networked sensor systems*, New York, NY, USA, 2006, SenSys '06, pp. 223–236, ACM.

[190] Harvard Medical School, "Harvard health letter: Counting every step you take," sep 2009.

[191] C. Tudor-Locke, Y Hatano, RP Pangrazi, and M. Kang, "Revisiting "how many steps are enough?"," 2008.

[192] K. Shashi Prabh and Jan-Hinrich Hauer, "Opportunistic packet scheduling in body area networks," in *Proc. of 8th European Conference on Wireless Sensor Networks (EWSN)*, Bonn, Germany, Feb. 2011.

[193] A. Kara and H.L. Bertoni, "Blockage/shadowing and polarization measurements at 2.45 GHz for interference evaluation between Bluetooth and IEEE 802.11 WLAN," *Antennas and Propagation Society International Symposium, 2001. IEEE*, vol. 3, pp. 376 –379 vol.3, 2001.

[194] Gregor Gaertner, Eamonn ONuallain, Andrew Butterly, Kulpreet Singh, and Vinny Cahill, "802.11 link quality and its prediction - an experimental study," in *Personal Wireless Communications*, Ignas Niemegeers and Sonia Heemstra de Groot, Eds. 2004, vol. 3260 of *Lecture Notes in Computer Science*, pp. 609–611, Springer Berlin / Heidelberg.

[195] M. Ghaddar, L. Talbi, and T.A. Denidni, "Human body modelling for prediction of effect of people on indoor propagation channel," *Electronics Letters*, vol. 40, no. 25, pp. 1592 – 1594, Dec. 2004.

[196] S. Obayashi and J. Zander, "A body-shadowing model for indoor radio communication environments," *Antennas and Propagation, IEEE Transactions on*, vol. 46, no. 6, pp. 920–927, jun. 1998.

[197] Alfonso Bahillo, Ruben M. Lorenzo, Santiago Mazuelas, Patricia Fernandez, and Evaristo J. Abril, *Assessment of the Shadow Caused by the Human Body on the Personal RF Dosimeters Reading in Multipath Environments*, Oct. 2009.

177

[198] P.S. Hall, M. Ricci, and T.W. Hee, "Characterization of on-body communication channels," in *Microwave and Millimeter Wave Technology, 2002. Proceedings. ICMMT 2002. 2002 3rd International Conference on*, Aug. 2002, pp. 770 – 772.

[199] R. D'Errico and L. Ouvry, "Time-variant BAN channel characterization," in *Personal, Indoor and Mobile Radio Communications, 2009 IEEE 20th International Symposium on*, sept. 2009, pp. 3000 –3004.

[200] Jean-Marie Gorce, Claire Goursaud, Guillaume Villemaud, Raffaele D'Errico, and Laurent Ouvry, "Opportunistic relaying protocols for human monitoring in BAN," in *PIMRC*. 2009, pp. 732–736, IEEE.

[201] Dino Miniutti, Leif Hanlen, Andrew Zhang, David Smith, Daniel Lewis, David Rodda, and Ben Gilbert, "Narrowband channel characterization for body area networks," IEEE 802.15.6 standard, July 2008.

[202] K. Shashi Prabh, Fernando Royo, Stefano Tennina, and Teresa Olivares, "BANMAC: An opportunistic mac protocol for reliable communications in body area networks," *Distributed Computing in Sensor Systems and Workshops, International Conference on*, vol. 0, pp. 166–175, 2012.

[203] P. Dutta, S. Dawson-Haggerty, Y. Chen, C.-J. Liang, and A. Terzis, "Design and evaluation of a versatile and efficient receiver-initiated link layer for low-power wireless," in *SenSys'10*, Zurich, Switzerland, Nov. 2010, pp. 1–14.

[204] C.-J. Liang, B. Priyantha, J. Liu, and A. Terzis, "Surviving Wi-Fi interference in low power ZigBee networks," in *SenSys'10*, Zurich, Switzerland, Nov. 2010, pp. 309–322.

[205] R.C. Shah, S. Wietholter, A. Wolisz, and J.M. Rabaey, "When does opportunistic routing make sense?," in *Pervasive Computing and Communications Workshops, 2005. PerCom 2005 Workshops. Third IEEE International Conference on*, 2005, pp. 350–356.

[206] 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 network sensor systems*, New York, NY, USA, 2008, SenSys '08, pp. 29–42, ACM.

[207] Sirajum Munir, Shan Lin, Enamul Hoque, S. M. Shahriar Nirjon, John A. Stankovic, and Kamin Whitehouse, "Addressing burstiness for reliable communication and latency bound generation in wireless sensor networks," in *Proceedings of the 9th ACM/IEEE International Conference on Information Processing in Sensor Networks*, New York, NY, USA, 2010, IPSN '10, pp. 303–314, ACM.

[208] Sushant Jain, Rahul C. Shah, Waylon Brunette, Gaetano Borriello, and Sumit Roy, "Exploiting mobility for energy efficient data collection in wireless sensor networks," *Mob. Netw. Appl.*, vol. 11, pp. 327–339, June 2006.

178

[209] Aaron Schulman, Vishnu Navda, Ramachandran Ramjee, Neil Spring, Pralhad Deshpande, Calvin Grunewald, Kamal Jain, and Venkata N. Padmanabhan, "Bartendr: a practical approach to energy-aware cellular data scheduling," in *Proceedings of the sixteenth annual international conference on Mobile computing and networking*, New York, NY, USA, 2010, MobiCom '10, pp. 85–96, ACM.

[210] Shuo Xiao, Vijay Sivaraman, and Alison Burdett, "Adapting radio transmit power in wireless body area sensor networks," in *Proceedings of the ICST 3rd international conference on Body area networks*, ICST, Brussels, Belgium, Belgium, 2008, BodyNets '08, pp. 14:1–14:8, ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering).

[211] Ashay Dhamdhere, Vijay Sivaraman, Vidit Mathur, and Shuo Xiao, "Algorithms for transmission power control in biomedical wireless sensor networks," in *Proceedings of the 2008 IEEE Asia-Pacific Services Computing Conference*, Washington, DC, USA, 2008, pp. 1114–1119, IEEE Computer Society.

[212] B. Moulton, L. Hanlen, June Chen, G. Croucher, L. Mahendran, and A. Varis, "Body-area-network transmission power control using variable adaptive feedback periodicity," in *Communications Theory Workshop (AusCTW), 2010 Australian*, Feb. 2010, pp. 139 –144.

[213] R. Kazemi, R. Vesilo, E. Dutkiewicz, and Ren Liu, "Dynamic power control in wireless body area networks using reinforcement learning with approximation," in *Personal Indoor and Mobile Radio Communications (PIMRC), 2011 IEEE 22nd International Symposium on*, sept. 2011, pp. 2203 –2208.

[214] D.B. Smith, L.W. Hanlen, and D. Miniutti, "Transmit power control for wireless body area networks using novel channel prediction," in *Wireless Communications and Networking Conference (WCNC), 2012 IEEE*, april 2012, pp. 684 –688.

[215] Hong Ying, C. Silex, A. Schnitzer, S. Leonhardt, and M. Schiek, "Automatic step detection in the accelerometer signal," in *4th International Workshop on Wearable and Implantable Body Sensor Networks (BSN 2007)*, R. Magjarevic, J. H. Nagel, Steffen Leonhardt, Thomas Falck, and Petri Máhónen, Eds. 2007, vol. 13 of *IFMBE Proceedings*, pp. 80–85, Springer Berlin Heidelberg.

[216] Michael Marschollek, Mehmet Goevercin, Klaus-Hendrik Wolf, Bianying Song, Matthias Gietzelt, Reinhold Haux, and Elisabeth Steinhagen-Thiessen, "A performance comparison of accelerometry-based step detection algorithms on a large, non-laboratory sample of healthy and mobility-impaired persons," aug 2008, pp. 1319 –1322.

[217] M. Ross, H. Shaffer, A. Cohen, R. Freudberg, and H. Manley, "Average magnitude difference function pitch extractor," *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 22, no. 5, pp. 353 – 362, oct. 1974.

179

[218] Eamonn J. Keogh and Michael J. Pazzani, "Scaling up dynamic time warping for datamining applications," in *Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining*, New York, NY, USA, 2000, KDD '00, pp. 285–289, ACM.

[219] Kannan Srinivasan and Philip Levis, "RSSI is under appreciated," in *In Proceedings of the Third Workshop on Embedded Networked Sensors (EmNets06)*, 2006.

[220] Rainer E. Burkard, Mauro Dell'Amico, and Silvano Martello, *Assignment Problems*, SIAM, 2009.

[221] G.B. Dantzig, "Linear programming and extensions," in *Landmarks in Physics and Mathematics*. 1998, Princeton University Press.

[222] Thomas S. Ferguson, "Linear programming - a concise introduction," http://www.math.ucla.edu/~tom/, 1995.

[223] George L. Nemhauser and Laurence A. Wolsey, *Integer and combinatorial optimization*, Wiley-Interscience, NY, USA, 1988.

[224] Tobias Achterberg, "SCIP: Solving constraint integer programs," *Mathematical Programming Computation*, vol. 1, no. 1, pp. 1–41, July 2009, http://mpc.zib.de/index.php/MPC/article/view/4.

[225] Douglas S. J. De Couto, Daniel Aguayo, John Bicket, and Robert Morris, "A high-throughput path metric for multi-hop wireless routing," in *Proceedings of the 9th annual international conference on Mobile computing and networking*, New York, NY, USA, 2003, MobiCom '03, pp. 134–146, ACM.

[226] I. Howitt and J.A. Gutierrez, "IEEE 802.15.4 low rate - wireless personal area network coexistence issues," *Wireless Communications and Networking, 2003. WCNC 2003. 2003 IEEE*, vol. 3, pp. 1481–1486 vol.3, March 2003.

[227] Wei Yuan, Xiangyu Wang, and J. P M G Linnartz, "A coexistence model of IEEE 802.15.4 and IEEE 802.11b/g," in *Communications and Vehicular Technology in the Benelux, 2007 14th IEEE Symposium on*, 2007, pp. 1–5.

[228] Dae Gil Yoon, Soo Young Shin, Wook Hyun Kwon, and Hong Seong Park, "Packet error rate analysis of IEEE 802.15.4 under IEEE 802.11b interference," *Vehicular Technology Conference, 2006. VTC 2006-Spring. IEEE 63rd*, vol. 3, pp. 1186–1190, May 2006.

[229] Wei Yuan, Xiangyu Wang, Jean-Paul Linnartz, and Ignas Niemegeers, "Experimental validation of a coexistence model of IEEE 802.15.4 and IEEE 802.11b/g networks," *International Journal of Distributed Sensor Networks*, 2010.

[230] M. Petrova, J. Riihijarvi, P. Mahonen, and S. Labella, "Performance study of IEEE 802.15.4 using measurements and simulations," in *Wireless Communications and Networking Conference, 2006. WCNC 2006. IEEE*, 2006, vol. 1, pp. 487–492.

[231] Nada Golmie, David Cypher, and Olivier Rébala, "Performance analysis of low rate wireless technologies for medical applications," *Computer Communications*, vol. 28, no. 10, pp. 1255–1275, 2005.

[232] M. Petrova, Lili Wu, P. Mahonen, and J. Riihijarvi, "Interference measurements on performance degradation between colocated IEEE 802.11g/n and IEEE 802.15.4 networks," in *Networking, 2007. ICN '07. Sixth International Conference on*, 2007, pp. 93–93.

[233] S. Pollin, I. Tan, B. Hodge, C. Chun, and A. Bahai, "Harmful coexistence between 802.15.4 and 802.11: A measurement-based study," *Cognitive Radio Oriented Wireless Networks and Communications, 2008. CrownCom 2008. 3rd International Conference on*, pp. 1–6, May 2008.

[234] Jennic Ltd, "Co-existence of IEEE 802.15.4 at 2.4 GHz - application note," http://www.jennic.com/.

[235] Gilles Thonet, Patrick Allard-Jacquin, and Pierre Colle, "ZigBee - WiFi coexistence, white paper and test report," Tech. Rep., Schneider Electric, 2008.

[236] Freescale Semiconductor, "Mc1319x coexistence - application note," http://www.freescale.com.

[237] Juyong Do, D.M. Akos, and P.K. Enge, "L and s bands spectrum survey in the san francisco bay area," in *Position Location and Navigation Symposium, 2004. PLANS 2004*, april 2004.

[238] S.A. Hanna and J. Sydor, "Distributed sensing of spectrum occupancy and interference in outdoor 2.4 GHz Wi-Fi networks," in *Global Communications Conference (GLOBECOM), 2012 IEEE*, 2012, pp. 1453–1459.

[239] Mo Sha, Gregory Hackmann, and Chenyang Lu, "Multichannel reliability and spectrum usage in real homes: empirical studies for home-area sensor networks," in *Proc. of IWQoS'11*, San Jose, CA, U.S.

[240] V. Valenta, Z. Fedra, R. Marsalek, G. Baudoin, and M. Villegas, "Analysis of spectrum utilization in suburb environment - evaluation of potentials for cognitive radio," in *ICUMT'09*, oct. 2009.

[241] L. Hanlen, D. Miniutti, D. Smith, D. Rodda, and B. Gilbert, "Co-channel interference in body area networks with indoor measurements at 2.4 GHz: distance-to-interferer is a poor estimate of received interference power," *International Journal of Wireless Information Networks*, 2010.

[242] MetaGeek, "Wi-Spy 2.4x," http://www.metageek.net/products/wi-spy.

[243] J. H. Hauer, V.Handziski, and A.Wolisz, "Experimental study of the impact of WLAN interference on IEEE 802.15.4 body area networks," in *Proc. of 6th European Conference on Wireless Sensor Networks (EWSN)*, Cork, Ireland, Feb. 2009.

[244] Sentilla Corporation, "Tmote sky datasheet," http://www.sentilla.com/pdf/eol/tmote-sky-datasheet.pdf, 2004.

[245] Jun Huang, Guoliang Xing, Gang Zhou, and Ruogu Zhou, "Beyond co-existence: Exploiting WiFi white space for ZigBee performance assurance," in *Network Protocols (ICNP), 2010 18th IEEE International Conference on*, 2010, pp. 305–314.

[246] Bin Zhen, Huan-Bang Li, Shinsuke Hara, and Ryuji Kohno, "Clear channel assessment in integrated medical environments," *EURASIP J. Wirel. Commun. Netw.*, vol. 2008, pp. 48:1–48:8, Jan. 2008.

[247] S. Myers, S. Megerian, S. Banerjee, and M. Potkonjak, "Experimental investigation of IEEE 802.15.4 transmission power control and interference minimization," in *Sensor, Mesh and Ad Hoc Communications and Networks, 2007. SECON '07. 4th Annual IEEE Communications Society Conference on*, 2007, pp. 294–303.

[248] F. Martelli, R. Verdone, and C. Buratti, "Link adaptation in IEEE 802.15.4-based wireless body area networks," in *Personal, Indoor and Mobile Radio Communications Workshops (PIMRC Workshops), 2010 IEEE 21st International Symposium on*, 2010, pp. 117–121.

[249] Dong Yang, Youzhi Xu, and Mikael Gidlund, "Wireless coexistence between IEEE 802.11- and IEEE 802.15.4-based networks: A survey," *International Journal of Distributed Sensor Networks*, 2011.

[250] Dong Yang, Youzhi Xu, and M. Gidlund, "Coexistence of IEEE 802.15.4 based networks: A survey," in *IECON 2010 - 36th Annual Conference on IEEE Industrial Electronics Society*, 2010, pp. 2107–2113.

[251] Luca Stabellini and Mohammad Mohsen Parhizkar, "Experimental comparison of frequency hopping techniques for 802.15.4-based sensor networks," in *The Fourth International Conference on Mobile Ubiquitous Computing, Systems, Services and Technologies UBICOMM 2010*, 2010, pp. 2107–2113.

[252] Xiuming Zhu, Song Han, Pei-Chi Huang, A.K. Mok, and Deji Chen, "Mbstar: A real-time communication protocol for wireless body area networks," in *Real-Time Systems (ECRTS), 2011 23rd Euromicro Conference on*, 2011, pp. 57–66.

[253] Linjia Yao, S.T. Ali, V. Sivaraman, and D. Ostry, "Decorrelating secret bit extraction via channel hopping in body area networks," in *Personal Indoor and Mobile Radio Communications (PIMRC), 2012 IEEE 23rd International Symposium on*, 2012, pp. 1454–1459.

[254] E.T. Yazdi, A. Willig, and K. Pawlikowski, "On channel adaptation in IEEE 802.15.4 mobile body sensor networks: What can be gained?," in *Networks (ICON), 2012 18th IEEE International Conference on*, 2012, pp. 262–267.

[255] Castalia simulator, URL: http://castalia.npc.nicta.com.au.

[256] Abdel-Ghant A. Dairaseh and Carl W. Baum, "Methods for packet combining in HARQ systems over bursty channels," vol. 2, pp. 213–224, Oct. 1997.

[257] Stefan Valentin, Tobias Volkhausen, Furuzan Atay Onat, Halim Yanikomeroglu, and Holger Karl, "Enabling partial forwarding by decoding-based one and two-stage selective cooperation," in *Proc. IEEE Cognitive and Cooperative Wireless Networks (CoCoNet) co-located with IEEE ICC*, Beijing, China, May 2008.

[258] Andreas Willig, "Memory-efficient segment-based packet-combining schemes in face of deadlines," *IEEE Transactions on Industrial Informatics*, vol. 5, no. 3, pp. 338–350, Aug. 2009.

[259] M. Demirbas, O. Soysal, and M. Hussain, "A singlehop collaborative feedback primitive for wireless sensor networks," April 2008, pp. 2047–2055.

[260] Joseph Polastre, Jason Hill, and David Culler, "Versatile low power media access for wireless sensor networks," in *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems*, New York, NY, USA, 2004, pp. 95–107, ACM.

[261] S. Rayanchu, A. Mishra, D. Agrawal, S. Saha, and Suman Banerjee, "Diagnosing wireless packet losses in 802.11: Separating collision from weak signal," in *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, 2008, pp. 735–743.

[262] Wendong Hu, D. Willkomm, M. Abusubaih, J. Gross, G. Vlantis, M. Gerla, and A. Wolisz, "Cognitive radios for dynamic spectrum access - dynamic frequency hopping communities for efficient IEEE 802.22 operation," *Communications Magazine, IEEE*, vol. 45, no. 5, pp. 80–87, 2007.

[263] Kristofer S. J. Pister and Lance Doherty, "TSMP: time synchronized mesh protocol," in *In Proceedings of the IASTED International Symposium on Distributed Sensor Networks (DSN08)*, 2008.

[264] X. Vilajosana and C. Pister, "Minimal 6TSCH configuration," Tech. Rep. draft-vilajosana-6tsch-basic-01, IETF Secretariat, July 2013.

[265] Ritesh Maheshwari, Shweta Jain, and Samir R. Das, "On estimating joint interference for concurrent packet transmissions in low power wireless networks," in *WiNTECH '08: Proceedings of the third ACM international workshop on Wireless network testbeds, experimental evaluation and characterization*, New York, NY, USA, 2008, pp. 89–94, ACM.

[266] M.Abusubaih, Berthold Rathke, and Adam Wolisz, "Packet loss discrimination in multi-cell 802.11 wireless LANs," Tech. Rep. TKN-08-010, Telecommunication Networks Group, Technische Universität Berlin, Oct. 2008.

[267] Paramvir Bahl, Ranveer Chandra, Thomas Moscibroda, Rohan Murty, and Matt Welsh, "White space networking with Wi-Fi like connectivity," in *SIGCOMM '09: Proceedings*

of the ACM SIGCOMM 2009 conference on Data communication, New York, NY, USA, 2009, pp. 27–38, ACM.

[268] Hang Liu, Hairuo Ma, Magda El Zarki, and Sanjay Gupta, "Error control schemes for networks: an overview," *Mob. Netw. Appl.*, vol. 2, no. 2, pp. 167–182, 1997.

[269] J. R. Norris, *Markov Chains*, Cambridge University Press, Cambridge, UK, 1997.