Interference Robustness and Security of Impulse-Radio Ultra-Wide Band Networks

THÈSE N^{O} 4698 (2010)

PRÉSENTÉE LE 11 JUIN 2010

À LA FACULTÉ INFORMATIQUE ET COMMUNICATIONS
LABORATOIRE POUR LES COMMUNICATIONS INFORMATIQUES ET LEURS APPLICATIONS 2
PROGRAMME DOCTORAL EN INFORMATIQUE, COMMUNICATIONS ET INFORMATION

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Manuel FLURY

acceptée sur proposition du jury:

Prof. C. Petitpierre, président du jury
Prof. J.-Y. Le Boudec, directeur de thèse
Prof. J.-P. Hubaux, rapporteur
Prof. L. Lampe, rapporteur
Prof. I. Oppermann, rapporteur

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE
Suisse
2010

Abstract

In this thesis we study Impulse-Radio Ultra-Wide Band (IR-UWB), a physical layer radio technology offering many features that make it a promising choice for future short-range wireless networks. The challenges in such networks are many, ranging from the cost-complexity constraints of devices, through the presence of interference created by other users, up to stringent security requirements imposed by sensitive applications.

Our main goal is to understand and show how a low-complexity IR-UWB receiver can be designed such that it is able to cope with the difficult environment that it will face in such networks. Although IR-UWB systems promise to provide a solution for some of the above-mentioned challenges, IR-UWB is not a panacea: More often than not, it will be able to live up to its promises only if the entire system is carefully designed.

One example is robustness to interference from concurrent users, which is the topic of the first part of this thesis. Short-range wireless networks are expected to be self-organized and uncoordinated rather than centrally organized. This in turn leads to uncontrolled interference due to concurrent transmissions from uncoordinated devices. Thanks to its large bandwidth, combined with low duty-cycle transmissions, IR-UWB should in theory be able to accommodate a large number of concurrent users while keeping multi-user interference (MUI) levels low. We show that, if not properly addressed, MUI can severely affect the performance of an IR-UWB receiver, making this benefit of IR-UWB void. This is especially true if low complexity architectures, such as the popular non-coherent energy-detection receiver, are used. Further, we show that MUI affects all aspects of physical layer packet reception and appropriate algorithms to deal with it are thus required at every level.

The first crucial step to receive an IR-UWB data packet is signal acquisition. We present a robust and low-complexity algorithm that allows for reliable signal acquisition with an IR-UWB energy-detection receiver in the presence of MUI, even in near-far scenarios.

After signal acquisition, the receiver performs a phase of channel estimation. Channel estimation is of particular importance for interference mitigation: it allows the receiver to dis-

tinguish the signal of the user of interest from MUI. In the case of energy-detection receivers that are compliant with the IR-UWB standard IEEE 802.15.4a, channel estimation is especially challenging because with this standard the signalling structure changes within a data packet. We introduce a novel receiver structure that takes this peculiarity into account and allows for the design of robust low-complexity receivers for IEEE 802.15.4a networks.

The final step in receiving a data packet is demodulation and decoding of the payload. We show that an adaptive thresholding scheme that uses the channel state information, obtained during channel estimation, can yield very good robustness against MUI. We also introduce more sophisticated algorithms that are based on statistical interference modeling and show that they yield an additional increase in robustness against MUI.

In the second part of this thesis we investigate clock-offset tracking for IR-UWB energy-detection receivers. Clock-offset tracking is needed because the oscillators driving the clocks of low-complexity receivers are of average quality at best. We show that the resulting de-synchronization between transmitter and receiver may lead to a huge performance degradation in the case of adaptive energy-detection receivers. To overcome this sensitivity to clock offsets, we present a clock-offset tracking algorithm that is constructed around the Radon transform, an image processing tool traditionally used to detect line features in images. Our algorithm is fully compatible with the IEEE 802.15.4a standard, does not increase the hardware complexity of the receiver and reduces the performance loss due to clock offsets to a marginal level.

In the third part of this thesis, we look at IR-UWB from the viewpoint of security. We evaluate to what extent IEEE 802.15.4a is vulnerable to distance-decreasing attacks on the physical layer. These attacks target the ranging mechanism that allows two wireless devices to estimate their mutual distance. Commonly, ranging is secured by secure ranging protocols employing cryptographic mechanisms that guarantee that the estimated distance is an upper-bound on the actual distance. However, a new breed of attacks bypasses these cryptographic mechanisms, introduced at higher communication layers, by directly attacking the physical layer. Understanding the impact of these attacks on IEEE 802.15.4a is of the utmost importance: its high precision ranging capabilities make IR-UWB a natural candidate for ranging applications, and IEEE 802.15.4a is the only wireless standard that has been specifically designed for ranging. Our analysis shows that IEEE 802.15.4a, does not automatically provide security against such attacks. We find that with the mandatory modes of the standard and no appropriate countermeasures in place, an external attacker can decrease the measured distance by more than one hundred meters with a very high probability.

Keywords

Ultra-wide band, UWB, impulse radio, IEEE 802.15.4a, multi-user interference, interference mitigation, energy-detection, synchronization, clock-offset tracking, wireless security, ranging.

Résumé

Dans cette thèse, nous étudions la radio par impulsions à bande ultra-large ou *Impulse-Radio Ultra-Wide Band* (IR-UWB). Cette technologie de couche physique radio offre de nombreuses fonctionnalités, qui en font un choix prometteur pour les futurs réseaux sans fils à courte portée. Les défis dans de tels réseaux sont nombreux : des contraintes de coût et de complexité des dispositifs, la présence d'interférences créées par d'autres utilisateurs ou encore des exigences de sécurité rigoureuses imposées par des applications sensibles.

Notre objectif principal est de comprendre et de montrer comment un récepteur IR-UWB à faible complexité doit être conçu de telle manière à ce qu'il soit en mesure de faire face aux environnements difficiles qu'il va rencontrer dans de tels réseaux. Bien que les systèmes IR-UWB promettent de fournir une solution pour certains des défis évoqués plus haut, la technologie IR-UWB n'est pas une panacée. En effet, la plupart du temps, elle ne sera en mesure d'honorer ses promesses que si l'ensemble du système est conçu avec soin.

Un exemple, qui est le sujet de la première partie de cette thèse, est la robustesse à l'interférence multi-utilisateurs (IMU). Plutôt que d'être organisés d'une manière centrale, les réseaux sans fils à courte portée vont être auto-organisés. Cela entraîne à son tour de l'interférence incontrôlée due aux transmissions simultanées par des dispositifs sans coordination. Grâce à sa grande largeur de bande, IR-UWB devrait en théorie être en mesure d'accueillir un grand nombre d'utilisateurs simultanés, tout en gardant l'IMU à un faible niveau. Nous montrons que ce bénéfice d'IR-UWB peut être complètement annulé si l'IMU n'est pas traitée correctement. Ceci est particulièrement vrai si des architectures à faible complexité, tels que le détecteur d'énergie, sont utilisées. Nous montrons que l'IMU affecte tous les aspects de la couche physique et que des algorithmes appropriés pour y faire face sont donc requis à tous les niveaux.

La première étape cruciale pour recevoir un paquet de données IR-UWB est l'acquisition du signal. Nous présentons un algorithme robuste et à faible complexité permettant une acquisition fiable du signal en présence d'IMU. Après l'acquisition du signal, le récepteur effectue une phase d'estimation du canal de propagation. L'estimation du canal est d'une importance par-

ticulière pour l'atténuation de l'interférence : elle permet au récepteur de distinguer le signal utile de l'IMU. Dans le cas des détecteurs d'énergie qui se conforment à la norme IR-UWB IEEE 802.15.4a, l'estimation du canal est particulièrement difficile parce qu'avec cette norme, le format du signal change à l'intérieur d'un paquet. Nous introduisons une structure de réception nouvelle qui prend cette particularité en compte et qui permet la conception de récepteurs à faible complexité pour des réseaux IEEE 802.15.4a.

La dernière étape dans la réception d'un paquet de données est la démodulation et le décodage de l'information transmise. Nous montrons qu'un simple seuillage adaptatif qui utilise de l'information sur l'état du canal obtenu lors de l'estimation de ce dernier, amène une excellente robustesse contre l'IMU. Nous introduisons également des algorithmes plus sophistiqués qui sont basés sur la modélisation statistique de l'IMU et nous montrons que ces algorithmes résultent dans une augmentation supplémentaire de la robustesse face à l'IMU.

Dans la seconde partie de cette thèse, nous étudions l'influence du décalage de l'horloge d'un détecteur d'énergie IR-UWB avec l'horloge de l'émetteur correspondant. Nous montrons que la désynchronisation résultante entre émetteur et récepteur peut, dans le cas d'un détecteur d'énergie adaptatif, entrainer une dégradation considérable de la performance. Nous présentons un algorithme de correction constante du décalage de l'horloge qui est entièrement compatible avec la norme IEEE 802.15.4a, n'augmente pas la complexité du récepteur au niveau hardware, et limite la perte de performance en raison de décalages d'horloge à un niveau marginal.

Dans la troisième partie de cette thèse, nous examinons IR-UWB du point de vue de la sécurité. Nous évaluons dans quelle mesure IEEE 802.15.4a est vulnérable à des attaques qui visent le mécanisme du "ranging". Le ranging est un mécanisme qui permet à deux dispositifs sans-fil d'estimer leur distance réciproque. Habituellement, la sécurité de tels mécanismes est garantie par des protocoles de ranging sécurisés qui emploient des mécanismes cryptographiques pour garantir que l'estimée de la distance constitue une borne supérieure de la distance réelle. Toutefois, un nouveau type d'attaques contourne ces mécanismes en s'attaquant directement à la couche physique. Comprendre l'impact de ces attaques sur IEEE 802.15.4a est de la plus haute importance car IR-UWB est un candidat naturel pour des applications de ranging grâce à sa capacité d'exécuter des mesures de distances avec une grande précision. En outre, IEEE 802.15.4a est la seule norme pour la communication sans fil qui a été spécifiquement conçue pour le ranging. Nos analyses montrent que IEEE 802.15.4a, ne fournit pas de sécurité systématique contre de telles attaques : avec les modes obligatoires de la norme IEEE 802.15.4a et sans contre-mesures appropriées en place, un attaquant externe peut diminuer la distance mesurée de plus d'une centaine de mètres et cela avec une probabilité de succès très élevée.

Mots clés

Ultra-large bande, radio impulsive, IEEE 802.15.4a, interférence multi-utilisateurs, atténuation de l'interférence, détection d'énergie, synchronisation, gestion du décalage des horloges, sécurité sans-fil, mesure de distances.

Acknowledgments

I would like to thank Professor Jean-Yves Le Boudec for accepting me in his research group and giving me the opportunity to evolve in such an enriching environment. I can hardly imagine a better advisor, both on a personal and on a scientific level. I am very grateful that during the last few years I could profit from his vast knowledge and experience, including areas extending far beyond IR-UWB networks.

I would also like to thank Professors Jean-Pierre Hubaux, Lutz Lampe, Ian Oppermann and Alain Wegmann, for agreeing to be part of my jury and for taking the time to evaluate this work.

I am much obliged to my colleagues Ruben Merz and Marcin Poturalski. Working together with Ruben on the energy-detection receiver was a great experience. Without his help and knowledge, this thesis would not exist in its present form. I will also not forget our various trips to far corners of the world, trying to convince people that interference is important, while discovering pomelos and other healthy foods. Working with Marcin was an equally pleasant and rewarding experience. I am very thankful to him for opening the world of wireless security to me. Had I not worked with him, this thesis would lack the entire third part and I would still be ignorant of the importance of margins in scientific publications.

Thanks to all the members of LCA, and in particular my officemates Olivier, Jochen and Pedram, for making this PhD a very pleasant journey. I was also very fortunate to encounter *the Lunchers* who made every lunch a special one, sporting everything from culinary highlights like Cappuccino through royal games to controversial discussions. I want to especially thank Senior Lunchers Irina and Dominique as well as their better halves who not only became close friends but also made me discover wonderful remote places like Morges and Grimentz.

Special thanks go to Angela, Danielle, Holly and Patricia for always being there to help with any imaginable non-scientific needs a PhD might possibly entail. I am also much obliged to our system administrators Hervé, Yves, and in particular to Marc-André, for keeping up with my ever increasing demand for more computing power.

I am especially grateful to my family, who is always there for me and supports me in all my plans.

Finally, I would like to thank Marie for her love and support. If it weren't for her, I would never have been able to finish this thesis.

List of Symbols and Abbreviations

$\hat{.}$	variable with a hat denotes estimate of corresponding variable without hat
$\mathbf{1}_{[\cdot]}$	indicator function
a_i	i-th transmitted BPSK data bit
B	bandwith of receiver bandpass filter
b_{ij}	IEEE 802.15.4a scrambling sequence
C	length of IEEE 802.15.4a preamble code
C_{mid}	number of medium level elements in preamble code cross-correlation
C_{NZ}	number of nonzero code symbols in preamble code
C_{peak}	number of peaks in preamble code cross-correlation
$c_{\text{THS},i}$	i-th element of random THS
C_{trough}	number of troughs in preamble code cross-correlation
d_i	i-th transmitted BPPM data bit
$E_{i,j}$	received energy during the j-th code symbol of the i-th preamble symbol
E_{N}	expected noise level of preamble symbol
E_{S}	expected signal level of preamble symbol
E_p	transmitted energy per pulse
$\mathbb{E}[.]$	expectation

${}_0F_1\left(;a;x\right)$	confluent hypergeometric limit function
$F^{-1}_{\text{BIN}}(x\|n,p)$	inverse CDF of binomial distribution with parameters n and p
$f_{\chi^2}(y\|\kappa)$	PDF of chi-square random variable with κ degrees of freedom
$F^{-1}_{\chi^2}(x\|\kappa)$	inverse CDF of chi-square random variable with κ degrees of freedom
$f_{\text{GMM}}(\mathbf{v}\|\mathbf{\Theta_v})$	PDF of GMM with parameters Θ_v
$f_{\text{HMM}}(\mathbf{v}\|\mathbf{\Theta_v})$	PDF of HMM with parameters Θ_v
$f_{\mathcal{N}}(v_n\|\sigma^2)$	PDF of normal distribution with zero-mean and variance σ^2
$f_{\text{NC}\chi^2}(y\|\kappa,\zeta)$	PDF of non-central chi-square random variable with κ degrees of freedom and non-centrality parameter ζ
$F^{-1}_{\text{NC}\chi^2}(x\|\kappa,\zeta)$	inverse CDF of non-central chi-square random variable with κ degrees of freedom and non-centrality parameter ζ
$f_{\Gamma}(y\|\kappa,\theta)$	PDF of gamma distribution with scale parameter κ and shape parameter θ
$g(y_m\|\hat{q}_m,\frac{\hat{N}_0}{2})$	non-linearity to reject MUI, depends on estimated weights $\hat{q}_m$ and noise PSD through threshold η_m
$h(t)$	channel impulse response
$\tilde{h}(t)$	impulse response of compound channel
$I(x,y)$	continuous function denoting a two-dimensional image
$I_v(z)$	v-th order modified Bessel function of the first kind
L	number of multipath components
$\text{LLR}(y\|\sigma,\kappa,\zeta)$	log-likelihood ratio of non-central chi-square distribution with non-centrality parameter ζ and central chi-square distribution, both with κ degrees of freedom and assuming that noise PSD is σ
L_{s}	spreading factor in IEEE 802.15.4a preamble
M	number of samples per preamble pulse / length of channel mask

M_{NZ}	number of nonzero elements of the channel mask
M_{T}	length of correlation template
$n(t)$	noise signal
N_{c}	number of chips per frame
N_{CH}	number of preamble symbols used for channel estimation
N_{cpb}	number of chips per burst
N_{data}	number of samples of payload
N_{err}	number of erroneous bits tolerated in a nonce
N_{f}	number of samples per frame
N_{G}	number of preamble symbols in correlation template
N_{g}	number of guard chips per frame
N_{h}	number of time-hopping positions per frame
N_{nonce}	number of bits in a nonce
N_{pay}	number of data symbols per IEEE 802.15.4a packet
N_{pre}	number of preamble symbols in IEEE 802.15.4a preamble
N_{s}	number of symbols per packet
N_{sfd}	number of preamble symbols in IEEE 802.15.4a SFD
N_{sync}	number of preamble symbols in SYNC part of IEEE 802.15.4a preamble
N_{train}	number of samples in training sequence
N_u	number of users
N_{V}	number of verification steps in coarse timing acquisition
P	model order
$p(t)$	transmitted pulse

$P_{\text{AWGN}}^{\text{FA}}$	FA probability that governs the detection threshold
$P_{\text{AWGN}}^{\text{FA,PICNIC}}$	FA probability that governs the PICNIC jump-back-and-search-forward threshold $\eta_{\text{picnic}}^{\text{jump}}$
$P_{\text{AWGN}}^{\text{FA,PID}}$	FA probability that governs the PID detection threshold
$P_{\text{AWGN}}^{\text{FA,PID,fine}}$	FA probability that governs the PID fine acquisition threshold
P_{guess}	security level of secure ranging protocol
$p_{m,i}$	coefficient derived from channel energy-delay profile
q_m	weighting coefficients for optimal energy-detection demodulation / channel coefficients for coherent demodulation
$Q(\boldsymbol{\Theta}, \hat{\boldsymbol{\Theta}}')$	complete-data log-likelihood
R	arrival rate of packets at transmitter queue
$R(\rho, \theta)[I(x, y)]$	continuous Radon transform of two-dimensional image $I(x, y)$
$r(t)$	received signal
$R_{i,j}$	two-dimensional Radon matrix
$\mathcal{R}_p(t)$	autocorrelation function of pulse $p(t)$
$r_{\text{pay}}(t)$	received signal during reception of UOI payload
R_{peak}	peak data rate
$r_{\text{pre}}(t)$	received signal during reception of UOI preamble
s_i	modulation coefficient of i-th IEEE 802.15.4a preamble symbol
$s_i^{(\text{sfd})}$	ternary SFD code
T	integration time of energy-detection receiver
t_m	correlation template
T_{ack}	duration reserved for reception of ACK

T_{b}	BPPM offset
T_{back}	duration of backoff timer
T_{burst}	duration of IEEE 802.15.4a burst of pulses
T_{c}	chip duration
t_{DET}	detection time of adversarial Rake receiver
t_{ED}	early detection delay
$t_{\mathrm{ED}}^{\mathrm{SFD}}$	early SFD detection delay
T_{f}	frame duration (pulse repetition period)
T_{Fix}	duration of fixed integration window
t_{LC}	late commit delay
$t_{\mathrm{LC}}^{\mathrm{SFD}}$	late SFD commit delay
t_{OOK}	OOK detection time
T_p	pulse duration
T_{pkt}	duration of a packet transmission
T_{psym}	duration of IEEE 802.15.4a preamble symbol
t_{relay}	relay time-gain
T_{spread}	channel delay spread
U_{tx}	utilization of transmitter
$v(t)$	aggregate of noise and MUI signal
$\bar{v}$	average noise energy per sample at energy-detection receiver output
W	length of search-back window in fine timing acquisition
$w(t)$	aggregated MUI interference signal
$w_m^{(l)}$	parameter estimates needed to calculate optimal energy-detection weights q_m

$x(t)$	transmitted signal
$\tilde{x}(t)$	contribution of UOI to received signal
$\tilde{x}_{\text{pay}}(t)$	contribution of UOI payload to received signal
$\tilde{x}_{\text{pre}}(t)$	contribution of UOI preamble to received signal
$y_{m,i}$	m-th sample of the i-th frame at integrator output
$y^{\text{pre}}_{m,j+iC}$	m-th sample of the j-th code symbol of the i-th IEEE 802.15.4a preamble symbol at integrator output
$y^{\text{pre},(l)}_{m,j+iC}$	m-th sample of the j-th code symbol of the i-th IEEE 802.15.4a preamble symbol at output of l-th energy-detection receiver branch
z_m	output of discrete correlation for timing acquisition
α_l	attenuation coefficient of l-th multipath component
$\Gamma(z)$	gamma function
$\gamma_p(n)$	posterior probability of sample n being generated by mixture component p
$\delta(t)$	Dirac delta function
δ_m	Kronecker delta
$\Delta\theta$	discretization step of Radon space θ-axis
$\Delta\rho$	discretization step of Radon space ρ-axis
ϵ	relative clock offset between transmitter and receiver
ϵ_r	rate of window expansion
η_{detect}	detection threshold
η_{mask}	channel mask threshold
η_m	MUI mitigation threshold depending on q_m
η^{sfd}_m	MUI mitigation threshold during SFD detection

$\eta_{\text{pid}}^{\text{fine}}$	PID fine timing acquistion threshold
η_{picnic}	PICNIC interference detection threshold
$\eta_{\text{picnic}}^{\text{jump}}$	threshold in PICNIC jump-back-and-search-forward procdure
η_{pid}	PID detection threshold
η_{radon}	threshold to reject noise in Radon transform
$\mathbf{\Theta}$	vector of model parameters
θ	coordinate of Radon space, angle between normal vector to a line and cartesian x-axis
$\mathbf{\Theta_c}$	channel parameters
$\mathbf{\Theta_v}$	parameters of noise/interference model
λ_p	pior probability of mixture component p
ν	LOS propagation delay
ν_0	TOA of UOI packet
π_{x_0}	initial state probability of HMM
ρ	coordinate of Radon space, distance of line from origin
σ_p^2	variance of mixture component p
τ_l	delay of l-th multipath component
$\tau_{\text{pkt},k}$	start of transmission of k-th packet
ϕ	angular clock-offset corresponding to ϵ
$\psi(i)$	starting point of integration for the i-th frame
$\psi_{x_{n-1}x_n}$	state transition probability of HMM
ACK	acknowledgment packet
ACQER	acquisition error rate

AcR	autocorrelation receiver
ARake	all-Rake receiver
ARX	adversarial receiver
ATX	adversarial transmitter
BER	bit error rate
BPAM	binary pulse amplitude modulation
BPF	bandpass filter
BPPM	binary pulse position modulation
BPSK	binary phase-shift keying
CCA	clear channel assessment
CDF	cumulative distribution function
CER	capture error rate
ED	early detection
EM	expectation maximization
FA	false alarm
FCC	Federal Communications Commission
FEC	forward error correction
FIFO	first-in first-out
GMM	Gaussian mixture model
HDR	high data rate
HMM	hidden Markov model
HPRF	high pulse repetition frequency
HRX	honest receiver

HTX	honest transmitter
IR-UWB	impulse-radio ultra-wide band
ISI	inter-symbol interference
LC	late commit
LDR	low data-rate
LLR	log-likelihood ratio
LOS	line of sight
LPRF	low pulse repetition frequency
MAC	medium access control
MB-OFDM	multi-band OFDM
MD	missed detection
MRC	maximum ratio combining
MUI	multi-user interference
NLOS	non line of sight
OFDM	orthogonal frequency-division multiplexing
OOK	on-off keying
PAM	pulse amplitude modulation
PDF	probability density function
PER	packet error rate
PHR	PHY header
PHY	physical layer
PPM	pulse position modulation
PRake	partial Rake receiver

PRF	pulse repetition frequency
RBE	rapid-bit-exchange
RFID	radio frequency identification
RMSE	root mean square error
RS	Reed-Solomon
SER	synchronization error rate
SFD	start frame delimiter
SNR	signal-to-noise ratio
SRake	selective Rake receiver
SV	Saleh-Valenzuela
THS	time-hopping sequence
TOA	time of arrival
TR	transmitted reference
UOI	user of interest
WPAN	wireless personal area network

Contents

Introduction and Related Work

Chapter 1

Introduction

1.1 Motivation

In the past few years we have seen a rapid emergence of short-range wireless networks. An increasing number of electronic devices are capable of communicating via such networks. Until recently, a typical mobile phone for example could only interact with the cellular network. Today it also communicates with other devices using short-range technologies such as WiFi or Bluetooth.

As we keep equipping more and more devices and even everyday objects with wireless networking capabilities to constitute an Internet of Things, several challenges arise.

First of all, although existing technologies such as WiFi may be suitable for mobile phones, they are not suitable for networked objects of the future, such as tiny sensor nodes, because of cost, complexity and energy constraints.

Second, as wireless short-range networks grow both in size and number, it becomes impractical, if not infeasible, to coordinate and manage them in the traditional centralized and infrastructure-based fashion. Rather, we can expect to have a large number of coexisting networks that are self-organized and uncoordinated. This in turn leads to an increase in uncontrolled interference due to concurrent transmissions from uncoordinated devices.

Third, as wireless networks grow in importance, they also become more attractive targets for hackers. At the same time, more and more applications emerge that require a device to be location aware. A location aware device is able to infer its own physical location by communicating with its environment. It may then base its actions on the obtained location information. This opens up a whole new space of possible attacks that target the inherently security sensitive

location information.

In this thesis we study Impulse-radio Ultra-wide Band (IR-UWB), a physical layer radio technology that has many features that make it a promising choice for future short-range wireless networks. The defining characteristic of IR-UWB is its large bandwidth that can range from several hundred megahertz to a few gigahertz. Thanks to this large bandwidth, combined with low duty-cycle transmissions, IR-UWB should in theory be able to accommodate a large number of concurrent users while keeping interference low. A large bandwidth also results in a fine temporal resolution that can be exploited to yield accurate ranging capabilities that enable location based services. Consequently, IR-UWB has been standardized in the form of the IEEE 802.15.4a amendment to the IEEE 802.15.4 standard for low data-rate wireless personal area networks (WPAN) with extensive battery life and very low complexity.

The main goal of this thesis is to understand and show how a low-complexity IR-UWB receiver can be designed such that it is able to cope with the challenging environment it will face in future short-range wireless networks. In our analysis we take an integral system-level approach where we consider the whole process of receiving data packets, covering physical layer functions from packet detection and timing acquisition up to and including payload decoding.

1.2 Dissertation Outline

We start this thesis with an introduction to IR-UWB in Chapter 2. We then introduce the physical layer models and our main assumptions that are used throughout this thesis and present common architectures for IR-UWB receivers. Finally, Chapter 2 also contains an overview of the IEEE 802.15.4a standard. In Chapter 3 we discuss the current state-of-the-art, covering work that is related to our work presented in subsequent chapters.

The main part of this thesis is organized in three parts. Part I is devoted to multi-user interference (MUI) and starts with Chapter 4. In Chapter 4 we evaluate the performance of the IEEE 802.15.4a standard in the presence of MUI with a low-complexity receiver that is based on energy-detection. We find that such a setting is highly vulnerable to MUI and that every aspect of physical layer data packet reception is affected. More robust solutions are thus needed for synchronization as well as data reception. We address data reception in Chapter 5 and propose a novel adaptive energy-detection receiver architecture that is more robust to MUI. Since our solution requires estimation of the channel energy-delay profile, robust estimation thereof is also covered in Chapter 5. Chapter 6 covers synchronization and presents several algorithms for robust packet detection and timing acquisition with an energy-detection receiver. Further,

we address robust algorithms for detection of the start frame delimiter, a special sequence that marks the beginning of the payload in IEEE 802.15.4a packets. In Chapter 8 we shortly leave energy-detection receivers and explore whether mitigation of MUI can be more efficiently performed in coherent receivers, if interference is modeled according to a non-Gaussian process. We derive a receiver that is built on this principle and indeed mitigates interference better than existing solutions. Of the two interference models that we explore, the Gaussian mixture model has a lower complexity and still yields a good performance. We show in Chapter 8 that this model can also be applied to the energy-detection receiver that we developed in Chapter 5 and Chapter 6, further improving its robustness to MUI.

Chapter 9 constitutes Part II of this thesis. In this chapter, we evaluate the effect of clock offsets between the transmitter and the receiver on architectures based on energy-detection. We find that especially adaptive receivers, like the ones we developed in the first part of this thesis, are very sensitive to such drifting clocks and we propose potential solutions to this problem.

Part III of this thesis deals with security in IR-UWB ranging. In Chapter 10 we show that the IEEE 802.15.4a standard is vulnerable to distance-decreasing attacks on the physical layer. Such attacks compromise the security of secure ranging protocols by leading two ranging devices to believe that they are closer to each other than they actually are. Having identified and quantified this threat, we evaluate possible countermeasures and make recommendations on how to perform ranging over IEEE 802.15.4a in a secure manner.

Finally, we conclude this thesis in Chapter 11 with a summary of the main findings and a discussion of possible directions for future work.

1.3 Contributions

The following is a list of the main contributions of this thesis.

- We demonstrate that interference mitigation at the IR-UWB physical layer is possible, even with low-complexity devices. This is a powerful result since it verifies one of the key assumptions underlying the design of optimal IR-UWB MAC protocols [1], which explicitly allow concurrent transmissions.

- We provide a comprehensive performance evaluation of a complete packet based IEEE 802.15.4a system based on energy-detection and evaluate how it is affected by MUI. The resulting catalog of causes for reduced performance can serve as a basis for more interference resistant receiver and network designs.

- We derive the optimal energy-detection receiver structure for the IEEE 802.15.4a standard and show how it can be adapted to yield a low-complexity receiver that maintains a good performance in the presence of MUI.

- We present algorithms for packet detection, timing acquisition and start frame delimiter detection with energy-detection receivers in the presence of MUI. Such robust synchronization mechanisms are a crucial building block for IR-UWB networks, since they enable concurrent transmissions.

- We show that interference mitigation can be improved by construction of receivers that assume MUI to follow a statistical, non-Gaussian model. For such receivers, we identify the need to classify MUI into different types and to deal with each of them in an appropriate manner. Finally, we show that in a realistic multipath environment, MUI is best characterized by a model with memory.

- We show that contrary to common belief, energy-detection receivers may be vulnerable to clock-offsets between transmitter and receiver; and we propose an elegant solution to alleviate this problem.

- We expose the vulnerability of ranging within the IEEE 802.15.4a standard to distance-decreasing attacks on the physical communication layer. We quantify the effect of these attacks in different scenarios and make recommendations on how to improve the security of IEEE 802.15.4a ranging.

Chapter 2

System Model and Assumptions

In this chapter, we introduce the IR-UWB physical layer (PHY) models, that are used throughout this thesis, as well as their underlying assumptions. We also provide an overview of possible receiver structures for IR-UWB with an emphasis on the energy-detection architecture, which is the architecture of choice in several chapters. Finally, we introduce the IEEE 802.15.4a standard for IR-UWB low data-rate networks.

This chapter is organized as follows: In Section 2.1 we give a general introduction to IR-UWB. The classical IR-UWB PHY is introduced in Section 2.2. Possible receiver structures for IR-UWB are presented in Section 2.3. Section 2.4 introduces the IEEE 802.15.4a standard and Section 2.5 gives a set of assumptions that are made throughout this thesis when conducting performance evaluations with the IEEE 802.15.4a standard.

2.1 IR-UWB, One Flavor of UWB

Ultra-Wide Band is a wireless communication technology that uses signals occupying a very large bandwidth. According to the common definition, a signal is required to either have an absolute bandwidth exceeding 500 MHz, or a fractional bandwidth of at least 20%, in order to qualify as UWB [2]. Strictly speaking, UWB has been around since the early days of radio: the very first attempts at wireless transmissions can be considered to be of UWB nature. More recently, renewed interest into UWB sparked in 2002, when the US Federal Communications Commission (FCC) issued a regulation, authorizing the unlicensed use of UWB in the frequency range from 3.1 to 10.6 GHz [2]. Similar regulations followed in Europe [3] and Japan. What all of these regulations have in common is that they impose strict limits on the allowable

emitted power levels in order to limit interference from UWB to coexisting systems that reside in the same frequency range. The FCC ruling, e.g., specifies that the power spectral density must not exceed -41.3 dBm/MHz in the 3.1 to 10.6 GHz range. These stringent limits on the transmitted power result in a relatively short communication range of UWB systems, usually in the order of several meters to a few tens of meters.

Current UWB systems can broadly be divided into two classes: Impulse-Radio Ultra-Wide Band (IR-UWB) based systems and Multi-Band OFDM (MB-OFDM) based systems. The focus of this thesis, and therefore also of the rest of this chapter, is on IR-UWB. IR-UWB has been chosen as a possible PHY for the IEEE 802.15.4 standard [4] for low data-rate (LDR) devices with an emphasis on very low complexity and power consumption wireless personal area networks (WPAN). IR-UWB owes its name to its pulse-based transmission scheme, transmitting very short pulses with a low duty cycle. MB-OFDM, on the other hand, uses orthogonal frequency-division multiplexing (OFDM) modulation and has been standardized in two ECMA standards for high data-rate (HDR) wireless personal area networks (WPAN)[5, 6]. It targets applications with data rates of several hundred Mb/s over a few meters and is, e.g., the technology that is underlying Wireless USB [7] that is meant to replace USB cables with a wireless link. MB-OFDM is not considered in this thesis.

2.2 The Classical IR-UWB Physical Layer

The classical IR-UWB PHY structure is based on the work by Win and Scholtz [8, 9, 10]. The main idea is to send a train of short pulses, each pulse having a width T_p in the order of a nanosecond or less, resulting in an UWB signal. Consecutive pulses have a separation much larger than T_p. Consequently, the resulting transmission has a low duty cycle, typically around only 1%. This in turn allows multiple users to access the wireless medium at the same time and to transmit concurrently. Still, with a uniform pulse separation, it may happen that a large number of pulses from signals of two users overlap. To avoid such *catastrophic collisions* the separation between pulses is randomized through a pseudo-random *time-hopping sequence* (THS).

2.2.1 Transmitted IR-UWB Signal

We assume that PHY transmissions occur in packets of N_s symbols. The transmitted signal corresponding to the transmission of a single packet is then given by

$$x(t) = \sum_{i=0}^{N_\mathrm{s}-1} a_i \cdot p(t - iT_\mathrm{f} - c_{\mathrm{THS},i}T_\mathrm{c} - d_iT_\mathrm{b}). \tag{2.1}$$

Here $p(t)$ denotes the waveform of the transmitted pulse with width T_p. Every data symbol is formed by one pulse. The main time unit is a *chip* of duration T_c and generally, $T_\mathrm{c} \geq T_p$. N_c consecutive chips form a *frame* of duration $T_\mathrm{f} = N_\mathrm{c}T_\mathrm{c}$. One symbol is sent per frame, thus resulting in a duty cycle of $T_\mathrm{p}/T_\mathrm{f}$. Within a frame, each pulse is shifted by a certain number of chips according to the THS. For the i-th pulse this offset is $c_{\mathrm{THS},i}T_\mathrm{c}$. The THS $\{c_{\mathrm{THS},i}\}$ is a pseudo-random sequence, whose elements are integers from the set $[0, N_\mathrm{h}-1]$. N_h thus denotes the number of chips that are available for time-hopping. N_h is usually smaller than N_c, leaving a *guard interval* of $N_\mathrm{g} = N_\mathrm{c} - N_\mathrm{h}$ chips to prevent inter-symbol interference (ISI).

To transmit information, two modulation formats can be used: *pulse amplitude modulation* (PAM) and *pulse position modulation* (PPM). In this thesis we restrict ourselves to binary PAM (BPAM) and binary PPM (BPPM) because they are more commonly used than their M-ary counterparts. With BPAM, the i-th information bit to be sent is conveyed by the amplitude a_i of the pulse (and consequently, $d_i = 0$). The pulse is sent with a different amplitude depending on whether a 0-bit or a 1-bit is being sent. The two most common forms of BPAM are binary phase-shift keying (BPSK) and on-off keying (OOK). With BPSK, 0-bits and 1-bits have opposite amplitudes, i.e.,

$$a_i \in \{-1, 1\}$$

With OOK, a 1-bit is signalled by the presence of the pulse, a 0-bit by its absence, i.e.,

$$a_i \in \{0, 1\}$$

With BPPM, the i-th information bit to be sent is conveyed by the position of the pulse through the offset d_iT_b (and consequently, $a_i = 1$). Usually, a 0-bit is signalled by sending the pulse with no offset and a 1-bit by sending the pulse with an offset of half the frame duration, i.e.,

$$d_i \in \{0, 1\}, \quad T_\mathrm{b} = T_\mathrm{f}/2$$

The frame is then divided into a 0-*block* and a 1-*block* . To signal a 0-bit, a pulse is sent in the 0-block, to signal a 1-bit, a pulse is sent in the 1-block.

BPSK and BPPM can also be used in combination, in which case every symbol carries two information bits, one conveyed by the amplitude of the pulse and one by the position of the pulse.

2.2.2 Multipath Channel

The transmitted signal $x(t)$ propagates through the wireless medium to reach the receiver. The received signal is given by

$$r(t) = x(t) * h(t - \nu) + n(t), \tag{2.2}$$

where $n(t)$ is zero-mean additive white Gaussian noise (AWGN) with two-sided power spectral density $N_0/2$, ν denotes the line-of-sight (LOS) propagation delay, $h(t)$ is the impulse response of the wireless channel and $*$ denotes convolution.

One fundamental concept here is *multipath propagation*, the fact that a signal reaches the receiver via different paths due to reflections by surrounding objects. This effect is especially pronounced indoors [11] resulting in a large number of replicas of the transmitted signal, called *multipath components*, to be received at the receiver. Depending on its propagation path, each multipath component arrives with a certain delay, attenuation and phase. In narrowband systems, multipath propagation leads to multipath fading if overlapping multipath components combine in a destructive way. Multipath fading can severely degrade the quality of the received signal. With UWB signals this effect is less pronounced because the multipath components remain resolvable due to the short pulse duration.

A widely used model for the impulse response of the UWB multipath channel $h(t)$ is the *tapped-delay-line model*, under which

$$h(t) = \sum_{l=0}^{L-1} \alpha_l \delta(t - \tau_l). \tag{2.3}$$

L is the number of multipath components corresponding to the different propagation paths. α_l and τ_l are the attenuation coefficient and the delay induced by the l-th multipath component, respectively. δ denotes the Dirac delta function. A lot of research has been devoted to measuring the UWB channel and building models for τ_l and α_l that characterize it well. For a recent survey on this research topic, the reader is referred to [12]. A popular model is the Saleh-Valenzuela (SV) model [13] that assumes that multipath components form clusters whose arrival times are

Poisson distributed. The SV model has also been adopted in the standardized IEEE 802.15.4a channel models [14, 15]. Throughout this thesis, the IEEE 802.15.4a channel models are used for simulations and performance evaluation.

The model in (2.3) is a relatively simple model that ignores frequency and path dependent distortions of the pulse. On the other hand, such effects mainly affect UWB signals with a large relative bandwidth [12]. UWB signals like the IEEE 802.15.4a signals considered in this thesis have a large absolute bandwidth only and are less affected by these effects.

Further, α_l and τ_l usually change over time because of variations of the environment. The environment varies due to, e.g., the motion of people. Still, throughout this thesis we assume the channel impulse response to be invariant for the duration of one packet. This is motivated by the rate of such changes in the environment that is very slow compared to the packet duration in the order of a millisecond.

2.2.3 Multiple Access and Multi-User Interference

A large part of this thesis is devoted to *multi-user interference* (MUI). By MUI we understand concurrent transmissions from other IR-UWB users that overlap in time with the transmission of the *user of interest* (UOI). We only consider MUI from users that also use an IR-UWB PHY. In particular, we are not concerned with coexistence of IR-UWB and narrowband systems; i.e., with interference from narrowband systems on IR-UWB, or vice versa. Coexistence is a research topics on its own and out of the scope of this thesis. For a recent survey on UWB coexistence, the reader is referred to [16] and the references therein.

Concurrent transmissions may stem from coexisting IR-UWB networks that are operating near-by. They may also originate from within the same network if the medium access control (MAC) protocol allows concurrent transmissions like it is, e.g., the case with IEEE 802.15.4a (Section 2.4.4). We further assume in general that concurrent transmissions are uncontrolled, i.e., that there is no power control mechanism preventing the *near-far problem* and that interfering signals may thus be received with a power exceeding the power of the UOI.

With N_u concurrent users present in the system, the signal at the receiver becomes a superposition of the received signals from the individual users and of the additive noise term:

$$r(t) = \sum_{u=0}^{N_u-1} \sum_{k} x_k^{(u)}(t) * h_k^{(u)}(t - \nu_k^{(u)}) + n(t). \tag{2.4}$$

Here, $x_k^{(u)}(t)$ denotes the transmitted signal corresponding to the k-th packet of the u-th user

that has a propagation delay $\nu_k^{(u)}$ and that is given by

$$x_k^{(u)}(t) = \sum_{i=0}^{N_\mathrm{s}-1} a_{k,i}^{(u)} \cdot p(t - iT_\mathrm{f} - c_{\mathrm{THS},i}^{(u)} T_\mathrm{c} - b_{k,i}^{(u)} T_\mathrm{f}/2 - \tau_{\mathrm{pkt},k}^{(u)}) \tag{2.5}$$

Apart from $\tau_{\mathrm{pkt},k}^{(u)}$, denoting the time instant at which the transmission of the k-th packet by the u-th user starts, this still corresponds to (2.1). The only difference is that we introduced appropriate sub- and superscripts for packet and user dependent variables.

The same goes for the channel impulse response $h_k^{(u)}(t)$ that differs between users and packets because we assume time invariance only for the duration of one packet (see Section 2.2.2) and because it depends on the users' physical positions:

$$h_k^{(u)}(t) = \sum_{l=0}^{L_k^{(u)}-1} \alpha_{k,l}^{(u)} \delta(t - \tau_{k,l}^{(u)}). \tag{2.6}$$

Without loss of generality, we assume that the UOI is the user with index $u = 0$. The explicit description of the full received signal model given by (2.4)-(2.6) is quite cumbersome. For most parts of this thesis, however, we can work with a simpler representation. We can make a first simplification by only considering the reception of a single packet of the UOI, e.g., the one with index $k = 0$. By further denoting the contribution of the UOI to the received signal as $\tilde{x}(t)$, given by

$$\tilde{x}(t) = x_0^{(0)}(t) * h_0^{(0)}(t - \nu_0^{(0)}), \tag{2.7}$$

and by additionally aggregating all MUI terms as

$$w(t) = \sum_{u=1}^{N_u-1} \sum_{k} x_k^{(u)}(t) * h_k^{(u)}(t - \nu_k^{(u)}), \tag{2.8}$$

we can get rid of the indices u and k altogether, resulting in the much simpler received signal model given by

$$\begin{aligned} r(t) &= \tilde{x}(t - \nu_0) + v(t) \\ &= \underbrace{\sum_{l=0}^{L-1} \alpha_l \sum_{i=0}^{N_\mathrm{s}-1} a_i \cdot p(t - iT_\mathrm{f} - c_{\mathrm{THS},i} T_\mathrm{c} - d_i T_\mathrm{f}/2 - \tau_l - \nu_0)}_{\tilde{x}(t-\nu_0)} + \underbrace{w(t) + n(t)}_{v(t)} \end{aligned} \tag{2.9}$$

where $v(t) = w(t) + n(t)$ denotes the aggregate of MUI and AWGN. Further, ν_0 denotes the time of arrival (TOA) of the UOI packet and is given by

$$\nu_0 = \nu_0^{(0)} + \tau_{\text{pkt},0}^{(0)}. \tag{2.10}$$

For convenience we also introduce the *compound channel impulse response* that combines the effect of the multipath channel and of the transmitted pulse

$$\tilde{h}(t) = p(t) * h(t). \tag{2.11}$$

This allows us to further simplify the expression of the received signal, which is now equal to

$$r(t) = \underbrace{\sum_{i=0}^{N_\text{s}-1} a_i \cdot \tilde{h}(t - iT_\text{f} - c_{\text{THS},i}T_\text{c} - d_iT_\text{f}/2 - \nu_0)}_{\tilde{x}(t-\nu_0)} + v(t). \tag{2.12}$$

2.3 Receiver Structures for IR-UWB

Receiver structures for IR-UWB can broadly be classified into correlation or *Rake receivers* and *autocorrelation receivers* (AcR). Rake receivers cross-correlate the received signal with a locally generated template that is matched to the received signal, while AcR receivers perform the correlation with a delayed version of the received signal. A special case of an AcR is the *energy-detection receiver* which performs a squaring operation instead of the autocorrelation. Rake receivers usually have a better performance but also a significantly higher complexity than AcR or energy-detection receivers.

2.3.1 Rake Receivers

The receiver with optimal single user performance, but also with the highest complexity, is the all-Rake (ARake) receiver using maximum ratio combining (MRC). The ARake needs to estimate the channel parameters as well as the TOA. We denote these estimates by $\hat{\alpha}_l$, $\hat{\tau}_l$ and $\hat{\nu}_0$. The ARake first correlates the received signal $r(t)$ with a template, $g(t)$, matched to the received pulse shape. We will here assume for simplicity that $r(t)$ is given by (2.9) with $v(t) = 0$, $N_\text{s} = 1$ and $c_0 = 0$; i.e., we neglect noise and interference by assuming $r(t) = \tilde{x}(t)$, consider detection of a single symbol only and assume no time-hopping. We further assume no distortion of the

pulse, such that $g(t) = p(t)$ [1]. This yields the following signal at the output of the correlator

$$y(t) = a_0 \sum_{l=0}^{L-1} \alpha_l \mathcal{R}_p(t - b_0 T_\mathrm{f}/2 - \tau_l - \nu_0), \tag{2.13}$$

where

$$\mathcal{R}_p(t) = \int_{-\infty}^{\infty} p(\tau) p(\tau - t)\, \mathrm{d}\tau \tag{2.14}$$

is the autocorrelation of the pulse waveform $p(t)$.

Two decision variables, $Y_{\hat{d}_0}$, are then formed, one for each of the hypotheses, $\hat{d}_0 = 0$ and $\hat{d}_0 = 1$. For each of these decision variables, the energy from the different multipath components is combined by sampling $y(t)$ L times, at time instants $t_{m,\hat{d}_0} = \hat{d}_0 T_\mathrm{f}/2 + \hat{\tau}_m + \hat{\nu}_0$, and by combining these samples according to

$$\begin{aligned} Y_{\hat{d}_0} &= \sum_{m=0}^{L-1} \hat{\alpha}_m \cdot y(t_{m,\hat{d}_0}) \\ &= a_0 \sum_{m=0}^{L-1} \hat{\alpha}_m \sum_{l=0}^{L-1} \alpha_l \mathcal{R}_p((\hat{d}_0 - b_0) T_\mathrm{f}/2 + \hat{\tau}_m - \tau_l + \hat{\nu}_0 - \nu_0) \end{aligned} \tag{2.15}$$

If all of the multipath components are resolvable (i.e., if $|\tau_l - \tau_m| \geq T_p, \forall l \neq m$) and if perfect channel state information is available (i.e., $\hat{\alpha}_m = \alpha_l$, $\hat{\tau}_m = \tau_l$ and $\hat{\nu}_0 = \nu_0$), then the ARake is able to collect the entire energy of the received signal, corresponding to $E_p \sum_{l=0}^{L-1} \alpha_l^2$, where $E_p = \mathcal{R}_p(0)$ is the transmitted energy per pulse.

Unfortunately, the optimal ARake receiver is hardly feasible in practice because of its high complexity. First of all, the number of resolvable multipath components that have to be accurately estimated and combined is potentially huge in UWB channels [17, 18]. Several sub-optimal receivers like the selective Rake (SRake) [19] or partial Rake (PRake) [20] try to circumvent this problem by only estimating a small subset of multipath components. Further, sub-optimal combining schemes like non-coherent square-law combining [21] remove the necessity to estimate α_l altogether. However, all of those receivers still share the stringent timing requirements of the ARake: We can see from (2.15) that even a slight timing offset of $\hat{\tau}_m$ or $\hat{\nu}_0$ can result in a huge energy loss because of the wide bandwidth of the UWB pulse. Rake receivers thus need highly accurate synchronization, clocks and channel delay estimation to prevent a performance loss due to timing impairments [22, 23, 24, 25]. All of this is only

1. In practice the received pulse is a distorted version of the transmitted pulse. This distortion either entails a lowered performance, if it is neglected, or an increased complexity, if it is corrected for.

possible with high sampling frequencies, thus increasing the complexity of the receiver.

2.3.2 Autocorrelation Receivers

Autocorrelation receivers are a low-complexity alternative to Rake receivers. Compared to Rake receivers, they do not only have lower sampling frequency requirements, but they also do not need any knowledge about the shape of the received pulse, nor do they need a down conversion stage to convert the signal from passband to baseband. Finally, they can operate with a simplified channel estimation or even completely omit channel estimation altogether.

AcRs correlate the received signal with a delayed version of itself. We distinguish transmitted reference (TR) receivers [26] where this delay is nonzero and energy-detection receivers where the received signal is simply squared instead of being correlated. [2]

In this thesis we only consider energy-detection receivers which are explained in the next section. The main reason for not considering TR receivers is that TR receivers require a specially designed transmitted signal: Every pulse that carries information is preceded by a reference pulses that serves as correlation template. This requirement on the form of the transmitted signal limits the applicability of TR receivers. In particular, it makes them incompatible with the IEEE 802.15.4a standard that does not foresee reference pulses. Further, the pulse and its reference are usually separated by a delay in the order of several tens of nanoseconds to account for the channel spread. Consequently, an analog delay line of the same length is required on the receiver side. Low-cost analog delay lines of such a length are in general difficult to realize [27].

2.3.3 Energy-detection Receivers

Unlike TR receivers, energy-detection receivers [28, 29] have no requirement for a specially designed transmitted signal. Further, they are explicitly supported by the IEEE 802.15.4a standard via the non-coherent BPPM modulation format. The decision to support non-coherent receivers in IEEE 802.15.4a paired with the simplicity of energy-detection receivers makes them a popular choice for low-complexity applications. The basic architecture of an energy-detection receiver is shown in Figure 2.1. After the antenna, the received signal is filtered with a bandpass filter (BPF) of bandwidth B to limit the effect of noise. Throughout this thesis we assume a perfect bandpass filter with a bandwidth matching the bandwidth of the transmitted

2. Squaring is equivalent to correlation at lag zero, so we can see energy-detection receivers as a special case of AcRs

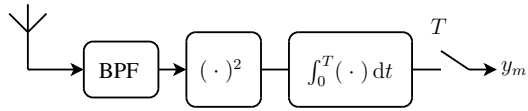

Figure 2.1: Energy-detection receiver architecture.

pulse $p(t)$. After filtering, the signal is squared in a squaring device and the result is integrated, yielding discrete samples. The duration of the integration, T, is an important system parameter and called *integration time*.

For simplicity we assume that T divides T_f and that a total of

$$N_\mathrm{f} = T_\mathrm{f}/T \tag{2.16}$$

samples are taken per frame. The discrete signal at the output of the integrator is then given by

$$y_{m,i} = \int_{mT+\psi(i)}^{(m+1)T+\psi(i)} [r(t)]^2\,\mathrm{d}t, \tag{2.17}$$

where $y_{m,i}$ denotes the m-th sample of the i-th frame, $r(t)$ is given by (2.12) and the starting point, $\psi(i)$, of the integration for the i-th frame corresponds to

$$\psi(i) = iT_\mathrm{f} + c_{\mathrm{THS},i}T_\mathrm{c} + \nu_0. \tag{2.18}$$

In its simplest form, the energy-detection receiver decides on what BPPM symbol was transmitted by comparing the received energy in the first and the second half of each BPPM frame. For the i-th frame, the two decision variables associated with the two hypotheses, $\hat{d}_0 = 0$ and $\hat{d}_0 = 1$, are then given by

$$Y_{i,\hat{d}_0} = \sum_{m=\hat{d}_0\cdot\frac{N_\mathrm{f}}{2}}^{\frac{N_\mathrm{f}}{2}(\hat{d}_0+1)-1} y_{m,i} = \int_{\hat{d}_0\cdot T_\mathrm{f}/2+\psi(i)}^{(\hat{d}_0+1)T_\mathrm{f}/2+\psi(i)} [r(t)]^2\,\mathrm{d}t \tag{2.19}$$

This receiver thus takes two samples per frame and has an integration time of $T = T_\mathrm{f}/2$. It then decides according to which part of the frame contains more energy

$$Y_{i,0} \underset{\hat{d}_i=1}{\overset{\hat{d}_i=0}{\gtrless}} Y_{i,1} \tag{2.20}$$

It can be shown that this simple receiver is the optimal receiver if no channel state information is available [3] [30, 31]. On the other hand, such a receiver also integrates a lot of noise, especially if T_f is large compared to the delay spread of the channel. In practice, it is therefore desirable to estimate the delay spread of the channel and to adapt the integration time appropriately [32].

Noise can be further limited and performance can be increased if partial channel state information in the form of the channel energy-delay profile is available. In this case, the optimal receiver estimates the channel energy-delay profile and weights the received signal accordingly; either in the analog domain before the integration [31], or in the digital domain after the integration [33].

Distribution of Energy-detection Receiver Output in Single User Case

If $r(t)$ does not contain MUI, i.e., $w(t) = 0$ in (2.12), the probability density function (PDF) of the receiver output $y_{m,i}$, given in (2.17), can be closely approximated [28] with

$$f(y_{m,i}|N_0/2, BT, \zeta_{m,i}) \approx \frac{1}{N_0/2} f_{\mathrm{NC}\chi^2}\Big(\frac{y_{m,i}}{N_0/2}\Big|2BT, \zeta_{m,i}\Big), \tag{2.21}$$

where

$$f_{\mathrm{NC}\chi^2}(y|\kappa, \zeta) = \frac{1}{2} e^{-\frac{(y+\zeta)}{2}} \Big(\frac{y}{\zeta}\Big)^{\frac{\kappa/2-1}{2}} I_{\kappa/2-1}(\sqrt{\zeta y}) \tag{2.22}$$

is the PDF of a non-central chi-square random variable with κ degrees of freedom and non-centrality parameter ζ. $I_\upsilon(z)$ is the υ-th order modified Bessel function of the first kind. In addition, the samples $y_{m,i}$ are assumed to be independent. This is a practical assumption, also made in the related work [33].

The non-centrality parameters $\zeta_{m,i}$ are given by

$$\zeta_{m,i} = \frac{p_{m,i}}{N_0/2} \quad \text{with} \quad p_{m,i} = \int_{mT+\psi(i)}^{(m+1)T+\psi(i)} [\tilde{x}(t)]^2 \, \mathrm{d}t. \tag{2.23}$$

The coefficients $p_{m,i}$ thus depend on the channel energy-delay profile of the UOI signal.

Accordingly, if the received signal consists purely of AWGN, i.e., $r(t) = n(t)$, we approximate the PDF of the receiver output $y_{m,i}$ with the commonly made approximation of

3. Or more precisely, if the only channel state information available is that the channel spread is smaller than $T_\mathrm{f}/2$.

independent chi-square random variables

$$f(y_{m,i}|N_0/2, BT) \approx \frac{1}{N_0/2} f_{\chi^2}\Big(\frac{y_{m,i}}{N_0/2}\Big|2BT\Big), \tag{2.24}$$

where

$$f_{\chi^2}(y|\kappa) = \frac{1}{2^{\kappa/2}\Gamma(\kappa/2)} y^{\kappa/2-1} e^{-y/2} \tag{2.25}$$

is the PDF of a chi-square random variable with κ degrees of freedom. The scaled chi-square distribution in (2.24) is equivalent to a gamma distribution. We can therefore alternatively write

$$f(y_{m,i}|N_0/2, BT) \approx f_\Gamma(y_{m,i}|BT, N_0), \tag{2.26}$$

where

$$f_\Gamma(y|\kappa, \theta) = \frac{1}{\theta^\kappa \Gamma(\kappa)} y^{\kappa-1} e^{-y/\theta} \tag{2.27}$$

is the PDF of the gamma distribution with scale parameter κ and shape parameter θ and $\Gamma(z)$ denotes the gamma function.

2.4 The IEEE 802.15.4a Standard

IEEE 802.15.4 [4] is a standard for low data rate WPAN. The main emphasis of IEEE 802.15.4 is on networks of low-cost and low-complexity devices with extensive battery life. IEEE 802.15.4 specifies the PHY and MAC layers for such networks. Other specifications, such as ZigBee [34], exist for the higher communication layers. Several PHYs can be used with IEEE 802.15.4, one of which is based on IR-UWB and specified in the IEEE 802.15.4a amendment [35] to the IEEE 802.15.4 standard.

In the following, we will introduce the main features of IEEE 802.15.4a that are used throughout this thesis.

2.4.1 Basic Concepts and Main Differences With Classical PHY

To date, IEEE 802.15.4a is the only standard for IR-UWB networks, and therefore, the IEEE 802.15.4a PHY is also our PHY of choice and predominantly used throughout this thesis. Still, the IEEE 802.15.4a PHY is sufficiently close to the classical PHY, introduced in Section 2.2, such that most of the concepts presented in this thesis are easily transferable and applicable to the classical PHY as well. Indeed, both PHYs share key concepts such as time-hopping to

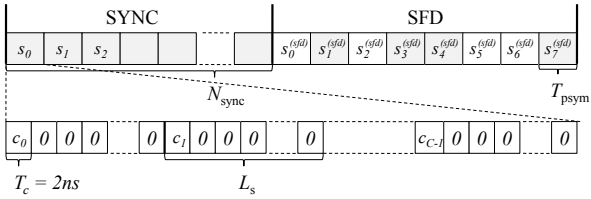

Figure 2.2: IEEE 802.15.4a preamble structure.

smooth the spectrum of the signal and to mitigate the possible impact of MUI, or the possibility for BPPM and/or BPSK modulation. The main difference of the two PHYs lies in the signalling format of the data frames. With IEEE 802.15.4a, instead of sending a single pulse per frame, a short, continuous *burst of pulses* with pseudo-random polarity is sent.

An IEEE 802.15.4a packet consists of a *preamble* followed by a *payload*. The preamble is a sequence derived from a *preamble code* that is known to the receiver. It is used for packet detection, timing acquisition and channel estimation. The payload carries the actual information bits to be transmitted. Time-hopping is used in the payload but not in the preamble. However, the preamble code used to construct the preamble sequence also defines the THS that is to be used during the payload transmission.

IEEE 802.15.4a operates in a number of 500 MHz channels in the frequency range from approximately 3 GHz to 10 GHz. Additionally, four channels with a higher bandwidth of $1-1.3$ GHz are also available in the same frequency range. To each channel, a set of preamble codes is allocated and every pair of a channel and a preamble code forms a logical channel. A standard compliant device needs to support two logical channels within one of the 500 MHz physical channels.

2.4.2 Preamble Structure

The structure of the IEEE 802.15.4a preamble is shown in Figure 2.2. The preamble comprises two parts: the SYNC part used for timing acquisition and channel estimation, followed by the start frame delimiter (SFD) that indicates the beginning of the payload.

In line with the classical IR-UWB PHY (Section 2.2), the basic time unit is a chip of duration T_{c}. During the preamble, pulses are sent at regular time intervals, every L_{s}-th chip. We make the standard assumption that *the spreading factor* L_{s} is large enough to prevent inter-pulse interference. Blocks of C consecutive pulses form a *preamble symbol*. Every pulse within a

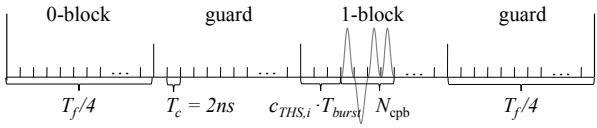

Figure 2.3: IEEE 802.15.4a payload symbol structure.

preamble symbol is modulated according to the ternary preamble code of length C. With the mandatory modes, there are two distinct preamble codes, forming two logical channels, allocated to each of the physical 500 MHz channels. The preamble codes have a perfect periodic auto-correlation and good cross-correlation properties, identical to the codes proposed in [36].

The repetition of N_{sync} preamble symbols builds the SYNC part of the preamble. The SFD is obtained by spreading a ternary *SFD code* of length N_{sfd} with a preamble symbol. The total length of the preamble thus corresponds to $N_{\text{pre}} = N_{\text{sync}} + N_{\text{sfd}}$ preamble symbols.

This results in the following received signal during the preamble:

$$r_{\text{pre}}(t) = \underbrace{\sum_{i=0}^{N_{\text{pre}}-1} s_i \sum_{j=0}^{C-1} c_j \cdot \tilde{h}(t-(j+iC)L_{\text{s}}T_{\text{c}}-\nu_0)}_{\tilde{x}_{\text{pre}}(t-\nu_0)} + v(t), \tag{2.28}$$

where $\tilde{x}_{\text{pre}}(t)$ is the contribution of the UOI and $v(t)$ accounts for noise and MUI. $c_j \in \{-1, 0, +1\}, j \in \{0, 1, \ldots, C-1\}$ is j-th *code symbol* of the ternary preamble code. Further, every preamble symbol is modulated by s_i, which is given by

$$s_i = \begin{cases} 1 & \text{if } i \in \{0, 1, \ldots, N_{\text{sync}} - 1\} \\ s^{(\text{sfd})}_{i-N_{\text{sync}}} & \text{if } i \in \{N_{\text{sync}}, \ldots, N_{\text{pre}} - 1\} \end{cases}, \tag{2.29}$$

where $s_i^{(\text{sfd})} \in \{-1, 0, +1\}, i \in \{0, 1, \ldots, N_{\text{sfd}} - 1\}$ is the ternary SFD code.

2.4.3 Payload Structure

The structure of an IEEE 802.15.4a payload symbol is shown in Figure 2.3. In contrast to the preamble signal, each symbol of an IEEE 802.15.4a payload is composed of a short, continuous burst of N_{cpb} pulses with pseudo-random polarity and time-hopping. As explained in Section 2.4.1, the THS depends on the preamble code that is used in the preamble. The same holds for the scrambling sequence that defines the pseudo-random polarity of the pulses; it also

is specific to the chosen preamble code. Further, time-hopping is performed on entire bursts. The time-hopping offset is thus a multiple of the duration of a burst $T_{\text{burst}} = N_{\text{cpb}} T_{\text{c}}$. The number of available hopping positions is restricted to $N_{\text{h}} = \frac{T_{\text{f}}/2}{2T_{\text{burst}}}$, half of the frame serves as guard interval against ISI.

Before modulation, the data is encoded using a mandatory $(55, 63)$ systematic Reed-Solomon (RS) code. Additional forward error correction (FEC) is available through a systematic rate $1/2$ convolutional code that is applied to the RS encoded data. The resulting encoded data bits are modulated using a combination of BPPM and BPSK. Information bits are modulated using BPPM which can be demodulated by both receiver types. Parity bits from the convolutional code are modulated using BPSK and a coherent receiver can thus take advantage of this additional information to increase robustness against noise.

The received signal during the payload is given by

$$r_{\text{pay}}(t) = \underbrace{\sum_{i=0}^{N_{\text{pay}}-1} a_i \sum_{j=0}^{N_{\text{cpb}}-1} b_{ij} \cdot \tilde{h}(t - iT_{\text{f}} - d_i T_{\text{f}}/2 - c_{\text{THS},i} T_{\text{burst}} - jT_{\text{c}} - \nu_0)}_{\tilde{x}_{\text{pay}}(t-\nu_0)} + v(t), \qquad (2.30)$$

where $\tilde{x}_{\text{pay}}(t)$ is the contribution of the UOI and $v(t)$ accounts for noise and MUI, N_{pay} is the number of symbols in the payload and T_{f} is the duration of a frame. As with the classical PHY, one symbol is sent per frame. $a_i \in \{-1, 1\}$ and $d_i \in \{0, 1\}$ denote the i-th BPSK and BPPM modulated data bits of the payload, respectively. $b_{ij} \in \pm\{1\}$ is the pseudo-random polarity of the j-th pulse of the i-th symbol specified by the scrambling sequence. Finally, $c_{\text{THS},i} \in [0, N_{\text{h}} - 1]$ denotes the THS.

The first 19 symbols of the payload are protected by a special single error correct double error detect (SECDED) code and constitute the PHY header (PHR). The PHR contains information about, e.g., the length of the packet. For simplicity, we ignore the PHR, thus treating it like it were part of the regular payload.

2.4.4 Medium Access Control

Medium access to the IR-UWB PHY of IEEE 802.15.4a is governed through a set of *clear channel assessment* (CCA) modes. The only mandatory mode is CCA mode 4 corresponding to pure ALOHA. With pure ALOHA, before sending a packet, the transmitter has to wait until a backoff timer of random duration expires. If after sending the packet, the transmitter does not receive an acknowledgement packet (ACK) from the intended receiver within a certain time

interval, the packet is retransmitted.

Two more sophisticated optional modes, CCA mode 5 and 6 exist, where the transmitter tries to determine whether the channel is idle before sending a packet. The transmitter decides on an idle or busy channel depending on whether it is able to detect the presence of a preamble. CCA mode 5 uses the standard packet structure introduced in Sections 2.4.2 and 2.4.3. With CCA mode 6, the payload is modified by multiplexing additional preamble symbols into the payload part of a packet.

2.5 Assumptions Made in Performance Evaluations

This section summarizes a set of assumptions that is common to most of the performance evaluations made in this thesis.

2.5.1 Mandatory Modes of the Standard

The IEEE 802.15.4a standard is very flexible in the sense that it allows for many different configurations that can be tailored to the application at hand. However, only few of the resulting parameter combinations need to be implemented by a device in order to be standard compliant. We argue that the majority of devices are likely to only implement the resulting mandatory modes and we therefore only consider these in our performance evaluations. In what follows, we introduce the mandatory modes considered and the associated systems parameters.

Even though the supported data rates in IEEE 802.15.4a show a great diversity, covering the wide range from 0.11 Mb/s up to 27 Mb/s, only a data rate of 0.85 Mb/s is mandatory. With this mandatory data rate, two transmission modes with different packet structure have to be supported. These two mandatory modes distinguish themselves in their mean *pulse repetition frequency* (PRF). The first one sends pulses at a low mean PRF of 3.9 MHz, the second one at a high mean PRF of 15.6 MHz. We will therefore denote these two modes as low PRF (LPRF) and high PRF (HPRF) modes, respectively. The HPRF mode, with a shorter spacing between successive pulses, is intended for coherent receivers and LOS channels where inter-pulse interference is less of an issue. The LPRF mode, on the other hand, is intended for non-coherent receivers and channels with a large delay spread.

Each of these two mandatory modes entails a different set of parameters that are summarized in Table 2.1. The most important differences between LPRF and HPRF are the different spreading factors used in the preamble and the different number of chips per burst in the pay-

load. Further, we simulate the 500 MHz physical channel number 3. The two preamble codes, IEEE 802.15.4a preamble codes 5 and 6, that are allocated to this channel are shown in Table 2.2, together with the mandatory SFD code sequence that is common to all channels.

2.5.2 Packet Generation at the Transmitter

In all our simulations we assume that packets at the transmitter are generated according to the same process that ultimately determines $\tau_{\text{pkt},k}^{(u)}$, the start of the PHY transmission of the k-th packet of the u-th user.

We assume that every transmitter has a first-in first-out (FIFO) queue of packets to be transmitted. Packets are generated by the upper layers according to a Poisson process with rate R packets/s and placed in the queue. For the packet at the front of the queue we start a backoff timer according to, e.g., the IEEE 802.15.4a backoff procedure. Once the timer expires, after a time T_{back}, the transmission of the packet on the channel starts. The transmission ends after a duration of T_{pkt}. The transmitter then waits for an additional duration T_{ack}, before starting the backoff timer for the next packet in the queue. This is done in order to leave room for the reception of an acknowledgement packet (ACK). For complexity reasons, we cannot simulate the full MAC protocol because this implies to simulate the reception and decoding of every single packet from *any* user at its destination, as well as the transmission and reception of the ACKs. We therefore neither simulate ACKs nor retransmission.

The utilization of the queuing system at the transmitter is

$$U_{\text{tx}} = R * (\mathbb{E}[T_{\text{back}}] + T_{\text{pkt}} + T_{\text{ack}}) \tag{2.31}$$

and the resulting *peak data rate* corresponds to

$$R_{\text{peak}} = 1/(\mathbb{E}[T_{\text{back}}] + T_{\text{pkt}} + T_{\text{ack}}) \tag{2.32}$$

With IEEE 802.15.4a, T_{ack} is assumed to correspond to the length of the preamble. For the backoff timer we assume the maximum allowable backoff exponent $BE = 8$. According to the IEEE 802.15.4a specification, T_{back} is then drawn uniformly at random from $[0, T_{\text{back}}^{\text{max}}]$ where

$$T_{\text{back}}^{\text{max}} = (2^{BE} - 1) \cdot 20 \cdot T_{\text{f}}. \tag{2.33}$$

The factor $20 \cdot T_{\text{f}}$ is given by the standard. The resulting peak data rates, assuming packets of the maximum allowable packet size for IEEE 802.15.4a of 1016 information bits (corresponding

Parameter	Description	LPRF	HPRF
	pulse repetition frequency	3.9 MHz	15.6 MHz
	transmitter duty cycle	$\approx 0.8\%$	$\approx 3\%$
	raw data rate	0.85 Mb/s	0.85 Mb/s
	channel number according to [35]	3	3
	indices of associated preamble codes [35]	5, 6	5, 6
B	bandwidth of channel and bandpass filter	500 MHz	500 MHz
C	length of preamble code	31	31
L_{s}	preamble spreading factor	64	16
N_{cpb}	number of chips per burst	4	16
N_{h}	number of time-hopping positions	32	8
N_{pre}	number of symbols in preamble	72	72
N_{sync}	number of symbols in SYNC part of preamble	64	64
N_{sfd}	length of SFD in preamble symbols	8	8
$N_{\text{pay}}^{\text{max}}$	maximum number of payload symbols	1208	1208
	corresp. number of information bits per packet	1016	1016
T_{c}	duration of a chip	2 ns	2 ns
T_{f}	duration of a frame	1024 ns	1024 ns
T_{psym}	duration of preamble symbol	3968 ns	992 ns
T_{sync}	duration of SYNC part	254 μs	63.5 μs
T_{sfd}	duration of SFD	31.8 μs	7.9 μs
T_{pre}	duration of preamble	285.8 μs	71.4 μs
R_{peak}	peak data rate with $N_{\text{pay}} = N_{\text{pay}}^{\text{max}}$(Section 2.5.2)	226 pkts/s	250 pkts/s
	corresp. peak throughput (information bits only)	230 kb/s	255 kb/s

Table 2.1: System parameters of mandatory IEEE 802.15.4a modes used in all simulations.

Identifier	Ternary Code Sequence
Preamble Code 5	$-1\,0\,1\,-1\,0\,0\,1\,1\,1\,-1\,1\,0\,0\,0\,-1\,1\,0\,1\,1\,1\,0\,-1\,0\,1\,0\,0\,0\,0\,-1\,0\,0$
Preamble Code 6	$1\,1\,0\,0\,1\,0\,0\,-1\,-1\,-1\,1\,-1\,0\,1\,1\,-1\,0\,0\,0\,1\,0\,1\,0\,-1\,1\,0\,1\,0\,0\,0\,0$
SFD Code	$0\,1\,0\,-1\,1\,0\,0\,-1$

Table 2.2: Mandatory, ternary code sequences used in all simulations.

to $N_{\text{pay}} = 1208$), are summarized in Table 2.1 as well.

Chapter 3

Related Work

The principal part of this thesis is concerned with the design of IR-UWB receivers that are robust to multi-user interference (MUI). Our main focus is on receivers that are based on the energy-detection architecture. As explained in Chapter 2, this architecture is becoming increasingly popular since it is supported by the IEEE 802.15.4a standard [35] and because it offers a great complexity reduction with respect to coherent receiver architectures. Section 3.1 gives an overview of the related work on MUI in IR-UWB networks. Related work on the energy-detection receiver architecture is discussed separately in Section 3.2.

One aspect of energy-detection receivers that is little explored is their robustness to timing impairments due to drifting clocks. The second part of this thesis is devoted to this topic and corresponding related work is discussed in Section 3.2.2.

Finally, the third part of this thesis looks at IR-UWB networks from a completely different perspective, namely from the viewpoint of security. Related work on security aspects of IR-UWB networks is treated in Section 3.3.

3.1 Multi-User Interference in IR-UWB Networks

Already the earliest research that proposes IR-UWB for wireless communication [8, 9, 10] is mainly motivated by the suitability of the IR-UWB PHY for multiple access communication. Multiple access with concurrent transmissions is possible in IR-UWB thanks to the time-hopping mechanism. [8, 9, 10] evaluate the IR-UWB system performance under the assumption of perfect LOS channel conditions without multipath fading, and by modeling MUI as a Gaussian random process. Later work showed that both of these assumptions underesti-

mate the effect of MUI on system performance: [37] shows that the interference robustness of IR-UWB due to time-hopping is significantly lowered in a multipath fading channel; The invalidity of the Gaussian approximation in practical scenarios is first established in [38, 39] and later confirmed, e.g., in [40, 41]. Due to the signalling structure with short, infrequent pulses, MUI in IR-UWB is impulsive in nature, rather than following a Gaussian distribution. In [42], an approximation of the exact distribution of MUI is given for the AWGN channel without multipath fading and assuming perfect power control, where all users are received with equal power. [43] shows that the distribution from [42] can be further approximated, yielding a Middleton Class A (MCA) model [44, 45].

The work discussed so far is concerned with the distribution of MUI in IR-UWB networks mainly because an accurate MUI model allows for better system performance evaluation through more accurate bit-error-rate (BER) calculations. Insight into the nature of MUI can however also be taken advantage of in order to design receivers that are more robust to MUI. This is usually achieved by modeling interference and noise terms with a non-Gaussian (impulsive) statistical model and by deriving optimal detection strategies for such impulsive noise environments. The resulting receiver designs mitigate the effect of interference by passing the received signal through a non-linearity prior to demodulation or decoding. This was already realized in [46, 47], where optimal non-linearities for narrowband receivers, detecting signals in MCA impulsive noise, are derived. Receivers based on this design are shown to outperform conventional receivers that assume the noise plus interference term to be Gaussian. This work also recognizes that such non-linearities need to adapt to changing interference conditions and, consequently, require an estimation of the parameters of the interference model. The MCA model is a Poisson-weighted mixture model formed by an infinite number of Gaussian component densities of increasing variance. [48] shows that truncating this model to its first two terms yields a simplified non-linearity that induces practically no performance loss. Additionally, the connection between this two term model and an empirical Gaussian Mixture Model (GMM) with two terms is made. Finally, suboptimal non-linearities that result in a thresholding structure are proposed and compared. With a thresholding structure, the receiver simply erases or limits samples, whose amplitudes lie above a certain threshold.

3.1.1 Thresholding Schemes

Thresholding schemes are simple and easy to implement. Naturally, they were therefore also among the first schemes that were proposed to mitigate MUI in IR-UWB receivers. A special

case of a thresholding scheme is hard limiting, which can be seen as equivalent to hard-decision decoding. Although the decision threshold in hard-decision decoding is usually independent of MUI, the effect of interference is still limited because individual samples cannot disproportionably contribute to the decision. [49, 50, 51] show that hard-decision decoding can yield good performance in IR-UWB networks that are subject to MUI.

[52, 53] model interference and noise in an IR-UWB network as a Gaussian-Laplacian mixture and propose a thresholding scheme, where samples that lie above a certain threshold are erased. Further, it is found that in an AWGN channel with MUI, this thresholding scheme outperforms hard limiting and performs close to a scheme with perfect blanking, i.e., where the receiver knows the time-hopping sequences of all users and erases chips suffering from pulse collisions.

[54] proposes interference mitigation with an erasure threshold as an enabling building block of an uncoordinated joint PHY/MAC protocol. A similar erasure thresholding scheme is later also proposed in [55]. In contrast to [54], not only the signal level of the user of interest, but also the signal level of the interferers is taken into account when setting the threshold.

Thresholding structures that limit high interference terms to a constant value are proposed in [51, 56, 57]. The soft-limiting structure in both, [51] and [57] is motivated by assuming that MUI is distributed according to a Laplace distribution.

3.1.2 Interference Mitigation Based on Non-Gaussian Models

Non-Gaussian models that lead to alternative receiver designs have also been proposed for IR-UWB. A popular non-Gaussian model for MUI is the aforementioned GMM, see e.g. [58]. In [59], the GMM is proposed as a model for MUI and noise in IR-UWB and it is shown how to perform channel and interference statistics estimation based on this model. In Chapter 7 we extend the use of the GMM to interference mitigation during the data decoding phase and show that the GMM substantially outperforms a thresholding mechanism that is similar to the one from [54]. Further, our proposal considers a realistic multipath environment, which is neglected in several of the related work cited hereafter. Following up on our work, [60] proposes a recursive formulation of the GMM that reduced the receiver complexity. This is similar to the recursive algorithm that we use for the energy-detection receiver in Chapter 8. In [61] the performance of an IEEE 802.15.4a compliant Rake receiver in the presence of MUI is evaluated and a semi-analytical BER model based on the GMM is developed. In [62] a two-term detector based on the truncated two-term MCA model is proposed. Further, the

assumption of equal power interferers made in the MCA is relaxed and it is shown that the resulting detector is equivalent to the one resulting from a two state GMM.

In [63, 64] receivers based on modeling MUI as a generalized Gaussian distribution are derived and it is explained how the model parameters can be estimated. In addition, [64] also proposes a receiver based on the Cauchy distribution. In [65] both the GMM and the generalized Gaussian models are generalized, yielding a generalized Gaussian mixture model.

[66] proposes a Gaussian-Laplacian mixture model for MUI and noise and derives a receiver with a structure similar to the soft-limiting receiver proposed earlier in [57]. Estimation of the model parameters of this model is addressed in [67]. Finally, [68] proposes a MUI model based on the symmetric-alpha-stable distribution and presents a corresponding receiver. In an AWGN channel without multipath fading, the proposed receiver is shown to outperform receivers such as the ones based on the generalized Gaussian, Gaussian-Laplacian or Cauchy distributions.

In general, it is an interesting question, how these models compare to each other and whether some are better suited to model MUI in IR-UWB than others. A first evaluation is made in [62, 69, 70], where detectors based on the GMM, Laplace, Cauchy and generalized Gaussian distributions are compared in both AWGN [62, 70] and multipath channels [69] and for both BPSK and BPPM modulation. In addition, an alternative non-linearity known as the α-detector [71] is proposed for IR-UWB. All of the resulting receivers are shown to significantly outperform conventional receivers. The receiver based on the Laplace distribution is shown to have inferior performance compared to the other receivers that mitigate interference. The remaining receivers are found to have a similar performance, with a slight advantage for the receiver based on the GMM. In [72, 73] the ability of the GMM, MCA and Laplace distribution to characterize MUI in an AWGN channel is evaluated and the GMM is again found to give the most accurate result.

An interesting observation is made in the recent article [65], where all receivers based on distributions that can be expressed as a mixture of generalized Gaussian distributions are compared in an IEEE 802.15.4a setting with combined BPPM/BPSK modulation. It is found that in a multipath channel with MUI, a GMM based receiver can be outperformed by both a generalized Gaussian based and a Laplace based receiver if interference mitigation is applied jointly to demodulation and decoding. This underlines once more the necessity to consider the effect of the multipath channel when assessing the system performance of IR-UWB networks that are subject to MUI. Also note that this result is not in contradiction to the results found in [69], since there interference mitigation was only applied after demodulation, at the output of the Rake receiver.

All of the models discussed so far assume that the underlying MUI or noise process is memoryless. However, because of the multipath channel, it is highly likely that consecutive samples are affected by MUI, thus leading to errors that occur in bursts. To model this behavior, we propose to model MUI in IR-UWB with a Hidden Markov Model (HMM) and we derive a corresponding receiver in Chapter 7. The HMM and in particular its simplest version, the two state Gilbert-Elliot model, are well-established models for burst-noise channels [74, 75] but were not proposed before to model MUI in IR-UWB networks. We are not aware of any other model for MUI in IR-UWB that takes correlation between samples into account.

Our HMM MUI model is extended in [76] to a two-phase two-state Markov model that better characterizes MUI if BPPM modulation is used. The extended model is then used to establish a semi-analytical framework allowing to evaluate the performance of an IEEE 802.15.4a compliant Rake receiver in the presence of MUI. It is shown that this model provides an excellent characterization of MUI. Further, it is found that the performance of the Rake receiver considerably degrades in the presence of strong interferers, emphasizing once more the need for interference mitigation schemes.

[77, 78] propose a receiver performing joint estimation and decoding under the HMM interference model. A similar approach is taken in [79] and it also resembles the algorithm our receiver for the HMM uses in Chapter 7. The difference to our algorithm is that in [77] all parameters of the HMM are estimated directly on the data sequence to be decoded, whereas in Chapter 7 we only estimate the noise state during decoding while the model parameters are estimated during a training sequence. Further, [77] compares the performance of the HMM receiver to a receiver using the memoryless GMM. A similar comparison is also done in [80] where different decoding metrics for coded and interleaved transmissions over HMM channels are analyzed. Both papers find that a receiver based on the HMM can significantly outperform a receiver that is based on the GMM. In contrast, we only identify a minor performance difference in Chapter 7. Note, however, that in [77, 80] interference is generated directly from a HMM and thus exactly matches the model of the receiver, whereas in Chapter 7 interference stems from simulated parallel IR-UWB transmissions. Further, [80] finds that the performance difference increases dramatically with interleaving, which we do not use in Chapter 7.

3.1.3 Interference Mitigation During Synchronization

Interference mitigation techniques are not limited to data decoding but they can also be applied during synchronization or for time-of-arrival (TOA) estimation in a multi-user environment.

Thresholding is applied for timing acquisition in IR-UWB with a coherent receiver in [81]. In the presence of near interferers, the proposed scheme is shown to significantly outperform traditional acquisition schemes that do not use thresholds. The thresholds must be adapted to varying channel and interference conditions. However, in [81] optimal thresholds are found through exhaustive simulations and it is left open how to adapt these thresholds in practice. We build on this work in Chapter 6 and propose a practical timing acquisition algorithm for energy-detection receivers. Our algorithm not only significantly reduces the impact of MUI, but it also solves the problem of adapting the involved thresholds. The algorithm from [81] is validated experimentally on a hardware platform in [82]. This work shows that concurrent transmissions during IR-UWB acquisition are possible, provided that appropriate interference mitigation is performed at the PHY.

To suppress MUI during TOA estimation, [83, 84] use non-linear filtering techniques known from digital image processing. Both works are based on energy-detection and show interesting parallels to our work. [83], e.g., uses robust statistics in the form of the median to suppress MUI, similar to what we propose for channel estimation in Chapter 5. Further, both papers represent the receiver output as an energy matrix that can be considered a digital image onto which image processing techniques are applied. We also take a similar approach in our clock-offset tracking algorithm in Chapter 9. Finally, their mechanisms could be combined with our synchronization algorithms presented in Chapter 6. Both papers only consider TOA estimation, assuming that packet detection and coarse synchronization on the strongest multipath component have already been achieved. Further, also here parameter adaptation, especially of the involved thresholds, is needed in practical deployments but no online adaptation mechanisms are proposed.

Statistical interference models can also be used for receiver operations other than data reception. [85] models MUI with a GMM, yielding a TOA estimator for IR-UWB ranging systems that operate in a multi-user environment. [86] devises a robust signal acquisition algorithm for transmitted reference IR-UWB receivers assuming impulsive noise modeled with a GMM. [87] estimates clock drifts and carrier frequency offsets in a coherent IR-UWB receiver, again under the assumption that MUI is modeled according to a GMM. For narrowband direct-sequence spread-spectrum systems, [88] proposes a signal acquisition scheme based on a non-linearity derived under the assumption that MUI follows a symmetric-alpha-stable distribution.

3.1.4 Multi-User Detectors

We would like to stress that all the techniques to handle MUI in IR-UWB that we reviewed so far are different from multi-user detection (MUD). MUD extends or adapts to IR-UWB, classical, well-established joint decoding techniques [89], which are also used in other systems like CDMA. They aim at canceling or suppressing interference by jointly estimating and decoding the signals of a large number of users. For example, a near-far interferer would be jointly received instead of being treated as interference. This annihilates the near-far effect and makes joint decoding potentially attractive. However, an optimal joint processing of all users [90] is mostly not possible due to its very high complexity. Therefore, suboptimal methods like minimum mean-square error MUD or receivers employing successive interference cancelation are used [91]. All of these methods share the common factor that the receiver has to acquire and actively decode each of the users. This might be perfectly suited for a centrally coordinated and synchronized system, where a base station communicates with a large number of users at the same time. However, with a distributed IR-UWB system, synchronizing the receiver with all the users is extremely complex and impractical. In addition, the complexity of the decoding operation is very high. For an overview on MUD in IR-UWB, we refer to [92].

3.1.5 Interference Management on the MAC Layer

MUI can also be managed on the MAC layer, through a mutual exclusion protocol (e.g., TDMA) and/or power control. However, neither of these approaches can prevent uncontrolled interference from coexisting non-coordinated piconets. Further, it has been shown that optimal MAC protocols for low data-rate IR-UWB networks explicitly allow concurrent transmissions without power control, while employing interference mitigation on the PHY [1, 93, 54, 94, 95, 96]. Many MAC protocols for IR-UWB [54, 97] are therefore uncoordinated, including the MAC protocol of the IEEE 802.15.4a standard [35].

3.1.6 Other Types of Interference in IR-UWB Networks

As already mentioned in Chapter 2, this thesis focuses solely on interference from concurrently transmitting IR-UWB devices. Coexistence with other systems and in particular narrowband interference are not considered. For a general survey on UWB coexistence, the reader is referred to [16] and the references therein. For an overview of recent research efforts on how to deal with narrowband interference in the case of non-coherent receiver architectures, we propose [98] as a starting point.

Throughout this thesis we also assume that the frame structure of IR-UWB PHY packets is such that intersymbol interference (ISI) does not occur. Relaxing this assumption allows for higher raw bit rates. However, the resulting ISI is not without impact on the IR-UWB receiver, especially in the case of non-coherent reception. For an overview on how to deal with ISI in this case, we also refer to [98]. We also point out [99] that makes use of the EM-algorithm together with an energy-detection receiver. In [99], the EM-algorithm is used to perform channel estimation under ISI, whereas we use it to estimate parameters of a MUI model in Chapter 8.

3.2 Energy-detection Receiver Architectures

The energy-detection architecture recently gained a lot of popularity since it is explicitly supported by the IEEE 802.15.4a standard [35]. However, this architecture has already been considered for signal detection at the end of the sixties [28]. In the context of IR-UWB, energy-detection is first mentioned in [100, 101], also for detecting the presence of a signal. For communication in IR-UWB, energy-detection is first proposed in [102, 103] in conjunction with BPPM modulation and it is shown that it can in theory achieve a very good performance in a fading multipath channel. This work is extended in [50] to a multi-user setting and energy-detection with OOK modulation. Energy-detection receivers with OOK are also proposed in [104, 105] where an approximation of the optimal OOK threshold is given. A generalization of this work to M-ary PAM is later presented in [106].

[107] proposes to use energy-detection receivers in IR-UWB sensor networks because of their good performance-complexity trade-off and relaxed synchronization requirements. These claims are verified in [108] where a first system analysis of an IR-UWB energy-detection receiver, encompassing both synchronization and data reception in a single user setting, is performed. Other papers that also take a similar system level approach in analyzing the performance of energy-detection receivers are [30, 109]. Further, [110] compares the data reception performance of a Rake and an energy-detection receiver in an IEEE 802.15.4a single user setting. [111, 112, 113] consider ranging with an energy-detection receiver in the same setting. [111] moreover describes a complete synchronization algorithm for an energy-detection receiver, including timing acquisition and SFD detection.

[32] finds that the integration time of an energy-detection receiver needs to be adapted to varying channel conditions to prevent excessive integration of too much noise. This is also recognized by [104, 105, 114, 30, 115, 109]. [114, 115] additionally propose algorithms that

allow the determination of the optimal integration time. [30] proofs that the energy-detection receiver is the optimal detector if no channel state information is available and that the optimal decision rule for BPPM in this case consists in comparing the energies in the 0-block and 1-block of a BPPM frame. This is also shown in [31]. [31] also derives the optimal ML receiver for BPPM in the case where partial channel state information, in the form of an average channel power-delay profile (APDP), is available. It is shown that this receiver corresponds to an energy-detection receiver where the received signal prior to integration is weighted with the APDP. [33] shows that such a weighting can also be performed in the digital domain, with weights corresponding to the estimate of the current channel energy-delay profile, and results in a performance improvement compared to [31]. [33] also shows how to estimate the channel energy-delay profile in a single user setting. The resulting energy-detection receiver architecture is the closest one to the architecture that we use throughout this thesis. What sets our receiver apart from the work discussed here is that we also address the calculation of the optimal weighting function in the case where the signaling between preamble and payload differs, such as it is the case with IEEE 802.15.4a. Further, our receiver addresses MUI, which is also not done by any of the related work discussed so far. Finally, we take an integral system-level approach including synchronization, whereas most of the related work solely focuses on data reception.

[116] and [117] devise weighting schemes similar to [33], but for energy-detection with OOK modulation. [118] derives the optimal detectors with and without ISI and with various amounts of channel state information that are available at the receiver. In the absence of ISI, a detector equivalent to the one from [33] is found if the energy-delay profile of the channel is available,[1] and the one from [31] is found if only an average in the form of the APDP is available.

A similar analysis can be made for TOA estimation. [119, 120] show that if no information about the channel is available, the optimal TOA estimator is an energy-detection receiver with a sliding integration window; if the APDP is available, the ML TOA estimator is a similar energy-detection receiver, where the received samples are weighted with a weighting factor that depends on the APDP.

Examples of hardware structures to implement an energy-detection receiver and evaluations thereof can be found in [121, 122, 123, 124].

1. Called "instantaneous power delay profile" in [118]. Further, the equivalence between the two detectors can be shown through our derivation of the optimal decision rule given in Chapter 5 and by noting that $\cosh(x) = {}_0F_1\left(;\frac{1}{2};\frac{x^2}{4}\right)$.

3.2.1 Energy-detection Receivers and Multi-User Interference

Apart from [50, 83, 84], none of the work cited so far considers energy-detection receivers in a multi-user environment. [50] determines the theoretically achievable rates in networks of energy-detection receivers that use hard-decision decoding. [83, 84] mitigate interference during TOA estimation through non-linear filtering. The lack of research on energy-detection receivers subject to MUI may partly be explained by the fact that early proposals for IR-UWB networks with energy-detection receivers such as [107] assume a MAC protocol that enforces mutual exclusion and therefore neglect MUI. However, at the latest since the adoption of a non-coordinated MAC protocol for IEEE 802.15.4a, this is an assumption that has to be revised. Further, even with a perfectly coordinated MAC protocol, interference from coexisting networks cannot be entirely prevented. Understanding the impact of MUI on energy-detection receivers and adapting their architecture such that they are able to cope with MUI is therefore of practical importance, and was, prior to this thesis, still a largely open problem. Recently, this has been more widely recognized and lead to several research efforts on energy-detection receivers in a multi-user environment.

[125] develops a framework to analyze the BER of an energy-detection receiver in the presence of uncoordinated interferers that are distributed according to a spatial Poisson process. Further, in their analysis [125] model MUI with a stable distribution.

[126] considers a weighted energy-detection receiver for OOK demodulation similar to [116] and proposes to mitigate MUI by modifying the likelihood ratio underlying the weighting function. This scheme has an effect similar to a thresholding scheme where high interference terms are erased such that they do not contribute to the decision.

[127, 128] propose an alternative modulation scheme where information is carried in the shift of orthogonal ternary sequences that are similar to those used in the IEEE 802.15.4a preamble. The scheme is suitable for reception with an energy-detection receiver and different sequences can be assigned to different users allowing for multiple access. Interestingly, the proposed scheme leads to a receiver architecture that is similar to the one we propose in Chapter 5 because it also tries to account for inter-pulse interference. Unlike in our work, cross-terms are neglected and consequently their architecture misses the additional branches found in our architecture.

In [129] the authors devise a methodology to analyze PHY effects on IR-UWB MAC protocols. They assume an energy-detection receiver and analyze the IEEE 802.15.4a MAC with various clear channel assessment modes. The authors demonstrate that PHY effects, and in

particular the possibility to have concurrent transmissions, have a significant impact on the performance of the MAC protocol.

In [130] an IR-UWB PHY model for network simulations is developed where the receiver is assumed to follow the energy-detection architecture.

3.2.2 Energy-Detection Receivers and Clock-offset Tracking

Energy-detection receivers with long integration times are robust to timing impairments such as clock offsets caused by differences between the frequencies of the transmitter and receiver oscillator. On the other hand, long integration times lead to noise enhancement. Consequently, the integration time is often shortened to a duration in the order of the channel spread or a weighting function is applied, as we have seen in Section 3.2. This in turn increases the sensitivity of energy-detection receivers to clock offsets, which was already recognized in [31]. However, to the best of our knowledge, this effect has neither been studied, nor have specific solutions been proposed.

For narrow-band physical layers, there are two well-established solutions for addressing clock offsets: the phase-locked loop (PLL) and early-late gate synchronizers [131]. Both solutions allow for continuous tracking of the transmitter clock by the receiver. Although applicable for coherent IR-UWB receivers, both solutions are not suited for IR-UWB energy-detection receivers. The PLL cannot be used because it either requires a constant carrier and/or the availability of phase information to create an appropriate error signal. Early-late gate synchronizers cannot be used because they require sampling of a symmetric correlator output.

For coherent IR-UWB receivers, which are very sensitive to timing impairments, [23, 132, 133, 134, 135, 136] address the issue of clock-offset tracking. In general, they use an adaptation to IR-UWB of one of the classical solutions for narrowband systems. Further, [87, 137] address frequency offset estimation with a coherent receiver, but not tracking.

Closest to our work in Part II is probably [138]. [138] describes a maximum likelihood estimator for clock-offset estimation that can be used in a coherent or non-coherent IR-UWB receiver. To estimate the drift, received signals are accumulated and arranged in a trellis. The best path on the trellis then yields the maximum likelihood estimate. Interestingly, a connection to our work can be made by noticing that the proposed trellis also yields a sort of two-dimensional representation of the received signal, similar to the energy matrix in Chapter 9. Contrary to our work, the estimator assumes perfect synchronization on the strongest multipath component and it also disregards signal contributions from other multipath components. Further, tracking and

compensating for the drift are not treated.

3.3 Security of IR-UWB

Security considerations were already made in the very early work on IR-UWB. The first papers on energy-detection receivers for IR-UWB are essentially security motivated and analyze the covertness of IR-UWB [100, 101]. Compared to narrowband systems, IR-UWB generally has a better covertness and a better robustness against jamming. The first is due to the inherently low signal levels and infrequent transmissions, the second is due to the wide bandwidth. On both topics there exists a quite large body of research, see e.g., [139, 140, 141, 142, 143].

Other work proposes to take advantage of the reciprocity of the rich UWB channel in order to establish shared secret keys between two devices communicating over IR-UWB [144, 145, 146].

Security in ranging has mainly focused on cryptographic secure ranging protocols, or distance-bounding protocols [147, 148, 149, 150, 151, 152, 153, 154, 155, 156, 157, 158, 159, 160]. These protocols are not specific to IR-UWB but most of them assume an accurate ranging mechanism that is based on round-trip-time measurements of radio signals. The secure ranging protocol then guarantees that the measured distance between two devices is an upper bound of their actual distance, even in the presence of an adversary that interferes with the ranging process. Secure ranging protocols thus prevent an adversary from decreasing the estimated distance.

Its potential for high precision ranging makes IR-UWB an ideal candidate technology to provide the ranging primitive required by secure ranging protocols. This has also been recognized by several proposals for secure ranging [150] or secure localization [152, 161] protocols. However, with the exception of [160], none of the protocols makes any specific assumptions about the IR-UWB PHY apart from assuming that it can provide accurate ranging measurements. [160] proposes a secure ranging protocol and evaluates it on an IR-UWB hardware platform. Except from interfacing the protocol with the IR-UWB ranging devices, also here the specific characteristics of the IR-UWB PHY are not taken into account. Consequently, to date the implications of performing secure ranging over an IR-UWB PHY remain unstudied.

However, the PHY can play an important role when it comes to the security of ranging mechanisms. This was discovered by [162] which points out attacks on the PHY as a possible attack vector against secure ranging protocols. These attacks make it possible for an adversary to decrease the estimated distance without breaking any cryptographic primitives or protocols.

Subsequently, [163] demonstrated with a proof-of-concept narrowband implementation that some of these attacks are indeed feasible for the ISO 14443 PHY and a compliant RFID receiver, as well as for 433MHz ASK/FSK modulation and a super-heterodyne receiver, common in wireless sensor networks. PHY attacks are by nature PHY-specific. Therefore, the results of [162, 163] cannot be mapped easily to other PHYs, including the IR-UWB PHY. We are not aware of any work, except for what we present in Chapter III, that addresses PHY attacks on IR-UWB ranging. Further, none of the attacks of the related work addresses synchronization.

Another aspect of security in ranging is location privacy. [156] analyzes location privacy of distance bounding protocols. It is found that an adversary overhearing the exchange of ranging messages may be able to infer information about the location of ranging devices that is not acceptable in certain security-sensitive applications. Again, this protocol is not specific to any particular PHY technology.

The need for location privacy and security was also recognized in the development of the IEEE 802.15.4a standard and led to the inclusion of an optional private ranging mode [35, 164]. The private ranging mode attempts to make overhearing harder by essentially letting devices choose at random, which of the proposed preamble codes they use in a ranging message exchange. Further, the private ranging mode gives devices the possibility for message authentication, thus preventing the injection of bogus messages.

The growing importance of security has also been recognized in a recent survey article on IR-UWB ranging, where security in IR-UWB ranging is mentioned as a future research direction needed in order for IR-UWB ranging to become widely accepted [165].

Part I

Robust IR-UWB Receiver Design

Chapter 4

Performance Evaluation of the IEEE 802.15.4a Physical Layer with Energy Detection and Multi-User Interference

One of the main goals of this thesis is to understand how an IR-UWB receiver performs in the presence of multi-user interference (MUI). Thanks to its wide bandwidth, combined with a low duty cycle time-hopping mechanism on the physical layer (see Section 2.2), IR-UWB should show a good robustness against MUI and thus allow for parallel transmissions of concurrent users.

This possibility to have parallel transmissions is appealing because of several reasons. First of all, it allows for the operation of several networks in close vicinity without coordination. Coordination of a large number of co-existing networks that are potentially being managed by different entities is difficult if not impossible in practice. Second, inside a single network, it allows for simple MAC protocols, such as ALOHA that is used in IEEE 802.15.4a. Finally, parallel transmissions increase spatial reuse and it can be shown that an optimal IR-UWB MAC protocol, achieving the highest network throughput, allows parallel transmissions [1].

In this chapter, we evaluate the performance of a complete IEEE 802.15.4a IR-UWB physical layer implementation with an energy-detection receiver and in the presence of MUI. Our goal is to assess whether such a standard compliant low-complexity receiver is robust enough to contain MUI at a level allowing for parallel transmissions.

In our evaluation we take all the specifics of IEEE 802.15.4a into account and we consider a complete system, covering all major operations needed at the receiver in order to receive an

IEEE 802.15.4a packet. We account for packet detection and timing acquisition, as well as detection of the start frame delimiter (SFD) and decoding of the payload. We further try to understand how each of these individual operations are affected by MUI.

We find that the performance of the energy-detection receiver is significantly reduced if the receiver is subject to MUI, even at low data rates. Further, the energy-detection receiver only shows a very limited *capture effect*, i.e., the ability to correctly receive a packet despite the presence of a concurrent transmission. Even with a single interferer, it performs close to a narrowband-like model of *destructive collisions*, where a packet is lost if another transmission is active at the same time. We conclude that using an energy-detection receiver in combination with IEEE 802.15.4a significantly diminishes one of the most appealing benefits of UWB, namely its robustness to MUI and thus the possibility to allow parallel transmissions.

Finally, we analyze the impact of MUI on the different receiver operations. This allows us to identify the main effects causing the above mentioned performance degradation. The insight gained can serve as a guideline for the design of more robust receivers.

This chapter is organized as follows: In Section 4.1 we describe the architecture of the energy-detection receiver that we use in our performance evaluation. The specific algorithms used to perform main receiver operations such as packet detection and timing acquisition are given in Section 4.2. Section 4.3 constitutes the performance evaluation of the complete system with the IEEE 802.15.4a PHY and in the presence of uncontrolled MUI. Conclusions of this evaluation are drawn in Section 4.4.

4.1 Receiver Architecture: Energy-detection Receiver with Adaptive Channel Mask

Throughout this performance evaluation we use the classical energy-detection receiver architecture introduced in Section 2.3.3. In addition to what was introduced in Section 2.3.3, the receiver considered here adapts the total integration time per symbol to varying channel conditions. Adaptation of the integration time is important because it ensures that only parts of the received signal containing useful information contribute to the decision, whereas parts consisting purely of noise are ignored [32].

The architecture of this energy-detection receiver is depicted in Figure 4.1. We assume that the receiver takes M samples per BPPM block. It adapts to varying channel conditions by weighting these received samples with a *channel mask*. The channel mask is a binary vector

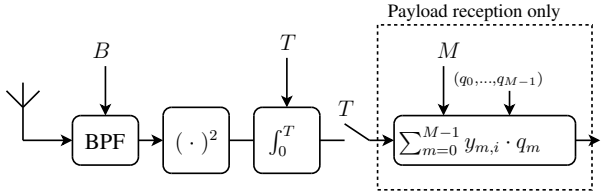

Figure 4.1: Architecture of the energy receiver. The antenna is followed by a bandpass filter, a squaring device, and the integrator. The signal at the output of the integrator is sampled every T seconds. Finally, for the i-th symbol, a scalar product is computed between a block of M samples $(y_{0,i}, \ldots, y_{M-1,i})$ and the channel mask $(q_0, \ldots, q_{M-1})$.

of length M, $(q_0,\ q_1, \ldots,\ q_{M-1})$, $q_m \in \{0, 1\}$, indicating which samples of the received signal contain useful information and which samples do not. The i-th symbol, d_i, of the payload is then demodulated according to the following decision rule:

$$\sum_{m=0}^{M-1} y_{m,i} \cdot q_m \underset{\hat{d}_i=1}{\overset{\hat{d}_i=0}{\gtrless}} \sum_{m=0}^{M-1} y_{m+N_\mathrm{f}/2,i} \cdot q_m \tag{4.1}$$

$y_{m,i}$ is given by (2.17), i.e.,

$$y_{m,i} = \int_{mT+\psi(i)}^{(m+1)T+\psi(i)} [r_{\text{pay}}(t)]^2 \, \mathrm{d}t \tag{4.2}$$

with

$$\psi(i) = iT_\mathrm{f} + c_{\text{THS},i} T_{\text{burst}} + \hat{\nu}_0. \tag{4.3}$$

and $r_{\text{pay}}(t)$ is given by (2.30). The receiver thus compares the energies in the first and second half of the BPPM frame after applying the channel mask. The binary output of the comparison in (4.1) is then fed to the Reed-Solomon decoder for error correction. Both, the channel mask and the TOA are estimated during the reception of the IEEE 802.15.4a preamble as we will see in the following. The length of the channel mask M is dictated by the minimum inter-pulse spacing in the preamble which equals $L_\mathrm{s}T_\mathrm{c}$. We assume that the maximum delay spread of any channel realization is shorter than $L_\mathrm{s}T_\mathrm{c}$ to prevent ISI. We therefore choose $M = \frac{L_\mathrm{s}T_\mathrm{c}}{T}$ and we assume that the integration time T is such that M is an integer.

4.2 Main Receiver Operations

In order to receive and demodulate the information that is contained in the payload of the packet, the receiver first has to detect the presence of the packet on the wireless channel. Further, it has to determine the beginning of the packet, or in other words estimate its TOA, ν_0. This process is commonly referred to as *synchronization*. Infrastructure-less packet based wireless networks typically lack global synchronization and synchronization has thus to be performed on a packet-per-packet basis. In IEEE 802.15.4a, this is done with the help of the known preamble sequence whose structure we have described in Section 2.4.2.

The receiver performs synchronization in two phases: *packet detection and timing acquisition* and *SFD detection*. In addition, in between these two phases, it performs *estimation of the channel mask*. In the following, we will discuss each of these receiver operations independently.

4.2.1 Packet Detection and Timing Acquisition

There are two steps for packet detection and timing acquisition: *coarse timing acquisition* and *fine timing acquisition*. The coarse timing acquisition algorithm tries to locate the starting time of one of the N_{sync} preamble symbols. It usually synchronizes on the strongest multipath component, which is not always the first in time. Coarse timing acquisition is therefore followed by a fine timing acquisition phase in order to improve the timing accuracy.

Coarse Timing Acquisition

The coarse synchronization algorithm is a classic timing acquisition algorithm using a correlation of the receiver output with a template derived from the known preamble code sequence of the UOI. The discrete time signal during the reception of the preamble is given by

$$y_m^{\text{pre}} = \int_{mT}^{(m+1)T} [r_{\text{pre}}(t)]^2 \, \mathrm{d}t, \tag{4.4}$$

where $r_{\text{pre}}(t)$ is given by (2.28).

A template t_l of length $M_{\text{T}} = N_{\text{G}} \cdot C \cdot M$ is formed by repeating the preamble code N_{G} times to obtain processing gain. The code symbols forming the template are squared due to

non-coherent reception. Hence

$$t_l = \sum_{k=0}^{N_{\mathrm{G}}-1} \sum_{j=0}^{C-1} c_j^2 \cdot \delta_{l-(j+kC)M} \tag{4.5}$$

where δ_m denotes the Kronecker delta.

The (discrete) correlation output that follows is

$$z_m = \sum_{l=0}^{M_{\mathrm{T}}-1} t_l \cdot y^{\mathrm{pre}}_{m-(M_{\mathrm{T}}-1)+l} \tag{4.6}$$

The preamble of the UOI is $L_{\mathrm{s}}T_{\mathrm{c}} \cdot C$-periodic. Consequently, z_m is $M \cdot C$-periodic if the UOI signal is present. Therefore, the algorithm processes the correlation output by blocks of MC consecutive samples. The i-th block is

$$\mathbf{z_i} = \{z_{iMC}, z_{iMC+1}, \ldots, z_{(i+1)MC-1}\} \tag{4.7}$$

Further, the algorithm has two steps: *detection* and *verification*. During detection, the presence of a signal is declared if at least one of the correlation output samples of the current block exceeds the *detection threshold* η_{detect}. The statistics of the receiver output are known if the received signal is AWGN only (see Section 2.3.3). Hence, we can set the threshold according to

$$\eta_{\mathrm{detect}} = \frac{N_0}{2} F^{-1}_{\chi^2}(1 - P^{\mathrm{FA}}_{\mathrm{AWGN}} | 2BT \cdot C_{\mathrm{NZ}} \cdot N_{\mathrm{G}}). \tag{4.8}$$

where $C_{\mathrm{NZ}} = \sum_{j=0}^{C-1} c_j^2$ denotes the number of nonzero code symbols of the UOI preamble code and $F_{\chi^2}(x|\kappa)$ is the cumulative distribution function (CDF) of a chi-square random variable with κ degrees of freedom. The product $C_{\mathrm{NZ}} \cdot N_{\mathrm{G}}$ corresponds to the number of samples that are combined due to the template. The threshold is set by fixing the design parameter $P^{\mathrm{FA}}_{\mathrm{AWGN}}$, which is the probability that AWGN only can exceed the threshold. When setting the threshold we assume that the noise PSD, $N_0/2$, is known. This assumption is made because thermal noise is generated by the receiver circuitry and can thus be calibrated. Further, we will show in Chapter 5 how it can also be estimated in a robust fashion.

If the presence of a signal is detected in the i-th block, we declare initial timing acquisition on the sample with i-th block index $j_m^{\max}$, having the highest correlation output value, i.e. $z_{iMC+j_m^{\max}} \geq z_{iMC+j}$, $\forall j \in \{0, \ldots, MC-1\}$. This selection criterion corresponds to the hybrid maximum selection and threshold crossing criterion that was proposed in [166].

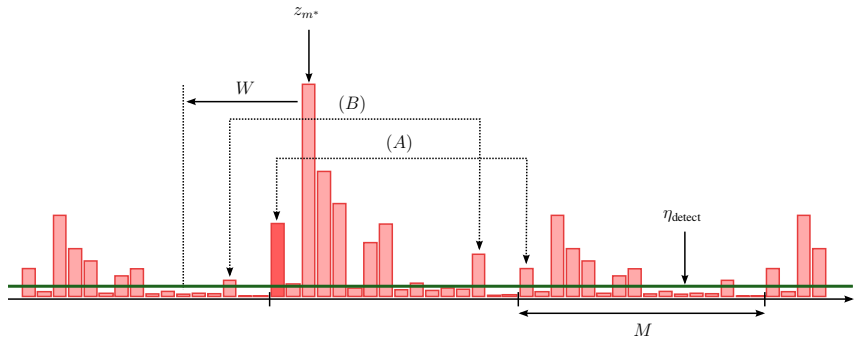

Figure 4.2: Illustration of the jump-back-and-search-forward algorithm. Starting from the result of the coarse synchronization z_{m^*}, a window of length W is searched. Although both (A) and (B) are above the threshold η_{detect}, only (A) fulfils the second condition given by equation (4.9) and is thus selected as first path.

In the verification phase, we require that for every block $i+k$, $k = 1, \ldots, N_{\text{V}}$, the maximum value $z_{(i+k)MC+j_{i+k}^{\max}} > \eta_{\text{detect}}$ and that its index $j_{i+k}^{\max}$ does not differ by more than the minimum inter-pulse distance M from the current synchronization point $j_{i+k-1}^{\max}$ of the preceding block. This ensures that both maxima stem from the same preamble pulse, the receiver thus verifies that the correlation output has the expected periodicity. If verification is fulfilled for N_{V} consecutive blocks, the verification phase succeeds and synchronization is declared. If one verification fails, synchronization starts anew with the detection phase.

Fine Timing Acquisition

After coarse timing acquisition we are synchronized on one of the multipath components of the first pulse of one of the preamble code repetitions. Usually, this will correspond to the strongest multipath component rather than the first in time. During fine timing acquisition we therefore perform a jump-back-and-search-forward algorithm to find the first sample with significant energy belonging to the same pulse. Our jump-back-and-search-forward algorithm is illustrated in Figure 4.2 and explained in what follows.

We limit our search-back window to $W = M/2$ samples, thus assuming that the first and the strongest path are not separated by more than the minimum inter-pulse distance. Assume that during coarse timing acquisition we synchronized on sample z_{m^*}. I.e., the synchronization point $m^* = (i + N_{\text{V}})MC + j_{i+N_{\text{V}}}^{\max}$ if the signal was first detected in the i-th block and all of

the N_{V} following verifications succeeded. The set of candidates for fine timing acquisition is then $\{z_{m^*-W}, \ldots, z_{m^*-2}, z_{m^*-1}, z_{m^*}\}$. Evidently, to be accepted, a candidate should be above the AWGN threshold η_{detect}. This is the standard criterion for jump-back-and-search-forward algorithms proposed in the literature (see, e.g., [167]). However, this condition is not sufficient because with non-coherent reception, the minimum of the autocorrelation of each preamble code is nonzero and equal to half its maximum value. The result is that even if the template is misaligned, it is still aligned with half of the pulses of the preamble code and the correlation is thus likely to exceed the AWGN threshold. We therefore use an additional heuristic to make sure that when searching back we do not choose samples that stem from the last pulse of the preceding preamble symbol. The heuristic adds the constraint that a candidate z_{m^*-k} needs to fulfil the condition

$$z_{m^*-k} \geq z_{m^*-k+M} \tag{4.9}$$

in order to be accepted. The idea here is that a correlation peak that stems from a fully aligned template will always be higher than a secondary peak that occurs with a misaligned template. This is illustrated in Figure 4.2. In fine synchronization we thus choose from the candidate set the sample with the lowest index that lies above the detection threshold η_{detect} and fulfils condition (4.9).

4.2.2 Estimation of the Channel Mask

Between fine synchronization and SFD detection, we estimate the channel mask. The channel mask is essentially a quantized version of the channel energy-delay profile: portions of the channel energy-delay profile that consist of noise only, will have a corresponding zero-valued entry in the channel mask; the other entries will be equal to one.

At this point, the receiver knows the starting point of a preamble symbol. Since it also knows the preamble code, it can deduct at what time instances pulses are being received. The received signal corresponding to a single pulse is sampled M times. To estimate the channel mask, the receiver can therefore accumulate blocks of M samples over several preamble symbols. A threshold similar to (4.8) is then applied to quantize the channel mask and distinguish between signal and noise. Assuming that $y_{m_{\text{fine}}}$ is the sample onto which we synchronized in the fine timing acquisition, this results in the following channel mask

$$q_m = \begin{cases} 1 & \text{if } \bar{y}_m \geq \eta_{\text{mask}} \\ 0 & \text{otherwise.} \end{cases}, m = 0, 1, \ldots, M-1 \tag{4.10}$$

where

$$\bar{y}_m = \sum_{i=0}^{N_{\text{CH}}-1} \sum_{j=0}^{C-1} c_j^2 \cdot y^{\text{pre}}_{m_{\text{fine}}+(j+iC)M+m} \tag{4.11}$$

The preamble code symbols c_j are squared here because we only want to take into account the C_{NZ} code symbols for which the preamble code is nonzero. N_{CH} are the number of preamble symbols over which we accumulate the received signal in order to get some processing gain. The threshold to distinguish signal from AWGN only is obtained the same way as (4.8) and equals

$$\eta_{\text{mask}} = \frac{N_0}{2} F^{-1}_{\chi^2}(1 - P^{\text{FA}}_{\text{AWGN}} | 2BT \cdot C_{\text{NZ}} \cdot N_{\text{CH}}). \tag{4.12}$$

4.2.3 Detection of the Start Frame Delimiter

After estimation of the channel mask, the receiver begins to look for the SFD sequence. SFD detection works according to a correlation procedure where the received signal is correlated with a template that is derived from the SFD code. During SFD detection the energy of samples belonging to the same preamble symbol is combined. This yields the sequence $\tilde{\mathbf{y}} = (\tilde{y}_0, \tilde{y}_1, \ldots)$ where each element contains the energy accumulated during one preamble symbol. When combining, both the preamble code and the channel mask are taken into account, i.e., we only combine samples where both the preamble code and the channel mask are nonzero. The i-th element of this sequence is then given by

$$\tilde{y}_i = \sum_{j=0}^{C-1} c_j^2 \sum_{m=0}^{M-1} y^{\text{pre}}_{m_{\text{fine}}+(j+iC+N_{\text{CH}}C)M+m} \cdot q_m \tag{4.13}$$

The offset $m_{\text{fine}} + N_{\text{CH}}CM$ accounts for the samples of the received signal y_m that were used up during timing acquisition and channel mask estimation.

During the SYNC part of the preamble, the expected received signal energy of an element is

$$E_{\text{S}} = \mathbb{E}[\tilde{y}_i] = \sum_{m=0}^{M-1} q_m \mathbb{E}\left[\sum_{j=0}^{C-1} c_j^2 \cdot y^{\text{pre}}_{m_{\text{fine}}+(j+iC+N_{\text{CH}}C)M+m} \right] = \frac{1}{N_{\text{CH}}} \sum_{m=0}^{M-1} q_m \cdot \bar{y}_m \tag{4.14}$$

where $\bar{y}_m$ is given by (4.11) and obtained during channel mask estimation. During reception of the SFD, some preamble symbols are absent due to the modulation with a zero-valued symbol of the ternary SFD code (Section 2.4.2). For these preamble symbols, $\tilde{y}_i$ consists of noise only and the expected noise level corresponds to the expected value of a chi-square random variable

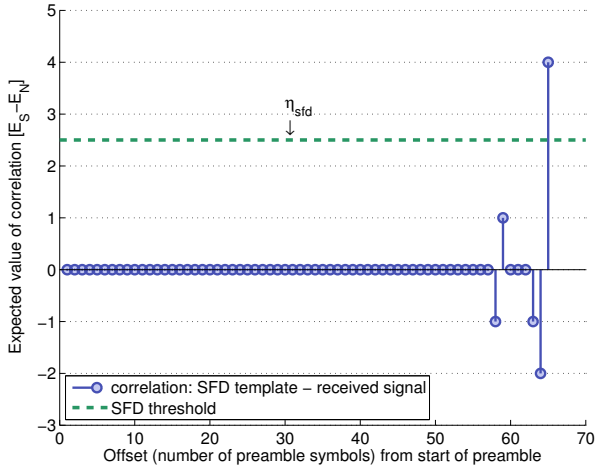

Figure 4.3: Expected value of the correlation between the bipolar SFD template and the received signal during an IEEE 802.15.4a preamble. E_{S} is the expected energy of preamble symbols containing a signal contribution, E_{N} of preamble symbols consisting of noise only. Detection of the SFD is declared if the correlation exceeds the SFD threshold $\eta_{\mathrm{sfd}} = 2.5 \cdot (E_{\mathrm{S}} - E_{\mathrm{N}})$.

(Section 2.3.3) and is given by

$$E_{\mathrm{N}} = \mathbb{E}[\tilde{y}_i] = 2BT \cdot M_{\mathrm{NZ}} \cdot C_{\mathrm{NZ}} \cdot \frac{N_0}{2} \tag{4.15}$$

where $M_{\mathrm{NZ}} = \sum_{m=0}^{M-1} q_m$ is the number of nonzero elements of the channel mask. For preamble symbols during the SFD where the SFD code is nonzero, the expected signal level is given by (4.14) like in the SYNC part.

With the mandatory IEEE 802.15.4a SFD code

$$\mathbf{s}^{(\mathbf{sfd})} = (0, 1, 0, -1, 1, 0, 0, -1), \tag{4.16}$$

the expected received sequence is then

$$\mathbb{E}[\tilde{\mathbf{y}}] = (\mathbb{E}[\tilde{y}_0], \mathbb{E}[\tilde{y}_1], \ldots) = (\underbrace{E_{\mathrm{S}}, E_{\mathrm{S}}, \ldots, E_{\mathrm{S}}}_{\text{SYNC part}}, \underbrace{E_{\mathrm{N}}, E_{\mathrm{S}}, E_{\mathrm{N}}, E_{\mathrm{S}}, E_{\mathrm{S}}, E_{\mathrm{N}}, E_{\mathrm{N}}, E_{\mathrm{S}}}_{\text{SFD}}), \tag{4.17}$$

where the length of the SYNC part is unknown.

To detect the SFD, the receiver correlates $\tilde{\mathbf{y}}$ with the bipolar template sequence

$$\mathbf{t}^{(\mathbf{sfd})} = (-1, 1, -1, 1, 1, -1, -1, 1) \tag{4.18}$$

$\mathbf{t}^{(\mathbf{sfd})}$ is obtained from (4.16) by setting nonzero elements of the SFD to 1 and zero valued elements to -1. The advantage of this bipolar sequence is that its correlation with the received signal is zero during the SYNC part. The entire correlation is shown in Figure 4.3. We observe that its expected maximum, corresponding to a perfect alignement with the SFD, has a value of $4(E_{\mathrm{S}} - E_{\mathrm{N}})$. The highest secondary peak of the correlation is at $E_{\mathrm{S}} - E_{\mathrm{N}}$. The receiver therefore declares detection of the SFD, if the output of the correlation exceeds the optimal threshold, η_{SFD}, that lies in the middle between these two values

$$\eta_{\mathrm{sfd}} = 2.5 \cdot (E_{\mathrm{S}} - E_{\mathrm{N}}) \tag{4.19}$$

4.3 Performance Evaluation

To evaluate the performance of the energy-detection receiver, we focus on the mandatory data rate of 0.85 Mbit/s with the LPRF mode and on the mandatory channel number 3 with preamble codes 5 and 6 (see Section 2.5.1).

The main parameters of the receiver are summarized in Table 4.1. The integration time is $T = 8$ ns, which corresponds to the length of a burst in the IEEE 802.15.4a LPRF mode. The parameters for synchronization and channel mask estimation were determined through extensive simulations and chosen such that the receiver has a good performance in a single user scenario. Further, the degrees of freedom when choosing these parameters are limited by the default length of the preamble, which is 64 preamble symbols.

We perform a packet-based simulation, assuming that a total of N_{u} users are transmitting packets. The $N_{\mathrm{u}}-1$ interferers have the same PHY as the UOI. Every receiver generates packets according to the procedure described in Section 2.5.2. We further assume that every transmitter sends packets of the maximum allowable size of 1016 information bits, corresponding to 1208 transmitted bits after RS encoding.

For the simulations with MUI, we consider two different scenarios. In scenario A, all the devices use the same preamble code. In scenario B, the user of interest uses preamble code 5 and the other users use preamble code 6. In both scenarios, all users generate packets at the same rate R. We distinguish a high traffic case where $R = 200$ packets/s, and a low traffic case

T	M	B	N_{G}	N_{V}	N_{CH}	$P_{\mathrm{AWGN}}^{\mathrm{FA}}$
8 ns	16	500 MHz	10	16	16	10^{-4}

Table 4.1: Main parameters of the energy-detection receiver

where $R = 10$ packets/s. Note that even in the high traffic case the network is not yet saturated [1] and that $R = 200$ packet/s ($R = 10$ packet/s, resp.) corresponds to an effective uncoded data rate of 241.6 kb/s (12.1 kb/s).

We use different metrics to assess the performance of the receiver, mainly the packet error rate (PER), the synchronization error rate (SER) and the bit error rate (BER). We assume a packet to be in error if it is missed during synchronization or if at least one bit is in error after RS decoding. The SER denotes the percentage of packets that are missed during synchronization. To calculate the BER, we only consider packets that were correctly acquired during synchronization. All confidence intervals shown are at the 95% level.

In all our simulations, we use the IEEE 802.15.4a residential non-line-of-sight (NLOS) channel model (CM2) [14]. Further, we define the signal to noise ratio (SNR) as $\mathrm{SNR} = \frac{E_p}{N_0}$ where E_p is the received energy *per pulse* (after the convolution of the pulse with the impulse response of the channel), and $N_0/2$ is the PSD of a zero mean, Gaussian noise process bandlimited to B. The PHY is simulated with an accuracy of 100 ps resulting in a *simulation* sampling frequency of 10 GHz. As the simulation sampling frequency is larger than $2B$, the Gaussian noise samples are correlated. We use the algorithm in [168] to generate the correlated noise samples.

4.3.1 Without MUI the Receiver is Well-balanced

In Figure 4.4, we show the PER and the BER obtained for a single user, without MUI. We also plot the percentage of packets that are missed by the synchronization procedure (SER). These curves will on the one hand serve as reference curves for the simulations with MUI. On the other hand, they also show that the synchronization procedure as well as the parameters given in Table 4.1 are appropriately chosen such that the receiver performance is well balanced between synchronization and the decoding of the data bits.

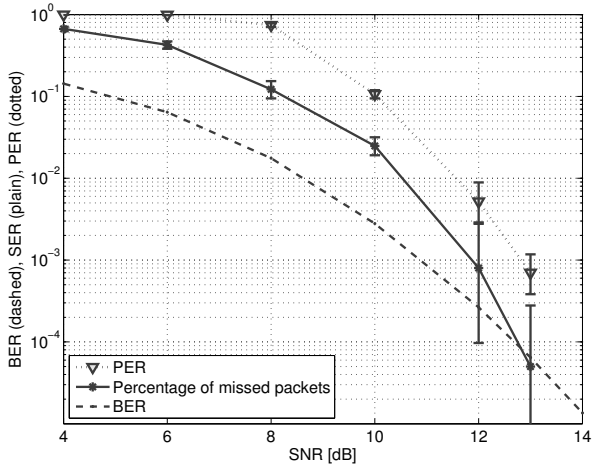

Figure 4.4: PER, BER, and percentage of packets missed (SER) due to synchronization for a single user, on a multipath channel, without MUI.

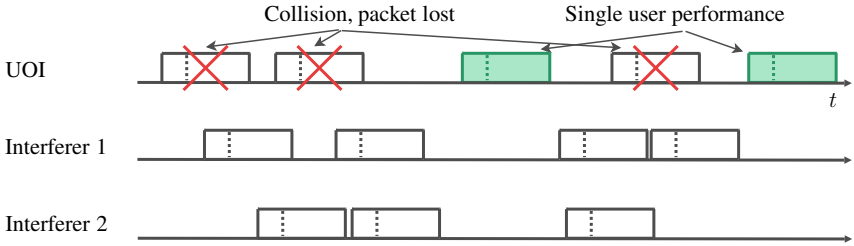

Figure 4.5: Illustration of the "Destructive Collisions" model. Packets experience single user performance if the channel is free from interference for the entire packet duration, otherwise packets are lost.

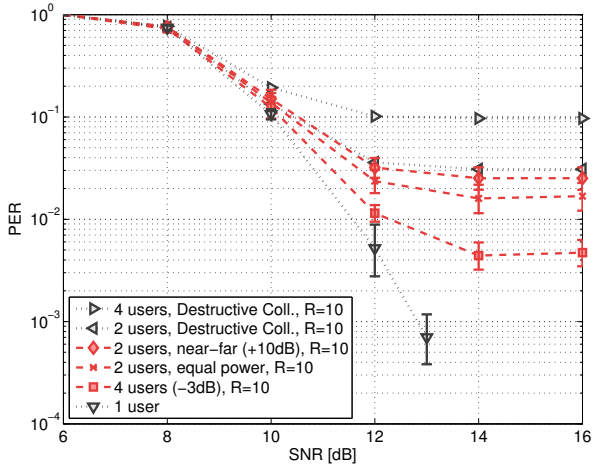

Figure 4.6: PER with MUI and different preamble codes (scenario B) in a low traffic case ($R = 10$ packet/s). Dashed lines from top to bottom: simulation results for two users, where the received power of the interferer is +10dB (near-far setting) or 0dB (equal power) with respect to the power of the UOI; four users where the received power of the interferers is -3 dB. Dotted lines from top to bottom: "Destructive Collisions" models for 4 and 2 users; single user performance. The performance in the near-far and equal power setting is close to the worst case given by the "Destructive Collisions" model.

4.3.2 Assessing Robustness to MUI: The "Destructive Collisions" Model

To assess the robustness of a receiver with respect to MUI, we compare its performance under MUI with the performance of a hypothetical model of *destructive collisions*. The "Destructive Collisions" model represents the worst case performance of a receiver that shows no capture effect at all. It thus corresponds to a typical model of the behavior of a narrowband network. The "Destructive Collisions" model is illustrated in Figure 4.5. Under this model, a packet is lost whenever there is more than one active transmission at the same time. If there is only one active transmission, the corresponding packet experiences single user performance.

4.3.3 With MUI Performance Approaches the Worst Case

Figure 4.6 shows the PER with MUI in scenario B (different preamble codes) with the low traffic rate of $R = 10$ packet/s. With $R = 200$ packets/s results are similar and therefore

1. Saturation occurs at $R = 226$ packets/s, see Section 2.5.2.

not shown. The single user performance from Figure 4.4 serves as a reference. We show the performance with one single interferer ($N_{\mathrm{u}} = 2$), once having a received power level equal to the UOI, once in a near-far setting where the power level of the interferer exceeds the power level of the UOI by 10 dB. Comparing these two curves with the corresponding curve obtained from the "Destructive Collisions" model shows that the energy-detection receiver is not robust against MUI. Indeed, in the near-far setting, the obtained performance is within the confidence interval of the worst case performance given by the "Destructive Collisions" model. In this case the receiver shows no capture effect at all. Even in the case of a single interferer with a power equal to the UOI, the capture effect is very limited, and performance is close to the worst case.

We also show results with three interferers ($N_{\mathrm{u}} = 4$) with a received power level 3 dB lower than the UOI. With this lower interference level, the receiver is finally able to handle concurrent transmissions, performing significantly better than the corresponding "Destructive Collisions" model.

Still, the analysis of the results with a single interferer demonstrate that with a conventional energy-detection receiver, concurrent transmissions in IEEE 802.15.4a are hardly possible and one of the main benefits of UWB over narrowband systems is therefore completely lost. To make things worse, this is the case even though the UOI and the interferers use different preamble codes that were meant to provide two separate logical channels.

In the following, we will try to understand what the reasons for this bad performance are. In particular, we hope to gain some insight into where the main weaknesses of the energy-detection receiver are and how it could be improved in order to make it more robust to MUI.

4.3.4 Taxonomy of Packet Errors and Interference Types

In a first step towards understanding the impact of MUI on the energy-detection receiver, we establish a taxonomy of all possible reasons for packet errors due to MUI. We find that every packet error must occur because of one of the following reasons:

1. synchronization error: the packet is missed during the synchronization phase.

   (a) *missed detection* (MD) during timing acquisition: the receiver tries to acquire timing of the packet but fails to do so.

   (b) *false alarm* (FA) during timing acquisition: the receiver is not trying to acquire timing because it wrongly assumes that timing has already successfully been acquired.

(c) MD in the SFD detection phase: the receiver acquires timing correctly but never detects the SFD.

(d) FA in the SFD detection phase: the receiver acquires timing correctly but declares detection of the SFD at the wrong time instant.

2. decoding error: the packet is received with more errors than the RS code is able to correct.

Further, the signalling format in IEEE 802.15.4a changes between the preamble and the payload (see Section 2.4). We can therefore additionally distinguish between the following interference types:

1. interference from an interfering preamble
2. interference from an interfering payload

4.3.5 Influence of the Preamble Code

Until now we have assumed that the interferers and the UOI use different preamble codes (scenario B). Now, if the interferers and the UOI were to use the same preamble code (scenario A), we would expect a much lower performance because of two reasons:

1. the preambles of the UOI and the interferers are indistinguishable. The receiver will therefore often synchronize on an interferer, leading to a lot of missed packets because of FAs.
2. the preamble in IEEE 802.15.4a determines the time-hopping sequence (THS) used in the payload. Identical preambles thus lead to an identical THS, which can lead to catastrophic collisions in the payload.

On the other hand, we have already seen that even in scenario B, the energy-detection receiver approaches worst case performance in terms of PER in the presence of a single interferer with a power level equal or higher than the UOI. Comparing the PER in scenarios A and B can therefore only lead to minor differences, which is verified in Figure 4.7 for an equal power interferer and both the low and high traffic case. Looking at the corresponding SER curves shown in Figure 4.8 reveals that indeed many more packets are misses during synchronization in scenario A than in scenario B. Not surprisingly, in scenario A, 98.4% of these packets are missed due to FAs, which is shown in Table 4.2. Why a lower rate of missed packets does not translate into a major performance improvement in terms of PER in scenario B can be seen

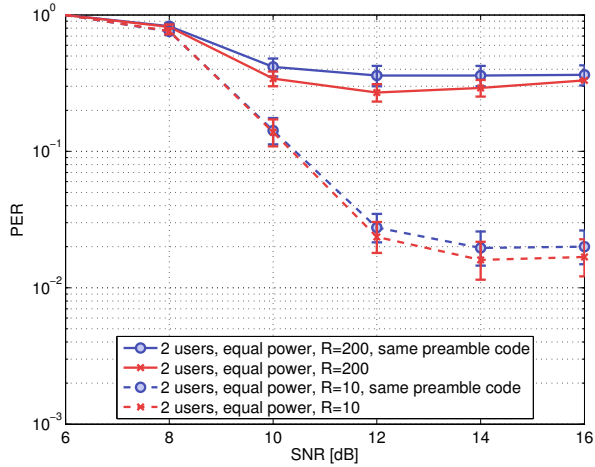

Figure 4.7: Comparison of the PER when two users use the same preamble code or different preamble codes. The two users have equal power at the receiver. We show a high ($R = 200$ packet/s) and a low traffic case ($R = 10$ packet/s). There is a negligible difference whether we use the same preamble code or a different preamble code.

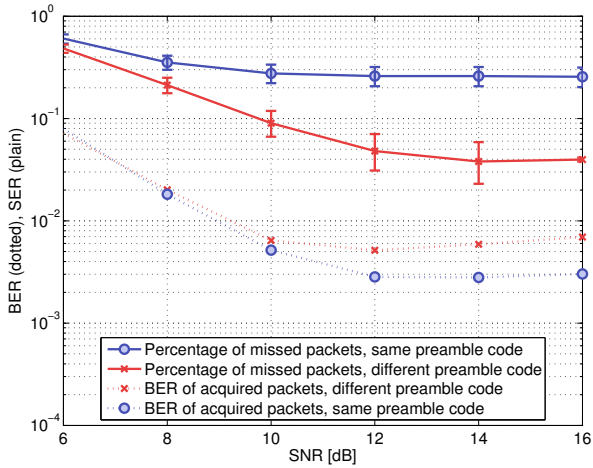

Figure 4.8: SER and BER (dashed lines) when two users use the same or a different preamble code. We consider the high traffic case ($R = 200$ packets/s). If different codes are being used, there are less packets missed. However, the packets additionally acquired, have generally more interference, which translates into a higher BER. The effect leading to the unusual trend of a slightly degrading BER performance with increasing SNR is explained in Section 4.3.7.

	Equal power		Near-far	
	Same Code	Diff. Code	Same Code	Diff. Code
Percentage of missed packets	25.6%	4.0%	34.6%	33.3%
Out of which missed due to FA	98.4%	63.8%	94.0%	19.1%

Table 4.2: Percentage of packets missed due to synchronization errors at a packet arrival rate of $R = 200\ \mathrm{packet/s}$, an SNR of 16 dB and one single interferer in the equal power or near-far setting. We compare scenario A where the interferer and the UOI share the same preamble code with scenario B where the preamble codes differ.

from the BER of packets that are correctly acquired, which is also shown in Figure 4.8. We notice that the acquired packets generally have more errors in scenario B than scenario A. We conclude that the packets additionally acquired by the receiver in scenario B are packets with interference and most of them are lost due to decoding errors.

We could therefore be tempted to conclude that the main reason for packet errors are decoding errors in the payload. However, a similar analysis for a near-far setting in scenario B shows that this is not true. With one strong interferer, the percentage of packets lost due to synchronization increases from 4% to 33.3% (Table 4.2). This leads us to our first insight:

Insight 1. *Both synchronization and payload decoding are significantly affected by MUI.*

A robust receiver will therefore have to address synchronization as well as payload decoding in order to achieve a noticeable performance gain with respect to the receiver presented here.

4.3.6 Synchronization Errors in Scenario With Different Preamble Codes

Although the origin of synchronization errors in scenario A may seem obvious, the situation is less clear in scenario B and further analysis is required. Figure 4.9 shows the classification of packet errors in scenario B according to the taxonomy introduced in Section 4.3.4. The numbers shown are for the interference limited case at an SNR of 16 dB with one equal power interferer and at a rate of $R = 200$ packets/s. The corresponding near-far setting is shown in Figure 4.10. We have already seen in the last section that both synchronization and decoding contribute substantially to packet errors. Looking at synchronization errors alone, we immediately see that the same is true for errors due to timing acquisition and SFD detection:

Insight 2. *Both timing acquisition and SFD detection are significantly affected by MUI.*

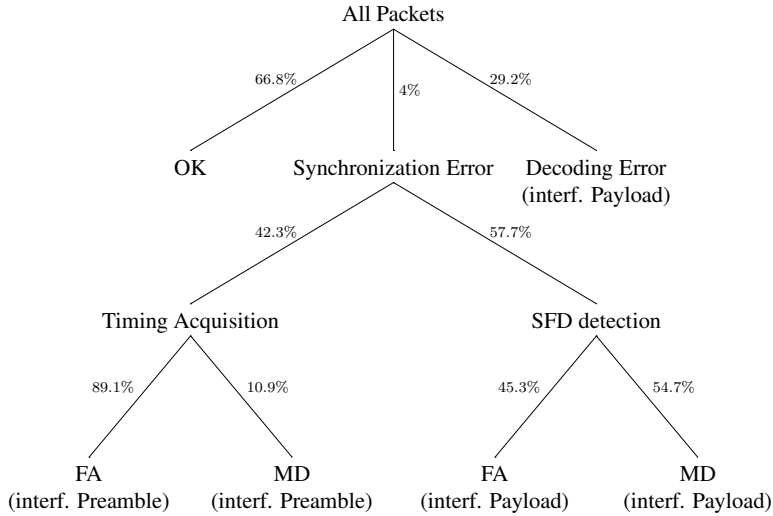

Figure 4.9: Classification of packet errors for one equal power interferer with different preamble code, at rate $R = 200$ packets/s and at an SNR of 16 dB. In parentheses, main interference type causing the error.

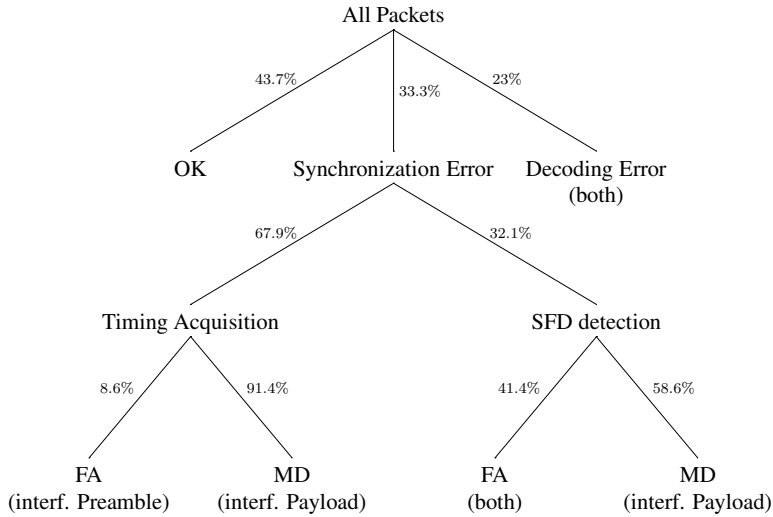

Figure 4.10: Same classification as in Figure 4.9 but for a strong near-far interferer.

We also see that in the equal power setting more errors occur during SFD detection, whereas in the near-far setting it is the other way round. Due to the stronger interference in the near-far setting, errors occur earlier in the packet reception process, already during packet detection and timing acquisition. We explain this by the fact that a strong interfering signal has a high likelihood of exceeding the detection threshold (that is solely based on the noise level, see Section 4.2.1), even if it is not perfectly aligned with the correlation template. This can generate MDs if the interfering signal introduces spurious maxima in the correlation output that make the verification fail. It can also lead to a FA with synchronization on an interfering signal if $N_{\mathrm{V}}+1$ consecutive maxima that are due to interference align such that the verification succeeds. Obviously, one would expect the former to be much more likely than the latter because it only requires producing one maximum in the correlation output at an arbitrary position.

We can observe this in the near-far setting where over 90% of timing acquisition errors (corresponding to over 60% of total synchronization errors) are due to MDs. We further find that 74.4% of these packets that are missed because of MDs in the timing acquisition phase have an overlap of more than 90% with an interfering payload. Due to the IEEE 802.15.4a signalling structure, an interfering payload burst in the simulated LPRF mode carries roughly four times the energy of an interfering preamble pulse. It is therefore highly likely that a small number of these high energy bursts is able to dominate the result of the correlation and make detection and verification of the periodic pattern induced by the UOI impossible. This is the major cause for timing acquisition errors in the near-far setting.

In the equal power setting, however, only relatively few packets (below 5% of total synchronization errors) are missed because of MDs in the timing acquisition phase. Further, the packets missed have only little overlap with interfering payloads, but surprisingly, 80.2% of them have an overlap of more than 90% with an interfering preamble. We conclude that unlike in the near-far setting, the interference level is not high enough for a few aligned bursts to change the result of the correlation. Another effect must play here because interfering preambles of lower energy cause most of the MDs.

What comes as an even bigger surprise is however the fact that almost 90% of timing acquisition errors are due to FAs in the equal power setting. Hence, it seems that even with different preamble codes, the receiver still synchronizes with the interferer often. A reason for this behavior becomes apparent if we look at the correlation between preamble codes shown in Figure 4.11(a). The figure shows the autocorrelation and cross-correlation properties of preamble codes 5 and 6 in the non-coherent case. Non-coherent here means that the code symbols are squared because of energy-detection, transforming the ternary preamble codes into binary

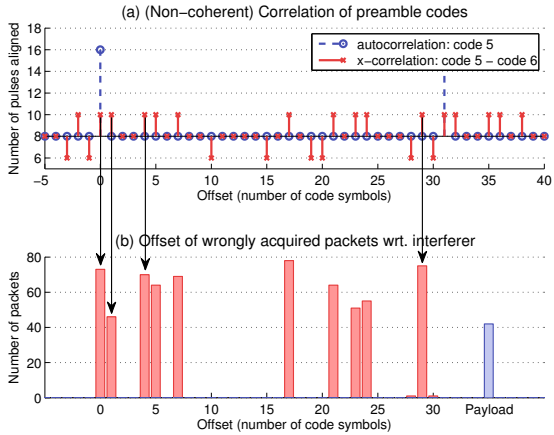

Figure 4.11: (a, top) Correlation of preamble code 5 with a periodic repetition of itself and with a periodic repetition of preamble code 6. (b, bottom) Classification of wrong timing acquisitions according to the offset with respect to the closest packet of an interferer. Correspondence of peaks in (a) and (b) suggests that wrong timing acquisitions are due to the correlation properties of the preamble codes.

ones. The autocorrelation is perfect in the sense that it only has a single peak per period of the code (of length $C = 31$. The peak occurs at offset 0, where all of the $C_{\text{NZ}} = 16$ nonzero code symbols of the preamble code are aligned. The cross-correlation of codes 5 and 6 on the other hand shows 10 peaks per period of the code. Each of the peaks corresponds to 10 code symbols that are aligned despite the fact that the codes differ. As a result, with non-coherent energy-detection, distinct peaks can be observed in the correlation output between the receiver template and an interfering preamble, even if different preamble codes are employed [2]. Further, these peaks are very likely to exceed the detection threshold and they occur at the same periodicity as the peaks in the autocorrelation. They are therefore the likely cause of the observed FAs. To verify this hypothesis, we classify FAs during timing acquisition according to whether they occur in the presence of an interfering preamble or an interfering payload. In the former case we further classify them according to the offset with respect to the beginning of the interfering preamble symbol closest in time. The result is shown in Figure 4.11(b). We have a perfect correspondence between the offsets of the cross-correlation peaks and the offsets of FAs, which verifies our hypothesis. We verified that the same is also true for the FAs during timing acquisition in the near-far setting. Further, this imperfect cross-correlation pattern also explains why an interfering preamble is able to cause MDs in the equal power setting whereas an interfering payload is not.

Insight 3. *During timing acquisition, we have two different effects. First, interfering preambles cause a lot of FAs and MDs because of the imperfect cross-correlation properties of the preamble codes with non-coherent reception. Second, interfering payloads may cause a lot of MDs if a few bursts dominate the output of the correlation.*

For SFD detection we see from Figure 4.9 that FA's and MD's roughly contribute equally to the errors in the equal power setting. The bipolar SFD detection template (4.18) has an equal number of $+1$s and -1s. Large interference terms are therefore equally likely to increase the correlation with the template leading to FAs or decreasing it leading to MDs.

Most of the FAs, namely 74.7%, occur because the SFD is declared 6 preamble symbols too early. This corresponds exactly to the position of the secondary peak of the SFD correlation (see Figure 4.3). Further, we confirm that FAs and MDs happen under strong interference: for MDs, 95% of the packets suffer from interference from a payload during more than 90% of the SFD part; for FAs this is the case for 93.9% of the packets.

2. Note that these peaks are not an anomaly of the IEEE 802.15.4a codes per se, but rather inherent to such sequences. See also Section 6.2.3

In the near-far setting, the situation is slightly different. Now, only 12.6% of FAs occur at the secondary peak of the SFD correlation. The others occur at arbitrary positions. We attribute this to the fact that due to the higher interference level, the SFD threshold can be exceeded even if the SFD sequence of the UOI is not present at all. 45.4% of the FA packets have an overlap of more than 90% with an interfering preamble, 48.6% with an interfering payload. In the near-far setting, we can thus attribute FAs during SFD detection to both interference types.

Insight 4. *Large interference terms can dominate the correlation with the bipolar SFD template causing both FAs and MAs. Many of the FAs occur at the secondary peak of the SFD correlation.*

For MDs during SFD detection, 41.5% of the packets suffer from interference from a payload during more than 90% of the SFD part, 16.5% of the packets suffer from interference from a preamble during more than 90% of the SFD part. The remaining packets suffer from interference with a payload during channel mask estimation. We conclude that in the near-far setting, roughly 40% of the MDs during SFD detection occur because the SFD threshold is overestimated because of interference with a payload.

Insight 5. *The estimation of the SFD threshold is not robust to strong MUI, leading to MDs.*

4.3.7 Analysis of Payload Decoding Errors

In this last section of the performance evaluation, we analyze errors that are due to decoding of the payload. We can see from Figures 4.9 and 4.10 that this corresponds to a significant percentage of the erroneous packets.

First, we want to understand to what extent decoding errors are due to interference during decoding itself, and to what extent they occur because of interference during channel mask estimation. To this end we ran simulations where we switched off MUI during channel mask estimation. Moreover, for these simulations we assumed a hypothetical receiver with perfect synchronization, i.e., where an oracle returns the exact beginning of the payload. Figure 4.12 shows the PER obtained in the equal power setting; in the near-far setting the results are similar. As a reference, the corresponding curves of the full simulation are shown as well. We can see that perfect synchronization and channel mask estimation without MUI only yield a small improvement in PER. We conclude that the main source of errors in payload decoding is interference during the payload and not an inaccurately estimated channel mask.

It can be observed in Figure 4.12 (and also in some of the performance figures we showed earlier) that the error-floor of the PER (and also of the BER, e.g., in Figure 4.8) shows an

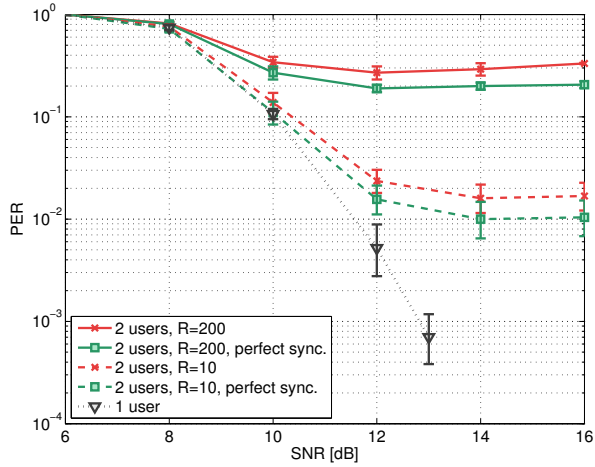

Figure 4.12: Comparison of the PER with a hypothetical receiver performing perfect synchronization and experiencing no MUI during channel mask estimation. The hypothetical receiver only shows a slight performance improvement. The curves shown are for the equal power setting.

	Avg. number of ones in channel mask					
	6 dB	8 dB	10 dB	12 dB	14 dB	16 dB
Correct packets	1.00	3.21	4.43	5.11	5.85	6.56
Erroneous packets	2.65	3.69	4.43	5.50	6.34	7.03

Table 4.3: Average number of ones in channel mask. Numbers shown are for the hypothetical receiver without MUI during channel mask estimation and with perfect synchronization. Simulations were performed with one equal power interferer and a rate of $R = 200$ packet/s)

increasing trend at high SNR when the signal becomes interference limited. This can be explained by Table 4.3 where the average number of ones in the channel mask is shown for different SNRs. The numbers shown are for the hypothetical receiver and for the high traffic case. The number of ones in the channel mask increases proportionally with the SNR. A large number of ones in the channel mask implies a higher likelihood of suffering from MUI as we integrate a larger amount of the received signal, which explains the increasing PER at higher SNRs. We can also see that packets with decoding errors are generally those that have a longer overall integration time.

Insight 6. *Limiting the integration time to signal parts containing UOI contributions and employing a time-hopping sequence are not sufficient to prevent a performance degradation due to MUI.*

In IEEE 802.15.4a, the RS code operates on blocks of 378 bits. Further, no interleaving is used. A packet is therefore in error if one of the RS blocks of 378 consecutive bits has more errors than the code can correct. One block corresponds to roughly 30% of the maximum payload size we simulated and, indeed, we find that 98.6% of the erroneous packets in the equal power setting are subject to interference during at least 30% of the payload. Further, 100% of these packets have some overlap with an interfering payload and in 91.8% this overlap exceeds 30%. 43.3% have no interference with an interfering preamble at all. In the near-far setting, interference from interfering preambles starts to get strong enough to provoke errors on its own: roughly 10% of decoding errors are solely due to interfering preambles. Still, the above results suggest that most of the packet errors are due to an interfering payload. This also makes sense from the perspective that a burst in the payload contains four times more energy than a pulse of the preamble. Furthermore, the payload is also about four times longer than the preamble, making a collision with a payload more likely.

4.4 Conclusion

We have evaluated and analyzed the effect of MUI on a complete IEEE 802.15.4a system that uses an energy-detection receiver. Although some papers hint at the fact that an energy-detection receiver with a long integration time might be vulnerable to MUI, we are not aware of any work prior to this one that evaluates and quantifies this effect. Our analysis shows that a simple energy-detection receiver is not only vulnerable to MUI, but that already at low traffic, it even exhibits close to worst case performance if its subject to uncontrolled MUI. This

makes one of the most appealing benefits of IR-UWB, namely its ability to support concurrent transmissions completely void. Further, we show that all receiver operations needed to receive a packet are severely affected by MUI.

Although this may seem to be a rather dim prospect for energy-detection in the scope of IEEE 802.15.4a, we also compile a catalog of insights into the main reasons for this bad performance. These insights show where there is room for improvement and they will serve as a guideline for the design of more robust receivers in the subsequent chapters.

4.5 Acknowledgments

We would like to thank Armin Wellig and Julien Zory from STMicroelectronics for all the discussions and comments that lead to the work forming this chapter.

Chapter 5

Robust IEEE 802.15.4a Energy-Detection Receiver Architecture With Adaptive Thresholding

In the previous chapter we have seen that the performance of an IEEE 802.15.4a compliant energy-detection receiver is severely degraded in the presence of MUI, resulting in a performance that is close to the worst case. We have also seen that both synchronization and data decoding are equally affected. In this chapter we focus on data decoding and try to answer the question of whether there are ways to improve upon the discouraging results found in Chapter 4.

As we have seen, limiting the total integration time and employing time-hopping are not enough to contain the effect of MUI at reasonable levels, even though such a receiver shows a good performance if no MUI is present. This also indicates that the situation may be similar to coherent reception, where receivers that were designed for the Gaussian single user case are known to perform badly in an impulsive, non-Gaussian multi-user environment (see Section 3.1 in the related work). Among the simplest interference mitigation techniques known to mitigate impulsive interference in coherent receivers are thresholding schemes (Section 3.1.1) that limit the contribution of high interference terms. If such a scheme could be adapted to energy-detection receivers, it could prove to be a viable solution because thresholding usually has a low implementation complexity.

We will see in this chapter that such an adaptation of thresholding to energy-detection receivers is indeed feasible at a moderate complexity increase. Our solution is fully adaptive to different channel conditions and results in a novel receiver architecture that takes the specific

signalling structure of IEEE 802.15.4a signals into account. In certain scenarios with MUI we find the packet error rate to be up to two orders of magnitude lower when compared to a classical energy-detection receiver designed without accounting for MUI.

This chapter is organized as follows: The architecture of the robust receiver is derived in Section 5.1. The adaptive thresholding scheme to mitigate MUI, as well as algorithms for robust parameter estimation are introduced in Section 5.2. The performance of the complete receiver is evaluated in Section 5.3. We conclude the chapter in Section 5.4.

5.1 Architecture of the Robust Receiver

The baseline receiver model used throughout this chapter is the same as in Chapter 4, i.e., we assume an IEEE 802.15.4a PHY and an energy-detection receiver with a sampling rate $1/T$. Further, the receiver produces $M = L_\mathrm{s}T_\mathrm{c}/T$ discrete samples per preamble code symbol. In contrast to Chapter 4, we assume perfect synchronization, i.e., the TOA as well as its estimate are assumed to be zero ($\nu_0 = \hat{\nu}_0 = 0$) throughout this chapter. Robust synchronization and timing acquisition are explained separately in Chapter 6.

For convenience, we write the received signal during the preamble as

$$y^{\text{pre}}_{m,j+iC} = y^{\text{pre}}_{m+(j+iC)M} = \int_{mT+(j+iC)L_\mathrm{s}T_\mathrm{c}}^{(m+1)T+(j+iC)L_\mathrm{s}T_\mathrm{c}} [r_{\text{pre}}(t)]^2 \,\mathrm{d}t, \tag{5.1}$$

where we introduced $y^{\text{pre}}_{m,j+iC}$ to denote the m-th sample of the j-th code symbol of the i-th preamble symbol. Further, $r_{\text{pre}}(t)$ is given by (2.28) and $m = 0, \ldots, M-1$.

To ease notation, but without loss of generality, we only consider the reception of the first payload symbol. For the same reason we also assume that the time-hopping index of the first payload symbol is equal to $c_{\text{THS},0} = 0$ and we neglect the BPSK modulation a_0 that does not matter in non-coherent reception. The received signal given by (2.30), and where we dropped index i that refers to the payload symbol, is then

$$r_{\text{pay}}(t) = \underbrace{\sum_{j=0}^{N_{\text{cpb}}-1} b_j \cdot \tilde{h}(t - d_0 T_\mathrm{f}/2 - jT_\mathrm{c})}_{\tilde{x}_{\text{pay}}(t)} + v(t) \tag{5.2}$$

where d_0 is the BPPM data bit of the first symbol and b_j are the elements of the scrambling sequence during the first symbol.

As was the case in Chapter 4, the receiver obtains M samples per BPPM block of the frame, yielding the signal vector $\mathbf{y} = (\mathbf{y^0}, \mathbf{y^1}) = (y_0, \ldots, y_{M-1}, y_{N_\mathrm{f}/2}, \ldots, y_{N_\mathrm{f}/2+M-1})$. The m-th received sample of the first symbol of the payload is given by (2.17), i.e.,

$$y_m = \int_{mT}^{(m+1)T} [r_{\text{pay}}(t)]^2 \, \mathrm{d}t \tag{5.3}$$

5.1.1 Optimal Decision Rule for Burst Transmissions

In the last chapter we have seen that in order to be robust to MUI, it is not enough for an IEEE 802.15.4a compliant energy-detection receiver to limit the integration time through, e.g., a binary channel mask (Chapter 4, Insight 6). In a first step towards a more robust receiver architecture, we derive the optimum decision rule for the energy-detection receiver in the case where the noise term $v(t)$ in (5.2) consists of AWGN only. We will later see that the optimal decision rule leads to a robust receiver that uses an adaptive thresholding mechanism to mitigate the effect of MUI.

According to Section 2.3.3, the samples y_m can be assumed to be distributed independently and according to a scaled noncentral chi-square distribution with $2BT$ degrees of freedom and non-centrality parameter $\zeta_{m,d_0} = \frac{p_{m,d_0}}{N_0/2}$ (2.22), where we explicitly stated the dependence of the non-centrality parameter on d_0. Plugging $\tilde{x}_{\text{pay}}(t)$ from (5.2) into (2.23) and assuming no ISI yields

$$p_{m,d_0} = \int_{mT}^{(m+1)T} \left[\sum_{j=0}^{N_{\text{cpb}}-1} b_j \cdot \tilde{h}(t - d_0 T_\mathrm{f}/2 - jT_\mathrm{c}) \right]^2 dt \tag{5.4}$$

Since $p_{m,1} = p_{m-N_\mathrm{f}/2,0}$, we can simplify the notation by introducing $q_m \doteq p_{m,0}$.

We observe that with T_f large enough to prevent ISI (in IEEE 802.15.4a this is achieved with proper guard intervals), the contribution of the UOI is confined to the first half of the samples of $\mathbf{y}$ if $d_0 = 0$ and to the second half if $d_0 = 1$. It follows that if $d_0 = 0$, $\mathbf{y^0}$ is distributed according to the non-central chi-square distribution and $\mathbf{y^1}$ according to a central chi-square distribution with $2BT$ degrees of freedom (2.24) (and vice versa if $d_0 = 1$). Consequently, the optimal decision rule according to the maximum likelihood criterion is found as

$$\sum_{m=0}^{M-1} \mathrm{LLR}(y_m | N_0/2, 2BT, q_m) \underset{d_0=1}{\overset{d_0=0}{\gtrless}} \sum_{m=0}^{M-1} \mathrm{LLR}(y_{m+N_\mathrm{f}/2} | N_0/2, 2BT, q_m) \tag{5.5}$$

with the log-likelihood ratio (LLR) given by

$$\begin{aligned}\mathrm{LLR}(y_m|N_0/2, 2BT, q_m) &= \ln\left[\frac{f_{NC\chi^2}\left(\frac{y_m}{N_0/2}\middle|2BT, \frac{q_m}{N_0/2}\right)}{f_{\chi^2}\left(\frac{y_m}{N_0/2}\middle|2BT\right)}\right] \\ &= \ln\left[{}_0F_1\left(;BT;\frac{q_m y_m}{N_0^2}\right)\right] - \frac{q_m}{N_0} \end{aligned} \tag{5.6}$$

The confluent hypergeometric limit function ${}_0F_1\left(;BT;\frac{q_m y_m}{N_0^2}\right)$ is given by

$${}_0F_1\left(;BT;\frac{q_m y_m}{N_0^2}\right) = 2^{BT-1}\Gamma(BT)\Big(\frac{N_0/2}{\sqrt{y_m q_m}}\Big)^{BT-1} I_{BT-1}(\frac{\sqrt{q_m y_m}}{N_0/2}) \tag{5.7}$$

where $I_\upsilon(z)$ is the υ-th order modified Bessel function of the first kind.

Plugging (5.6) and (5.7) into (5.5) yields

$$\sum_{m=0}^{M-1} \ln\left[\frac{I_{BT-1}\left(\frac{\sqrt{y_m q_m}}{N_0/2}\right)}{\sqrt{y_m^{BT-1}}}\right] \underset{d_0=1}{\overset{d_0=0}{\gtrless}} \sum_{m=0}^{M-1} \ln\left[\frac{I_{BT-1}\left(\frac{\sqrt{y_{m+N_\mathrm{f}/2} q_m}}{N_0/2}\right)}{\sqrt{y_m^{BT-1}}}\right] \tag{5.8}$$

Finally, (5.8) can be linearly approximated [33] resulting in a simpler decision rule

$$\sum_{m=0}^{M-1} y_m \cdot q_m \underset{d_0=1}{\overset{d_0=0}{\gtrless}} \sum_{m=0}^{M-1} y_{m+N_\mathrm{f}/2} \cdot q_m \tag{5.9}$$

Hence, the optimal detector applies a weighting function with coefficients q_m prior to comparing the energies in the 0-block and the 1-block of a BPPM frame. With $N_{\text{cpb}} > 1$, (5.4)-(5.9) give us a generalization of the result found in [33]. With $N_{\text{cpb}} = 1$ the weighting function reduces to the one found in [33]. How to estimate the weights q_m is shown in the next section.

5.1.2 Estimation of Weighting Coefficients and Noise Power Spectral Density

We have seen in the previous section that in order to optimally demodulate the received symbol, the receiver needs to estimate the weights q_m. We shall see later, that estimates of q_m are also needed in order to effectively mitigate the impact of MUI. Further, mitigation of MUI also requires an estimate of the noise PSD $N_0/2$.

To show how the weighting coefficients q_m can be estimated from the preamble during a

channel estimation phase, we rewrite (5.4)

$$\begin{aligned} q_m &= \sum_{j=0}^{N_{\text{cpb}}-1} b_j^2 \int_{mT}^{(m+1)T} \tilde{h}^2(t - jT_{\text{c}})dt \\ &+ 2 \sum_{j=0}^{N_{\text{cpb}}-1} \sum_{k=j+1}^{N_{\text{cpb}}-1} b_j b_k \int_{mT}^{(m+1)T} \tilde{h}(t - jT_{\text{c}})\tilde{h}(t - kT_{\text{c}})dt. \end{aligned} \tag{5.10}$$

By applying the change of variables $s \doteq t - jT_{\text{c}}$ and introducing $K \doteq T_{\text{c}}/T$, we obtain

$$\begin{aligned} q_m &= \sum_{j=0}^{N_{\text{cpb}}-1} \int_{(m-jK)T}^{(m-jK+1)T} \tilde{h}^2(s)ds \\ &+ 2 \sum_{j=0}^{N_{\text{cpb}}-1} \sum_{k=j+1}^{N_{\text{cpb}}-1} b_j b_k \int_{(m-jK)T}^{(m-jK+1)T} \tilde{h}(s)\tilde{h}(s - (k-j)T_{\text{c}})ds \\ &= \sum_{j=0}^{N_{\text{cpb}}-1} w_{m-jK}^{(0)} + 2 \cdot \sum_{j=0}^{N_{\text{cpb}}-1} \sum_{k=j+1}^{N_{\text{cpb}}-1} b_j b_k \cdot w_{m-jK}^{(k-j)} \end{aligned} \tag{5.11}$$

where $w_m^{(l)}, l = 1, \ldots, N_{\text{cpb}} - 1$ is given by

$$w_m^{(l)} = \int_{mT}^{(m+1)T} \tilde{h}(t)\tilde{h}(t - l \cdot T_{\text{c}})dt \tag{5.12}$$

To estimate q_m, we thus have to estimate the parameters $w_m^{(l)}$. Note that $w_m^{(0)}$ represents the energy-delay profile of the channel and corresponds to the weight applied in [33].

Assuming $v(t)$ is AWGN only, the expected value of a discrete preamble sample is given by

$$\begin{aligned} \mathbb{E}[y_{m,j+iC}^{\text{pre}}] &= \mathbb{E}\left[\int_{mT+(j+iC)L_{\text{s}}T_{\text{c}}}^{(m+1)T+(j+iC)L_{\text{s}}T_{\text{c}}} [\sum_{i=0}^{N_{\text{pre}}-1} s_i \sum_{j=0}^{C-1} c_j \cdot \tilde{h}(t - (j+iC)L_{\text{s}}T_{\text{c}}) + v(t)]^2 \, \mathrm{d}t \right] \\ &= s_i^2 c_j^2 \underbrace{\int_{mT}^{(m+1)T} \tilde{h}^2(t)dt}_{w_m^{(0)}} + \underbrace{\mathbb{E}\left[\int_0^T v^2(t)dt \right]}_{\bar{v}} \end{aligned} \tag{5.13}$$

where the first equality is obtained by plugging (2.28) into (5.1). In the second equality we used the fact that $v(t)$ is stationary and zero-mean if it consists of AWGN only. Further, we assumed that there is no inter-pulse interference between consecutive pulses of the preamble

signal.

Because both sequences, s_i and c_j, are known, equation (5.13) suggests that we can get an estimate $\hat{\bar{v}}$ of the noise energy $\bar{v}$ by averaging samples $y^{\text{pre}}_{m,j+iC}$ for which $c_j = 0$:

$$\hat{\bar{v}} = \frac{1}{N_{\text{CH}}(C - C_{\text{NZ}})M} \sum_{i=0}^{N_{\text{CH}}-1} \sum_{j=0}^{C-1} \sum_{m=0}^{M-1} \delta_{c_j} \cdot y^{\text{pre}}_{m,j+iC} \tag{5.14}$$

N_{CH} are number of preamble symbols that are used for channel estimation, C_{NZ} are the number of nonzero code symbols of the preamble code and δ_m denotes the Kronecker delta.

The PSD $N_0/2$ of the noise process can then be estimated from $\hat{\bar{v}}$. Given, that for $c_j = 0$, $y^{\text{pre}}_{m,j+iC}$ is distributed according to the chi-square distribution with $2BT$ degrees of freedom (see equation (2.24)), we have that

$$\frac{\hat{N}_0}{2} = \frac{\hat{\bar{v}}}{2BT} \tag{5.15}$$

Further, an estimate $\hat{w}^{(0)}_m$ of $w^{(0)}_m$ can be obtained by averaging over samples $y^{\text{pre}}_{m,j+iC}$ of the SYNC part of the preamble for which $c_j \neq 0$ [1] and by subtracting $\hat{\bar{v}}$ from the result:

$$\hat{w}^{(0)}_m = \frac{1}{N_{\text{CH}} C_{\text{NZ}}} \left[\sum_{i=0}^{N_{\text{CH}}-1} \sum_{j=0}^{C-1} c_j^2 \cdot y^{\text{pre}}_{m,j+iC} \right] - \hat{\bar{v}} \tag{5.16}$$

Estimating $w^{(0)}_m$ as well as the noise PSD $N_0/2$ is all that has to be done during the channel estimation phase if there are no bursts, i.e., $N_{\text{cpb}} = 1$. However, if $N_{\text{cpb}} > 1$ we additionally need to estimate the parameters $w^{(l)}_m$, with $l \in \{1, \ldots, N_{\text{cpb}}-1\}$, in order to calculate the optimal weights given by (5.11). Unfortunately, there is no way for $w^{(l)}_m$, $l \neq 0$, to be estimated by a classical energy-detection receiver architecture. Consequently, equations (5.11) and (5.12) not only define a new weighting function but also show the necessity for a new receiver structure that allows for the estimation of the parameters $w^{(l)}_m$.

A possible receiver structure that overcomes the limitations of the classical energy-detection receiver is depicted in Figure 5.1. With respect to a classical energy-detection receiver, this new receiver employs $N_{\text{cpb}} - 1$ additional branches. The l-th additional branch delays the received signal by lT_{c} and multiplies the received signal with this delayed version. The resulting signal is then integrated and sampled to yield the samples $y^{\text{pre},(l)}_{m,j+iC}$. However, the additional branches are only needed during channel estimation in the preamble, in order to estimate the parameters $w^{(l)}_m$. During the other phases of packet reception, synchronization and decoding, the additional

1. Remember that for all preamble symbols of the SYNC part we have $s_i = 1$, see Section 2.4.2.

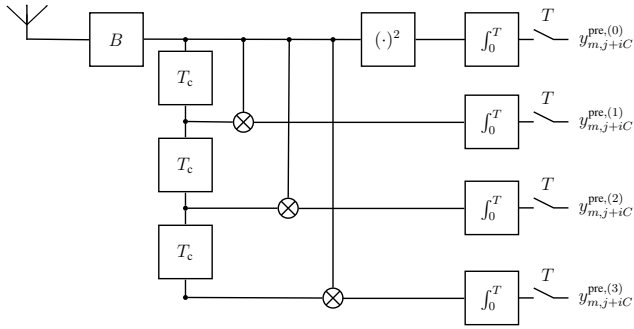

Figure 5.1: Proposed receiver structure in the case of payload signaling with bursts of four pulses ($N_{\text{cpb}} = 4$) and with an integration time of $T \leq T_{\text{c}}$. The additional branches are needed *only* during estimation of the parameters $w_m^{(l)}$ given in (5.12). They are not needed for data decoding where only the upper branch is required. Hence, their impact on power consumption is minimal.

circuitry is *not* used. The added complexity and power consumption should thus be moderate. This also limits the additional memory requirements of this more sophisticated receiver.

In an analogous manner to the classical receiver architecture, $\hat{w}_m^{(0)}$ can be obtained from the branch with no delay, $y^{\text{pre},(0)}_{m,j+iC}$, according to (5.16). The same is true for the noise PSD. $\hat{w}_m^{(l)}, l \in \{1, \ldots, N_{\text{cpb}} - 1\}$ can be obtained in a similar way from the l-th branch, thanks to the observation that

$$\mathbb{E}[y^{\text{pre},(l)}_{m,j+iC}] = s_i^2 c_j^2 \underbrace{\int_{mT}^{(m+1)T} \tilde{h}(t)\tilde{h}(t - l \cdot T_{\text{c}})dt}_{w_m^{(l)}} \quad , \quad l \in \{1, \ldots, N_{\text{cpb}} - 1\} \tag{5.17}$$

resulting in

$$\hat{w}_m^{(l)} \;=\; \frac{1}{N_{\text{CH}} C_{\text{NZ}}} \sum_{i=0}^{N_{\text{CH}}-1} \sum_{j=0}^{C-1} c_j^2 \cdot y^{\text{pre},(l)}_{m,j+iC} \quad , \quad l \in \{1, \ldots, N_{\text{cpb}} - 1\} \tag{5.18}$$

From the parameters $\hat{w}_m^{(l)}, l \in \{0, \ldots, N_{\text{cpb}} - 1\}$, an estimate $\hat{q}_m$ of the weights q_m can be directly calculated using (5.11) under the condition that $K = T_{\text{c}}/T$ is an integer greater than or equal to one, or in other words that $T \leq T_{\text{c}}$.

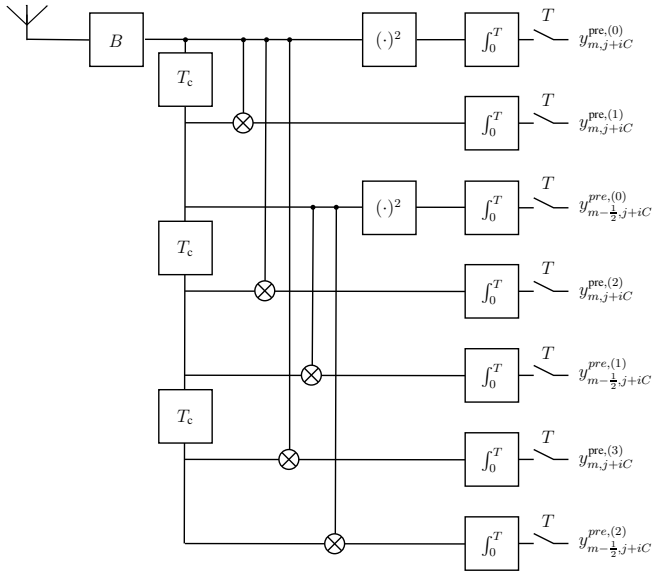

Figure 5.2: Receiver structure allowing to calculate the optimal weights q_m in the case of payload signaling with bursts of four pulses ($N_{\text{cpb}} = 4$) and with an integration time of $T = 2T_{\text{c}}$. As this integration time exceeds the duration of a chip T_{c}, three additional branches are needed with respect to Figure 5.1 in order to be able to estimate all of the needed parameters.

Trade-off Between Sampling Frequency and Hardware Complexity

For integration times that exceed the duration of a chip we cannot obtain all of the required parameter estimates from the receiver structure introduced above and shown in Figure 5.1. This is because with $K < 1$, some of the parameter estimates $\hat{w}^{(l)}_{m-jK}$, for which jK becomes a fraction rather than an integer, cannot be expressed through any of the parameter estimates given by (5.18). Nevertheless, a receiver structure that allows us to obtain all of the parameter estimates $\hat{w}^{(l)}_{m-jK}$ required to calculate q_m also exists in this case. It can be obtained from the new receiver structure introduced so far, by adding a maximum of $\frac{N_{\text{cpb}}(N_{\text{cpb}}-1)}{2}$ additional branches to calculate the missing parameter estimates. The resulting optimal structure for $N_{\text{cpb}} = 4$ and $T = 2T_{\text{c}}$ is shown in Figure 5.2 and its derivation is given as an example in the following. With an integration time of $T = 2T_{\text{c}}$, we have $K = 1/2$. We see from (5.11) that in this case we need to estimate the following parameters in order to calculate q_m: $w^{(0)}_m$, $w^{(0)}_{m-\frac{1}{2}}$, $w^{(0)}_{m-1}$,

$w^{(0)}_{m-\frac{3}{2}}$, $w^{(1)}_{m}$, $w^{(1)}_{m-\frac{1}{2}}$, $w^{(1)}_{m-1}$, $w^{(2)}_{m}$, $w^{(2)}_{m-\frac{1}{2}}$, $w^{(3)}_{m}$. Out of these ten parameters, six can be readily estimated according to (5.18), namely $w^{(0)}_{m}$, $w^{(0)}_{m-1}$, $w^{(1)}_{m}$, $w^{(1)}_{m-1}$, $w^{(2)}_{m}$ and $w^{(3)}_{m}$. The remaining four parameters $w^{(0)}_{m-\frac{1}{2}}$, $w^{(0)}_{m-\frac{3}{2}}$, $w^{(1)}_{m-\frac{1}{2}}$ and $w^{(2)}_{m-\frac{1}{2}}$ cannot be estimated by the receiver structure in Figure 5.1. However, we can see that an estimate of $w^{(0)}_{m-\frac{3}{2}} = w^{(0)}_{(m-1)-\frac{1}{2}}$ can be obtained from the estimate of $w^{(0)}_{m-\frac{1}{2}}$. We therefore need to add three additional branches to the receiver structure of Figure 5.1 in order to estimate the three remaining parameters $w^{(0)}_{m-\frac{1}{2}}$, $w^{(1)}_{m-\frac{1}{2}}$ and $w^{(2)}_{m-\frac{1}{2}}$, which leads to the optimal receiver structure shown in Figure 5.2.

For receivers that calculate the optimal weights q_m, there is thus a tradeoff between a lower sampling frequency and the additional circuitry that is needed if the sampling frequency is decreased.

5.2 Mitigation of MUI with an Adaptive Thresholding Mechanism

The decision rule derived in Section 5.1.1 is optimal if the weights q_m are known and if the interference and noise term $v(t)$ is AWGN. In the last section we have seen that the first condition can be fulfilled since the weights q_m can be estimated. However, it is well-known that in the presence of MUI, the Gaussian assumption for $v(t)$ does generally not hold ([39] or see Section 3.1). Rather, MUI tends to be impulsive with a few strong interference terms of high energy. With the decision rule in (5.9) it could therefore happen that few samples y_m suffering from high interference dominate the decision leading to a decoding error, even though they may have a low associated weight q_m.

On the other hand, the estimated weights $\hat{q}_m$, combined with the knowledge about the distribution of the receiver output under AWGN, give us a statistical model for the distribution of samples y_m if no MUI is present. A large deviation of y_m from the model suggests that it is subject to MUI, and its contribution to the decision should be limited.

To detect a deviation from the AWGN model, the receiver can calculate the threshold

$$\eta_m = \frac{\hat{N}_0}{2} F^{-1}_{\mathrm{NC}\chi^2}(1 - P^{\mathrm{FA}}_{\mathrm{MUI}} | 2BT, \frac{\hat{q}_m}{\hat{N}_0/2}) \tag{5.19}$$

by inverting the CDF of the non-central chi-square distribution defining the AWGN model. The sensitivity of the threshold can be adjusted via the small false alarm probability $P^{\mathrm{FA}}_{\mathrm{MUI}}$. The threshold adapts to different channel conditions through the associated estimated weights $\hat{q}_m$.

In order to limit the contribution of high interference terms, the receiver applies a non-linearity governed by η_m to the received samples prior to the decision process. Different non-linear operations are possible. The one we found to work best is to set samples, above the threshold, to the value of the corresponding weight $\hat{q}_m$

$$g(y_m|\hat{q}_m, \hat{N}_0/2) = \begin{cases} y_m & \forall m : y_m \leq \eta_m \\ \hat{q}_m & \forall m : y_m > \eta_m \end{cases} \tag{5.20}$$

The resulting decision rule, including the adaptive threshold is given by

$$\sum_{m=0}^{M-1} g(y_m|\hat{q}_m, \hat{N}_0/2) \cdot \hat{q}_m \underset{d_0=1}{\overset{d_0=0}{\gtrless}} \sum_{m=0}^{M-1} g(y_{m+N_f/2}|\hat{q}_m, \hat{N}_0/2) \cdot \hat{q}_m \tag{5.21}$$

Through the threshold η_m, the non-linearity $g(y_m|\hat{q}_m, \hat{N}_0/2)$ depends on the weights $\hat{q}_m$ as well as on the estimated thermal noise level. We have seen in Section 5.1.2 how to estimate these quantities in AWGN. In the next section, we will show how to estimate them in a robust fashion under MUI.

5.2.1 Robust Parameter Estimation Using Order Statistics

Estimation of the parameters $\hat{w}_m^{(l)}$ and $N_0/2$ using sample averages according to (5.14), (5.15), (5.16) and (5.18) is not robust to MUI. The reason is similar to what we have observed in the previous section where we treated payload decoding: a few samples $y_{m,j+iC}^{\text{pre}}$ suffering from high interference may dominate the sample average and lead to a biased estimation. Unlike in payload decoding, during the estimation phase we do in general not yet have any knowledge about the expected signal level available. It is therefore not possible to devise a thresholding scheme to mitigate interference similar to the one proposed in the last section. Instead we will resort to order statistics. It is well-known that the median, e.g., is more robust to outliers than the mean. We propose to replace the sample mean in (5.14), (5.16) and (5.18) with the sample median, yielding e.g.,

$$\hat{w}_m^{(0)} = \text{median}\left\{y_{m,j+iC}^{\text{pre}} : (i,j) \in \{0,\ldots,N_{\text{CH}}-1\} \times \{0,\ldots,C-1\}, c_j \neq 0\right\} - \hat{\bar{v}} \tag{5.22}$$

The median of a chi-square distribution with κ degrees of freedom can be approximated by

$\kappa - 2/3$. Consequently, we calculate the estimate of the noise PSD according to

$$\frac{\hat{N}_0}{2} = \frac{\hat{\tilde{v}}}{2BT - 2/3} \tag{5.23}$$

if the sample median is used instead of the sample mean.

Performing parameter estimation with the sample median as explained above can still be vulnerable to interference from an interfering preamble. The reason is the imperfect cross-correlation of the preamble codes in IEEE 802.15.4a if non-coherent reception is used (see Insight 3 of Chapter 4 and the corresponding Figure 4.11). From Figure 4.11 it becomes apparent that even if an interfering preamble is constructed from a different preamble code than the one used by the UOI, on average half of the pulses (eight out of sixteen) are still aligned; in the worst case even ten out of sixteen. For channel estimation, this means that with a single interfering preamble, more than half of the code symbols used in the averaging process may have a (potentially strong) contribution from the interferer. Consequently, the median will not be able to provide a robust estimate in this case.

To overcome this problem we propose to trade processing gain for robustness against interference. Due to the squaring operation in the energy-detection receiver, interference is mainly additive. Therefore, the more energy a code symbol contains, the more likely it is that interference is present. We therefore propose to only consider half of the code symbols in the averaging process, namely those that have the lowest received energy.

The received energy during the j-th code symbol of the i-th preamble symbol is given by

$$E_{i,j} = \sum_{m=0}^{M-1} y^{\text{pre}}_{m,j+iC} \tag{5.24}$$

Further, we treat code symbols for which $c_j = 0$ separately from code symbols for which $c_j \neq 0$, yielding two sets of energy values. We then sort each of these two sets. Prior to calculating the median, we discard all samples $y^{\text{pre}}_{m,j+iC}$ that belong to a code symbol whose corresponding energy value ends up in the upper half of one of the sorted sets.

Note that the receiver can only start to partition code symbols into a set of code symbols with low energy (to be kept) and a set of code symbols with high energy (to be discarded) after at least half of the $N_{\text{CH}}C$ code symbols used in channel estimation have been received. This is impractical as it requires a lot of samples to be stored before they can be processed and thus entails a large cost in terms of memory. To circumvent this problem we can resort to a simple

Algorithm 1: Heuristic for online calculation of the partition of set S into disjoint subsets S_1 and S_2, where S_1 contains the $\frac{|S|}{2}$ smallest elements of S and S_2 the $\frac{|S|}{2}$ largest elements of S

Input: Set $S = \{s_0, s_1, \ldots, s_{n-1}\}$ of n observations.
Output: Partition of S into $S1$ and $S2$, where $|S_1| \approx \frac{n}{2}$ and S_1 contains the majority of elements s_i of S for which $s_i \leq \text{median}\{S\}$.

```
begin
    S_1 ← s_0;
    S_2 ← ∅;
    V ← s_0;
    for i ← 1 to n − 1 do
        m ← median{V};
        if s_i ≤ m then
            S_1 ← S_1 ∪ s_i
        else
            S_2 ← S_2 ∪ s_i
        end
        V ← V ∪ s_i
    end
end
```

heuristic algorithm (shown in Algorithm 1) that calculates an approximate partition into high and low energy code symbols in an online fashion.

Similar considerations can be made for the median that does not lend itself as well to online updating as the mean. Still, there exist online algorithms, such as the remedian [169], for calculating estimates of the median. Such algorithms can be used in order to cut down on memory consumption and yield a computational complexity in the same order as the mean.

5.3 Performance Evaluation

The setup to evaluate the performance of the robust receiver is roughly the same we used to evaluate the impact of MUI on the simple energy-detection receiver in Chapter 4. We again assume the LPRF mode with the mandatory channel number 3. Further, we simulate a high and a low traffic case, where N_u users generate packets at rates $R = 20$ packets/s and $R = 100$ packets/s, according to the procedure in Section 2.5.2. The transmitted packets are assumed to be of the maximum allowable length of 1208 bits/packet.

We already mentioned that, in contrast to Chapter 4, we here assume perfect synchronization and SFD detection, and consequently, only channel estimation and payload decoding are

simulated. Since synchronization is not performed, we also concentrate on the scenario where the UOI and the interferers use different preamble codes. We again assign preamble code 5 to the UOI, and preamble code 6 to all interferers.

Remember that in our simulations, the signal to noise ratio (SNR) is defined as $\mathrm{SNR} = \frac{E_p}{N_0}$ where E_p is the received energy *per pulse* (after the convolution of the pulse with the impulse response of the channel) and where $N_0/2$ is the PSD of the AWGN noise process that is assumed to be bandlimited to $B = 500$ MHz. To model the wireless channel, we again use the IEEE 802.15.4a residential non line of sight (NLOS) channel model (CM2) [14] and the length of the channel estimation phase corresponds to $N_{\mathrm{CH}} = 16$ preamble symbols.

All confidence intervals shown are at the 95% level.

Throughout performance evaluation, we will denote an energy-detection receiver with a structure consisting of a single branch as conventional receiver or ED_{CONV}. An energy-detection receiver that adapts to the IEEE 802.15.4a burst transmissions by calculating the optimal weights according to (5.11), is denoted as ED_{OPT}. ED_{OPT} thus has a structure like it is depicted in Figures 5.1 or 5.2. If in addition ED_{OPT} employs a mechanism to mitigate interference, we call it robust energy-detection receiver or ED_{ROB}.

5.3.1 Enhanced Robustness Against MUI

Compared to the receiver considered in Chapter 4, the robust receiver, ED_{ROB}, proposed here shows a significantly increased robustness against MUI. This becomes evident in Figure 5.3, where we show the PER for the robust receiver in the same scenarios considered in Figure 4.6 of Chapter 4. Note, however, that the rate shown here ($R = 20$ packets/s) is twice as high as in the previous chapter.

Three different interference settings are shown for the robust receiver ED_{ROB}: a near-far setting with one interferer that has a power level exceeding the one of the UOI by 10 dB, an equal power setting with one interferer having the same received power as the UOI, and a setting with three weak interferers that are received with a power level of -3 dB with respect to the UOI. In each of these settings, we simulate ED_{ROB} using the thresholding mechanism described in Section 5.2, with $P_{\mathrm{MUI}}^{\mathrm{FA}} = 0.01$, and the robust parameter estimation of Section 5.2.1. Further, we for now restrict ourselves to the simplest receiver architecture in Figure 5.1, thus implying an integration time of $T = T_{\mathrm{c}} = 2$ ns (or equivalently a sampling frequency of 500 MHz.

For comparison we also show the performance of the optimal receiver ED_{OPT} (also with integration time $T = T_{\mathrm{c}} = 2$ ns) if no MUI is present. Finally, we show the corresponding

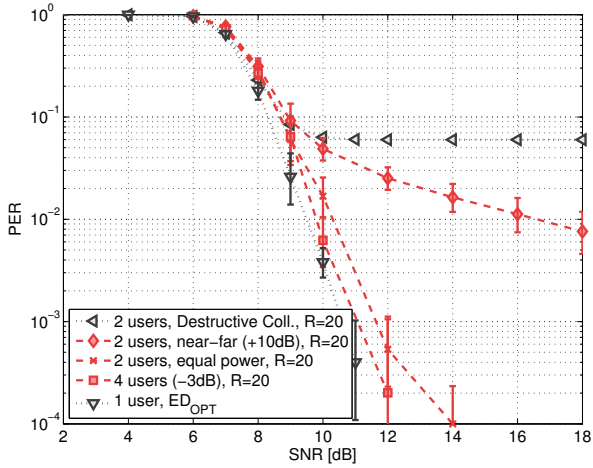

Figure 5.3: PER with MUI in a low traffic case (R = 20 packet/s) for the robust receiver ED_{ROB} with $T = T_{\text{c}} = 2$ ns. With equal or lower power interferers, performance is close to single user. Even in a harsh near-far scenario the receiver shows some capture effect when compared to the worst-case "Destructive Collisions" model.

curve of the "Destructive Collisions" model (see Section 4.3.2) that allows us to assess whether the receiver shows any capture effect and thus robustness to MUI.

For both, the weak and the equal power interference setting, performance of ED_{ROB} is close to the optimal single user performance. Even in the near-far setting, there is still a significant performance increase compared to the worst case "Destructive Collisions" curve. Performance of the receiver from Chapter 4 on the other hand was close to or even coincided with the "Destructive Collisions" for both the equal power and the near-far setting. We can thus already see that the mechanisms proposed in the current chapter, result in a significantly enhanced robustness of energy-detection receivers against MUI.

5.3.2 Limits of Conventional Energy-Detection Architectures and Impact of Robust Parameter Estimation

We have seen in Section 5.1.2 that a conventional energy-detection receiver structure consisting of a single branch cannot estimate the optimal weights. These weights are, however, required to optimally demodulate the bursts that are sent during an IEEE 802.15.4a payload. Further, they are essential for the MUI mitigation technique proposed here, because they form the basis for

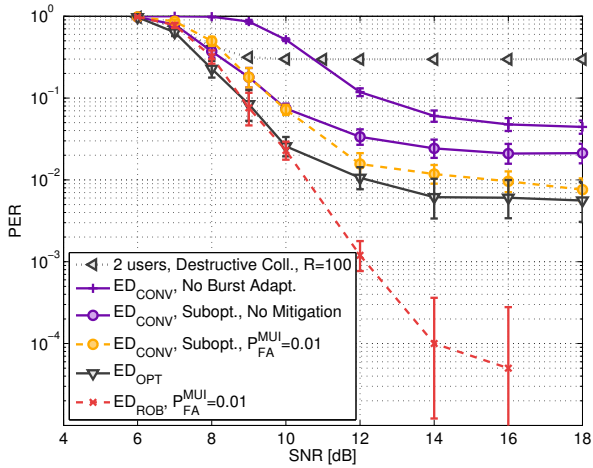

Figure 5.4: PER in the presence of a single interferer with power equal to the UOI. Packets are generated at rate R = 100 packet/s. Suboptimal receivers based on the conventional energy-detection architecture with a single branch perform worse than ED_{OPT} that calculates the optimal weights but does not mitigate MUI. Using thresholding and robust parameter estimation improves the PER by up to two orders of magnitude with respect to ED_{OPT}.

calculating the threshold used to reject strong interference terms. One option for a conventional receiver is to not adapt to burst transmissions at all, thus assuming that the signalling structure does not change between an IEEE 802.15.4a preamble and an IEEE 802.15.4a payload. Such a receiver would result in the use of a set of suboptimal weights $q_m = w_m^{(0)}$. Another option for a conventional receiver, ED_{CONV}, is to calculate the optimal weights according to (5.11), but to set parameters that it cannot calculate to zero. For a receiver with $T = T_c$ this would result in omitting the cross-terms that require additional branches, yielding $q_m = \sum_{j=0}^{N_{\text{cpb}}-1} w_{m-jK}^{(0)}$. A threshold to reject MUI can then be calculated according to (5.19) based on these suboptimal weights. Unfortunately, none of these strategies yields a performance that is close to ED_{ROB}. This can be seen in Figure 5.4 where we show the performance of different receivers in an equal power setting with one interferer and for a rate of $R = 100$ packets/s. All of the receivers compared have an integration time of $T = T_c = 2$ ns.

Figure 5.4 shows that a conventional receiver that calculates the suboptimal weights as explained above and that employs a threshold as well as robust parameter estimation (ED_{CONV}, Subopt., $P_{\text{MUI}}^{\text{FA}} = 0.01$), performs significantly better than the worst case given by the "Destructive Collisions" model [2]. However it only improves little over a conventional receiver that calculates suboptimal weights and does not use any form of interference mitigation (ED_{CONV}, Subopt., No Mitigation) or over a conventional receiver that does not adapt to the bursts of the payload (ED_{CONV}, No Burst Adapt.). Finally it even performs a little bit worse than ED_{OPT} that calculates the optimal weights but employs no interference mitigation.

Using ED_{ROB} on the other hand, can yield a PER up to two orders of magnitude lower, compared to ED_{OPT}. This is equivalent to an improvement of four orders of magnitude with respect to the worst case.

The PER in the corresponding near-far setting with a single strong interferer and a rate of $R = 100$ packets/s is shown in Figure 5.5. Due to the stronger interference all receivers that do not mitigate MUI perform close to or equal to the worst case. The PERs for ED_{CONV} being worse, we only show the curve for ED_{OPT}. In the near-far case, ED_{ROB} achieves an improvement of up to one order of magnitude with respect to ED_{OPT}. Figure 5.5 also illustrates that robust parameter estimation is needed. To this end we show two additional curves: the PER if ED_{ROB} only employs thresholding but parameter estimation is based on the sample mean and thus not robust (ED_{ROB}, Thld Only); and the PER if robust parameter estimation

2. Note that the "Destructive Collisions" model is based on the single user performance of the optimal receiver ED_{OPT}. Receivers with an inferior single user performance may thus perform worse than the "Destructive collisions" model at low SNR

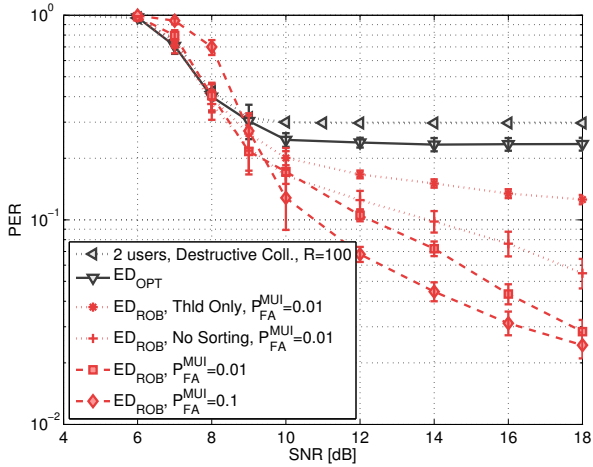

Figure 5.5: PER in the near-far setting with one interferer. Packets are generated at rate R = 100 packet/s. ED_{OPT} performs close to the worst case. Using the robust receiver can yield a gain of up to one order of magnitude. ED_{ROB} is not too sensitive to the choice of threshold. However, if parameter estimation is not performed in a robust fashion (ED_{ROB}, Thld. Only; ED_{ROB}, No Sorting), the full performance gain cannot be realized.

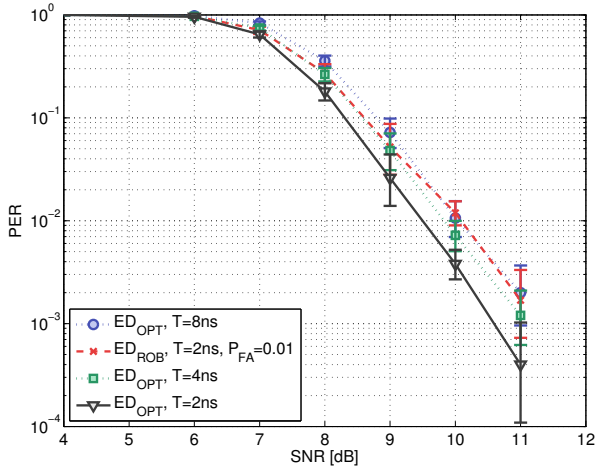

Figure 5.6: Performance without MUI. Decreasing the sampling rate leads to a moderate performance degradation because it leads to more noise being integrated. Performance is also degraded if thresholding is used, because the threshold also rejects some useful signal contributions.

uses the median but sorting of code symbols to reject those with high energy does not take place (ED_{ROB}, No Sorting). These curves clearly indicate that robust parameter estimation based on the median as well as rejection of high energy samples in the estimation process are required to get the optimal performance in terms of PER. However, using the online algorithm in Algorithm 1 or sorting code symbols only once all samples needed for parameter estimation are received, does not result in any noticeable performance difference and we therefore only show one curve for the full receiver. Finally, for the full receiver, we compare different values of the probability $P_{\text{MUI}}^{\text{FA}}$ that governs the threshold. In general the receiver is not too sensitive to the choice of threshold. A more aggressive threshold of $P_{\text{MUI}}^{\text{FA}} = 0.1$ yields a slightly improved PER at high SNRs where the signal is interference limited. At low SNRs the opposite is true, a more aggressive threshold leads to a slight performance decrease because it discards useful signal information.

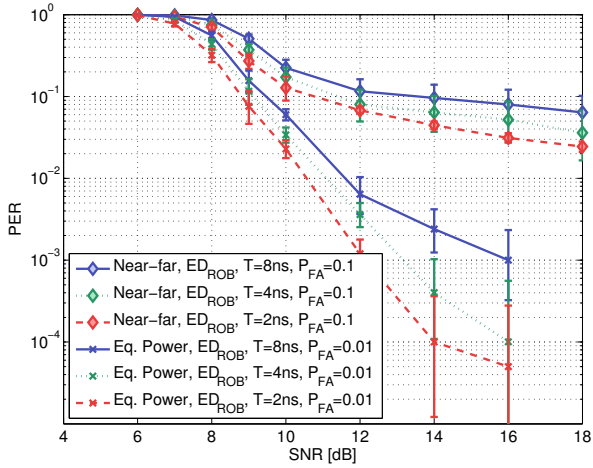

Figure 5.7: Reducing the sampling rate also decreases the robustness against MUI. While still moderate at $T = 4$ ns, the performance degradation starts to become significant at $T = 8$ ns.

5.3.3 Single User Performance and Impact of Increasing the Integration Time

We now evaluate the impact of the integration time on the performance of the receiver. It was shown in Section 5.1.2 that a longer integration time reduces the sampling rate, but necessitates additional branches and thus leads to a more complex hardware implementation of the receiver.

We compare receivers with different integration times of $T = T_c = 2$ ns, $T = 4$ ns and $T = 8$ ns. For $T = 2$ ns, calculating the optimal weights requires an architecture with 4 branches during channel estimation (shown in Figure 5.1). For $T = 4$ ns, 6 branches are needed (shown in Figure 5.2) and finally, $T = 8$ ns requires the maximum of $\frac{N_{\text{cpb}}(N_{\text{cpb}}+1)}{2} = 10$ branches.

Figure 5.6 shows the performance without MUI for ED_{OPT}. We see that increasing the integration time, decreases the performance. We attribute this to the fact that with longer integration times, more noise is integrated. However, the performance difference is not striking, it amounts to approximatively 0.5 dB between $T = 2$ ns and $T = 8$ ns.

A similar performance decrease can also be observed between ED_{OPT} and the robust receiver ED_{ROB} with $T = 2$ ns. Increasing the robustness against MUI thus comes at the cost of a slight degradation of the single user performance. This is due to the fact that both, robust parameter estimation using the median and demodulation using the thresholding non-linearity

are not optimal under AWGN only.

Figure 5.7 shows the performance of ED_{ROB} with different integration times in the near-far and equal power settings with one interferer and a rate of $R = 100$ packets/s. Also here, increasing the integration time results in a worse performance. Especially an integration time of $T = 8$ ns yields significantly worse performance than $T = 2$ ns. Due to the coarser estimate of the UOI channel energy-delay profile, mitigation of MUI is less effective for longer integration times.

5.4 Conclusion

We have shown that the performance of an energy-detection receiver can be dramatically improved in the presence of MUI through a simple, yet effective thresholding mechanism. Our thresholding scheme necessitates estimation of the channel energy-delay profile, which is already an integral part of many energy-detection receiver architectures (see Section 3.2 or, e.g., [33]). We can therefore consider the required additional complexity with respect to such receivers to be minimal: It consists in applying a non-linearity such as (5.20) during demodulation. This non-linearity depends on a threshold that can be tabulated. Further, we proposed a method based on order statistics to estimate the channel energy-delay profile in a manner that is robust to MUI. Our method only entails a slight increase in complexity with respect to classic energy-delay profile estimation that uses simple averaging with the mean.

Apart from energy-detection receivers that weight the received signal with the estimated energy-delay profile, it is also possible to design very simple energy-detection receivers, employing a long, fixed integration window. These simple receivers can perform most operations entirely in the analog domain, which is favorable in terms of complexity. However, they are very vulnerable to both noise and interference. Receivers that use a weighting function have a higher complexity, but as we have seen in this chapter, they allow for designs that are robust to MUI. With respect to complex Rake receivers, they still share many of the advantages of simple energy-detection receivers, such as lower sampling frequencies, less stringent timing requirements or no need for down conversion (see also Section 2.3 in Chapter 2), allowing for implementations of lower complexity.

Finally, we have shown in this chapter that in the case of signalling structures with scrambled bursts of pulses, like they are employed in IEEE 802.15.4a, estimation of the channel energy-delay profile requires a novel receiver architecture, resulting in an additional moderate increase in complexity. In the analog domain, this manifests itself in the need for additional

branches, mixing the received signal with a delayed version of itself; in the digital domain, the outputs of these branches need to be combined to yield the appropriate weighting function. In terms of complexity compared to sophisticated Rake receivers, also this receiver retains the advantages of an energy-detection receiver described above. Moreover, the additional complexity compared to a receiver that does not adapt to burst transmissions is only needed during channel estimation. Further, this new architecture allows for new trade-offs between interference robustness, hardware complexity and sampling rate requirements.

Chapter 6

Robust Non-Coherent Synchronization Algorithms for IEEE 802.15.4a IR-UWB Networks

IEEE 802.15.4a networks are packet-based and lack global synchronization. The first step towards the correct reception of a packet is therefore synchronization. It consists in detecting the presence of the packet on the channel and in finding the time reference of the source. Such a timing acquisition generally comprises a first coarse acquisition followed by a finer acquisition [170]. Only then can the destination begin to look for the start frame delimiter (SFD). The SFD is a specifically crafted data sequence that marks the end of the synchronization header and the beginning of the payload. Once the SFD is detected, the destination can finally recover the payload data by demodulating and decoding the received signal.

In the previous chapter we proposed a thresholding scheme for an IEEE 802.15.4a energy-detection receiver in order to mitigate impulsive interference during decoding of the payload data. Although the proposed receiver drastically improves the performance, our results from Chapter 4 indicate that such an improvement is not enough for an energy-detection receiver to reach reasonable performance levels in the presence of MUI. The reason is that not only payload decoding but also the synchronization phase are significantly affected by MUI. To have a system that is robust to MUI, we therefore also need to address synchronization, which is what we will do in this chapter.

Chapter 4 also indicates that interference is a serious issue for both timing acquisition and SFD detection (Insight 2 of Chapter 4). Here again, we thus have to address both mecha-

nisms for the system to become robust against MUI. In this chapter we propose PICNIC, a suite of robust and low-complexity algorithms tailored to IEEE 802.15.4a that allow for reliable synchronization with an IR-UWB energy-detection receiver in the presence of multi-user interference (MUI), even in near-far scenarios. For robust timing-acquisition, PICNIC employs an adaptive thresholding scheme, similar to what was proposed in [81] for coherent receivers. It further uses a simplified interference cancelation algorithm to alleviate cross-code interference between different IEEE 802.15.4a preamble codes. For robust SFD detection, we propose a novel algorithm that we call DESSERT. This algorithm transforms SFD detection into a decoding problem, allowing us to use thresholding techniques to mitigate interference.

This chapter is organized as follows: Section 6.1 reviews parts of the system model that are essential to understand this chapter. Packet detection and timing acquisition algorithms are presented in Section 6.2, while SFD detection is addressed in Section 6.3. The performance of our algorithms is assessed and compared to other solutions in Section 6.4. Section 6.5 concludes the chapter.

6.1 System Model and Assumptions

In this chapter, we again assume that the receiver front-end corresponds to the classical energy-detection receiver architecture introduced in Section 2.3.3 of Chapter 2. Like in the previous chapters, we focus on the IEEE 802.15.4a PHY described in Section 2.4.

With this receiver and PHY, the discrete time signal at the output of the energy-detection receiver during the reception of the preamble is

$$y_m^{\text{pre}} = \int_{mT}^{(m+1)T} [r_{\text{pre}}(t)]^2 \, \mathrm{d}t, \tag{6.1}$$

Like before, we assume that the integration time T is such that $M = L_{\text{s}} T_{\text{c}}/T$ discrete samples are produced per preamble code symbol.

As a reminder, the received signal during the preamble of the UOI, $r_{\text{pre}}(t)$, equals

$$r_{\text{pre}}(t) = \underbrace{\sum_{i=0}^{N_{\text{pre}}-1} s_i \sum_{j=0}^{C-1} c_j \cdot \tilde{h}(t - (j+iC)L_{\text{s}}T_{\text{c}} - \nu_0)}_{\tilde{x}_{\text{pre}}(t-\nu_0)} + v(t), \tag{6.2}$$

where $\tilde{x}_{\text{pre}}(t)$ is the contribution of the UOI and $v(t)$ accounts for noise and MUI. c_j is the

ternary preamble code sequence of length C. Further, every preamble symbol is modulated by s_i, which is given by

$$s_i = \begin{cases} 1 & \text{if } i \in \{0, 1, \ldots, N_{\text{sync}} - 1\} \\ s^{(\text{sfd})}_{i - N_{\text{sync}}} & \text{if } i \in \{N_{\text{sync}}, \ldots, N_{\text{pre}} - 1\} \end{cases}, \tag{6.3}$$

where $s^{(\text{sfd})}_i \in \{-1, 0, +1\}, i \in \{0, 1, \ldots, N_{\text{sfd}} - 1\}$ is the ternary SFD code of length N_{sfd} that marks the end of the preamble and the beginning of the payload. For additional details, please refer to Chapter 2.

6.2 Packet Detection and Timing Acquisition Algorithms

We compare three packet detection and timing acquisition algorithms with increasing degree of robustness to MUI. The "baseline" algorithm uses correlation with a known template. This algorithm corresponds to the algorithm that we introduced in the performance evaluation of the energy-detection receiver in Chapter 4. We have already seen that this approach is vulnerable to MUI. The "power-independent detection" (PID) enhances the baseline algorithm using thresholding (PID was developed in [81] for coherent reception). Finally, "power-independent detection and preamble code interference cancelation" (PICNIC) adds an interference cancelation (IC) scheme tailored to IEEE 802.15.4a.

All three algorithms proceed in two phases, coarse timing acquisition and fine timing acquisition. During coarse timing acquisition, they usually synchronize on the strongest multipath component, which is not always the first in time. Coarse timing acquisition is then followed by a fine timing acquisition to improve the timing accuracy. Note that in principle, our coarse timing acquisition algorithms can be combined with any fine timing acquisition algorithms and has not exclusively to be used in conjunction with the fine timing acquisition algorithms proposed here.

6.2.1 Baseline Algorithm

The baseline algorithm is a classic timing acquisition using a correlation of the receiver output with a template derived from the known preamble code sequence of the UOI. The full algorithm is explained in Section 4.2.1 of Chapter 4. The reader is referred there for a detailed description of the algorithm, here we merely give a short reminder, outlining its main principles.

The received signal given by (6.1) is correlated with the following template sequence

$$t_l = \sum_{k=0}^{N_{\mathrm{G}}-1} \sum_{j=0}^{C-1} c_j^2 \cdot \delta_{l-(j+kC)M} \tag{6.4}$$

where δ_m denotes the Kronecker delta, c_j is the preamble code sequence of length C and N_{G} is the number of preamble code symbols that are repeated in the template for processing gain.

The resulting discrete correlation output sequence z_m is $M \cdot C$-periodic if the UOI signal is present. Therefore, the algorithm processes the correlation output by blocks of MC consecutive samples, the i-th block being

$$\mathbf{z_i} = \{z_{iMC}, z_{iMC+1}, \ldots, z_{(i+1)MC-1}\} \tag{6.5}$$

Coarse Timing Acquisition

The baseline algorithm starts with coarse timing acquisition which involves two steps: detection and verification. During detection, the presence of a signal is declared if at least one of the correlation output samples of the current block exceeds the threshold η_{detect} given by (4.8). If a signal is detected, verification starts. During verification, the receiver verifies that the correlation output has the expected periodicity. It does so by checking that the maxima of two consecutive blocks are both well-aligned and above the threshold. If this holds for N_{V} consecutive blocks, coarse timing acquisition succeeds.

We saw in Chapter 4, that the baseline method works well in a single user scenario but that it does not take MUI into account. A strong interfering signal has a high likelihood of exceeding the threshold (that is solely based on the noise level), even if not perfectly aligned with the template. This can generate missed detections (MD) if the interfering signal introduces spurious maxima in the correlation output that make the verification fail. It can lead to a false alarm (FA) with synchronization on an interfering signal if $N_{\mathrm{V}} + 1$ interfering maxima are aligned.

Fine Timing Acquisition

Coarse timing acquisition is then followed by fine timing acquisition. During fine timing acquisition the jump-back-and-search-forward algorithm of Section 4.2.1 is applied in order to improve the TOA estimate. With this algorithm, a search-back window of $W = M/2$ samples preceding the sample found during coarse timing acquisition is searched. Fine timing acqui-

sition selects the first sample in time that is within the search-back window, lies above the threshold η_{detect} and exceeds the secondary correlation peaks that are due to the periodicity of the signal.

6.2.2 Power-Independent Detection Using Thresholding

The general concept of the PID was introduced for coherent receivers in [81]. We here show how it can be practically applied to IEEE 802.15.4a with energy-detection or more precisely to the baseline timing acquisition algorithm described in Sections 6.2.1 and 4.2.1.

Coarse Timing Acquisition

The PID prevents large interference terms in (6.1) from dominating the result of the correlation. This is achieved by applying a threshold check to the received signal at the *input* of the correlation (in contrast to the the baseline method where it is performed on the output). Samples above the threshold are set to 1, samples below are set to 0. The template corresponds to the one in the baseline algorithm and is given by (6.4). The correlation output becomes

$$z_m = \sum_{l=0}^{M_{\text{T}}-1} t_l \cdot \mathbf{1}_{\left[y^{pre}_{m-(M_{\text{T}}-1)+l} > \eta_{\text{pid}}\right]} \tag{6.6}$$

where $\mathbf{1}_{[\cdot]}$ denotes the indicator function. The threshold η_{pid} is given by

$$\eta_{\text{pid}} = \frac{N_0}{2} F^{-1}_{\chi^2}(1 - P^{\text{FA,PID}}_{\text{AWGN}} | 2BT) \tag{6.7}$$

and parameterized by $P^{\text{FA,PID}}_{\text{AWGN}}$, the desired probability of a pure noise signal exceeding the threshold η_{pid}.

Except for the threshold check against η_{detect}, which is omitted, the remaining steps of the baseline method are unchanged. In contrast to the original PID description in [81], where the involved thresholds were obtained through extensive simulations, our algorithm uses the explicit threshold computation of equation (6.7) [1].

1. The so-called elementary threshold from [81] corresponds to η_{pid} and the main threshold is omitted as we simply track the maximum over blocks of the correlation output.

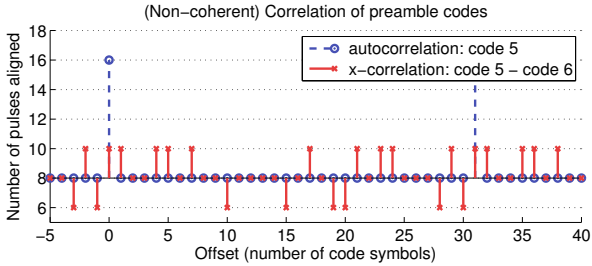

Figure 6.1: Auto-correlation and cross-correlation of the two IEEE 802.15.4a preamble codes 5 and 6 of length $C = 31$ when non-coherent reception is used. The cross-correlation shows 10 peaks per period that may cause false alarms.

Fine Timing Acquisition

Fine synchronization is identical to the one for the baseline method (Sections 6.2.1 and 4.2.1) with the only difference that a different threshold check needs to be applied. The output of the correlation (6.6) is now distributed according to a binomial distribution with parameters $C_{\text{NZ}} \cdot N_{\text{G}}$ and $P_{\text{AWGN}}^{\text{FA,PID}}$, yielding the following threshold that replaces η_{detect} during fine synchronization

$$\eta_{\text{pid}}^{\text{fine}} = \frac{N_0}{2} F_{\text{BIN}}^{-1}(1 - P_{\text{AWGN}}^{\text{FA,PID,fine}} | C_{\text{NZ}} N_{\text{G}}, P_{\text{AWGN}}^{\text{FA,PID}}). \tag{6.8}$$

$P_{\text{AWGN}}^{\text{FA,PID,fine}}$ is the probability that a pure noise signal exceeds the fine timing acquisition threshold $\eta_{\text{pid}}^{\text{fine}}$, C_{NZ} are the number of nonzero code symbols in the preamble code and $F_{\text{BIN}}^{-1}(x|n,p)$ is the inverse of the CDF of a binomial distribution with parameters n and p.

6.2.3 Preamble Code Interference Cancelation

The two IEEE 802.15.4a preamble codes used per frequency band do not have a perfect cross-correlation as we have shown in Chapter 4 (see Insight 3 of Chapter 4). This is shown once more in Figure 6.1 for the code sequences 5 and 6, of length $C = 31$. While code 5 has a periodic auto-correlation with only one peak per period, its cross-correlation with code 6 shows 10 peaks per period. These cross-correlation peaks can generate FAs or MDs for a receiver using code 5. The resulting performance loss is significant, as we will see in Section 6.4.

Note that such cross-correlation peaks are not an anomaly of the IEEE 802.15.4a preamble codes per se, but rather inherent to any pseudo-random binary sequences with perfect autocorrelation. Although the design of optimal code sequences is a vast research area on its own and

therefore out of the scope of this thesis (and the codes of the standard have anyway already been fixed), we note that the cross-correlation of, e.g., codes 5 and 6, is already optimal in the sense that it is only three-valued and that the correlation peaks are of minimal value. For an in-depth introduction to pseudo-random sequences and their properties, we refer the reader to [171].

The PICNIC algorithm attempts to detect and cancel out interference that is due to these cross-correlation peaks by looking for the pattern of the cross-correlation and subtracting it from the correlation output if it is found present. PICNIC essentially pre-processes each block $\mathbf{z_i}$, obtained from the PID correlation output in (6.6), before handing it over to the coarse synchronization. For convenience, we omit the index i from here on. If interference is present, $\mathbf{z}$ contains C_{peak} sub-blocks of length M with high energy, corresponding to the C_{peak} peaks in the cross-correlation. It also contains C_{trough} sub-blocks with low energy, corresponding to the C_{trough} troughs in the cross-correlation. The remaining $C_{\text{mid}} = C - C_{\text{peak}} - C_{\text{trough}}$ blocks have a medium energy level. The algorithm proceeds in three steps: 1) compare the positions of the high-, medium- and low-energy sub-blocks with the cross-correlation and decide whether interference is present 2) if present, find the exact beginning of the sub-blocks such that 3) sub-blocks with similar energy levels can be averaged yielding an estimate of the channel-energy delay profile that can be subtracted. These steps are detailed in the following and illustrated in Figure 6.2. In what follows, the mandatory frequency band 3 with codes 5 and 6 serves as an example. However, the method equally applies to the other frequency bands with other codes. Also, as IEEE 802.15.4a allows two codes per frequency band, the knowledge of a single cross-correlation per band is sufficient.

Detecting the Presence of an Interfering Preamble Code

PICNIC tries to identify sub-blocks with energy levels corresponding to the cross-correlation pattern. Two ternary vectors of length C, $\mathbf{x}^{\text{tern}}_{\text{cross}}$ and $\mathbf{z}^{\text{tern}}$ are correlated, representing the different energy levels of the cross-correlation and of the C sub-blocks of $\mathbf{z}$, respectively.

$\mathbf{z}^{\text{tern}}$ is constructed from $\mathbf{z}$ (see Figure 6.2-1 for an illustration) by first determining the maximum over every sub-block of length M yielding vector $\mathbf{z}^{\text{max}}$ with elements

$$z_j^{\text{max}} = \max(z_{jM}, z_{jM+1}, \ldots, z_{(j+1)M-1}), \quad j \in \{0, \ldots, C-1\}$$

Second, $\mathbf{z}^{\text{max}}$ is converted to the ternary vector $\mathbf{z}^{\text{tern}}$ by replacing its C_{peak} highest values with "+1", its C_{trough} lowest values with "-1" and the rest with "0". The cross-correlation (shown in

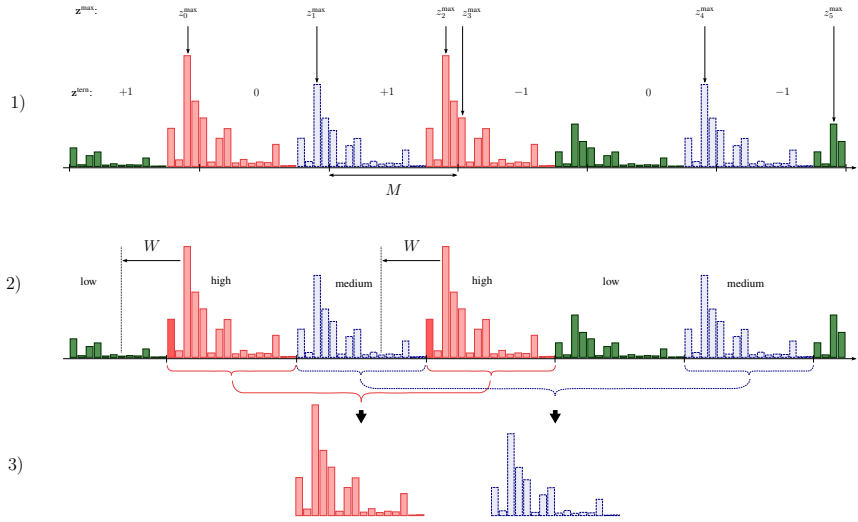

Figure 6.2: Three steps of the PICNIC algorithm to cancel the effect of interfering code: 1) Interference is detected by matching cross-correlation pattern to high-, mid- and low-energy blocks in the correlation output. (2) Time-base is aligned on the interferer to find beginning of blocks via a search-back algorithm. 3) Channel energy-delay profile to be subtracted is calculated separately for high-, mid- and low-energy blocks via robust method based on order statistics.

Figure 6.1 for codes 5 and 6) is mapped to a ternary vector $\mathbf{x}_{\text{cross}}^{\text{tern}}$ following the same procedure.

To detect interference, $\mathbf{z}^{\text{tern}}$ is correlated with $\mathbf{x}_{\text{cross}}^{\text{tern}}$ and the maximum of the correlation is compared with the interference detection threshold

$$\eta_{\text{picnic}} = \lfloor \frac{C_{\text{peak}} + C_{\text{trough}}}{2} \rfloor + 1$$

i.e., we test that more than half of the peaks and troughs of the cross-correlation correspond to the peaks and troughs of the sample vector. If the maximum is above η_{picnic}, we assume that interference is present and continue the algorithm. Otherwise, we continue the timing acquisition according to the PID method.

Note that all of the IEEE 802.15.4a preamble codes have a structure similar to codes 5 and 6 whose cross-correlation is shown in Figure 6.1. The same holds for the respective cross-correlations. For every preamble code pair that is assigned to the same frequency band, we can thus identify C_{peak} peaks and C_{trough} troughs. The method described here thus not only applies to the mandatory codes 5 and 6 but can directly be applied to other IEEE 802.15.4a codes as well.

Determination of the First Multipath Component

If interference is present, it needs to be subtracted from the vector $\mathbf{z}$. Hence, a rough estimate of the channel energy-delay profile of the interfering signal must be obtained. This implies that the first multipath component of the interfering signal must be found. We use a jump-back-and-search-forward procedure similar to the one used in the fine timing acquisition algorithm of the PID (Section 6.2.2).

First, using a majority vote on the indices of the samples $z_j^{\max}$ of the C_{peak} high energy sub-blocks corresponding to a "+1" in $\mathbf{z}^{\text{tern}}$, the index of the strongest path into a sub-block of size M is determined. Second, for each of the C_{peak} high-energy sub-blocks, we start from the strongest path and search in a window of length $W = M/2$ the first path above the noise threshold given by

$$\eta_{\text{picnic}}^{\text{jump}} = \frac{N_0}{2} F_{\text{BIN}}^{-1}(1 - P_{\text{AWGN}}^{\text{FA,PICNIC}} | C_{\text{NZ}} N_{\text{G}}, P_{\text{AWGN}}^{\text{FA,PID}}). \tag{6.9}$$

The binomial distribution with parameters $C_{\text{NZ}} \cdot N_{\text{G}}$ and $P_{\text{AWGN}}^{\text{FA,PID}}$ corresponds to the distribution of the correlation output (6.6) if only AWGN is present. The threshold is set by fixing $P_{\text{AWGN}}^{\text{FA,PICNIC}}$, the probability that AWGN can exceed the threshold. Finally, the first path index is

the lowest one found by more than half of the C_{peak} individual search procedures.

Interference Cancelation by Subtraction of the Estimated Channel Energy Delay Profile

When aligned with the interfering signal, $\mathbf{z}$ is split up into C_{peak} high-energy sub-blocks, C_{trough} low-energy sub-blocks and C_{mid} medium-energy sub-blocks. The signal is wrapped around if needed (see Figure 6.2 for the first low-energy sub-block). Then, the energy-delay profile $\mathbf{q}^{\text{high}} = \{q_0^{\text{high}}, q_1^{\text{high}}, \ldots, q_{M-1}^{\text{high}}\}$ is estimated for the high-energy sub-blocks, similar to the robust channel estimation algorithm using the median that was presented in Section 5.2.1 of Chapter 5.

Let $z_{j,m}^{\text{high}}$ denote the m-th sample of the j-th high-energy sub-block. We find $\mathbf{q}^{\text{high}}$ according to

$$q_m^{\text{high}} = \text{median}\left\{ z_{j,m}^{\text{high}} : j \in \{0, 1, \ldots, C_{\text{peak}} - 1\} \right\} - \bar{v}_{\text{AWGN}}^{\text{PICNIC}} \tag{6.10}$$

where the median is used instead of the mean to be robust to outliers (which might include e.g. the signal of the UOI that we do not want to subtract) and $\bar{v}_{\text{AWGN}}^{\text{PICNIC}} = C_{\text{NZ}} \cdot N_{\text{G}} \cdot P_{\text{AWGN}}^{\text{FA,PID}}$ is the expected noise level at the output of the correlation (6.6). To cancel interference, we can now subtract $\mathbf{q}^{\text{high}}$ from all the high-energy sub-blocks in $\mathbf{z}$. We then proceed similarly for the medium- and low-energy sub-blocks.

6.3 Start Frame Delimiter Detection Algorithms

After timing acquisition, the receiver knows the beginning of a preamble symbol. It may then perform channel estimation, e.g., according to the robust procedure that we introduced in Section 5.2.1 of Chapter 5. Still, the receiver does not know exactly how many preamble symbols where used up during timing acquisition [2]. Assuming that packet detection and timing acquisition were performed perfectly, the TOA is now of the form $\nu_0 = N_\Delta \cdot C L_{\text{s}} T_{\text{c}}$, i.e., the only uncertainty on the TOA is the number of preamble symbols N_Δ used during timing acquisition. To simplify notation we will in what follows make the equivalent assumption that $\nu_0 = 0$ but that the number of preamble symbols N_{sync} sent in the SYNC part of the preamble is unknown.

To provide the receiver with a means to detect the end of the preamble and the beginning of the payload, the end of the preamble is marked with the special SFD sequence (see Section 2.4.2 or the system model in this chapter). After channel estimation, the receiver starts to look for

2. Due to changing channel conditions or interference, the number of preamble symbols required to acquire timing may differ significantly.

the SFD. Once the SFD has been detected, decoding of the data bits of the payload can start.

In the following, we will introduce different SFD detection algorithms, that we will later compare in both a single user and in a multi user setting.

6.3.1 Algorithm Based on Correlation with Bipolar Template

In Chapter 4 we introduced an SFD detection method that is based on correlation of the received signal y_{mm}^{pre} with a bipolar template that is derived from the SFD code $s_i^{(\text{sfd})}$. Although this algorithm yields an acceptable performance if no MUI is present (Section 4.3.1) it is not robust to MUI due to multiple reasons (Insight 2, Insight 4 and Insight 5 of Section 4.3).

In the following, we will see that even in a single user setting without MUI, there are several algorithms that significantly outperform the correlation based algorithm. We will therefore not revisit this algorithm here, but rather refer the interested reader to Section 4.2.3 of Chapter 4 for a detailed description of the correlation based SFD detection method.

6.3.2 Algorithms Based on Sequential Decoding of Preamble Symbols

An alternative to detecting the SFD through a correlation procedure is to look at SFD detection as a *decoding problem*. With this approach, the receiver decodes N_{sfd} consecutive received preamble symbols and tries to determine whether they correspond to the squared SFD sequence $\mathbf{s}^{\mathbf{2}(\text{sfd})} = (s_0^{2(\text{sfd})}, \ldots, s_{N_{\text{sfd}}-1}^{2(\text{sfd})})$. The square operation occurs because we have a non-coherent receiver. Further, we know from Chapter 5 that thresholding can effectively mitigate interference during data decoding and we can therefore hope to reuse similar concepts for robust SFD detection.

Online Algorithm

We propose an online algorithm that sequentially processes preamble symbols as they are received. For each received preamble symbol, the algorithm calculates a series of likelihood-ratio tests. Based on the results of these tests, it decides whether to stop (if it estimates that it detected the SFD) or whether to continue receiving preamble symbols (if it estimates that the SFD is still to come). Because of the way our algorithm works we call it DESSERT for "detection of SFD through sequential ratio tests".

The DESSERT algorithm proceeds by sequentially considering blocks of N_{sfd} preamble

symbols. The k-th block can be represented as a vector of consecutive samples

$$\mathbf{y_k} = (y^{\text{pre}}_{kCM}, y^{\text{pre}}_{kCM+1} \ldots, y^{\text{pre}}_{(k+N_{\text{sfd}})CM-1}) \tag{6.11}$$

Every sample of the energy-detection receiver output is distributed according to a scaled non-central chi-square distribution with $2BT$ degrees of freedom and non-centrality parameter $\frac{p_m}{N_0/2}$ (see Section 2.3.3 in the system model chapter) with p_m given by

$$p_m = \int_{mT}^{(m+1)T} [\tilde{x}_{\text{pre}}(t)]^2 \, \mathrm{d}t. \tag{6.12}$$

Plugging (6.2) into (6.12) and assuming no inter-pulse interference (i.e., $\tilde{h}(t) = 0, \forall t < 0, \forall t > L_{\text{s}}T_{\text{c}}$), we obtain

$$p_m = s^2_{\lfloor \frac{m}{CM} \rfloor} \cdot c^2_{\lfloor \frac{m}{M} \rfloor \mod C} \cdot q_{m \mod M} \tag{6.13}$$

with

$$q_m = \int_{mT}^{(m+1)T} [\tilde{h}(t)]^2 \, \mathrm{d}t, \quad m \in \{0, 1 \ldots, M-1\}. \tag{6.14}$$

The coefficient q_m represents the energy-delay profile of the channel and can be estimated in a robust fashion following the procedure proposed in Chapter 5 [3].

Assuming independence between samples y^{pre}_m, it follows that for the k-th block of N_{sfd} consecutive preamble symbols

$$f(\mathbf{y_k}|\mathbf{\Theta}, \mathbf{s^2_k}) = \prod_{i,j,m} \frac{1}{N_0/2} f_{\text{NC}\chi^2}\Big(\frac{y_{(k+i)CM+jM+m}}{N_0/2}\Big| 2BT, \frac{s^2_{k+i}c^2_j q_m}{N_0/2}\Big), \tag{6.15}$$

where $i \in \{0, \ldots, N_{\text{sfd}}-1\}$, $j \in \{0, \ldots, C-1\}$ and $m \in \{0, \ldots, M-1\}$. $\mathbf{\Theta} = (N_0/2, BT, q_m)$ contains only quantities that are either known or can be estimated robustly. Finally, $\mathbf{s^2_k}$ is defined as

$$\mathbf{s^2_k} = (s^2_k, \ldots, s^2_{k+N_{\text{sfd}}-1}) \tag{6.16}$$

Due to the structure of the preamble given by (6.3), only $N_{\text{sfd}} + 1$ sequences are possibly

3. During the preamble we are in the case $N_{\text{cpb}} = 1$ and q_m thus corresponds to the coefficient $w^{(0)}_m$. This coefficient can be estimated in an interference robust way from the energy-detection receiver output according to equation (5.22)

observable for $\mathbf{s_k^2}$, namely

$$\mathbf{s_k^2} = \begin{cases} \mathbf{s^{2(\text{sfd})}} & \text{if } k = N_{\text{sync}} \\ (1, s_0^{2(\text{sfd})}, \ldots, s_{N_{\text{sfd}}-2}^{2(\text{sfd})}) & \text{if } k = N_{\text{sync}} - 1 \\ (1, 1, s_0^{2(\text{sfd})}, \ldots, s_{N_{\text{sfd}}-3}^{2(\text{sfd})}) & \text{if } k = N_{\text{sync}} - 2 \\ \ldots & \ldots \\ (1, 1, \ldots, 1, s_0^{2(\text{sfd})}) & \text{if } k = N_{\text{sync}} - N_{\text{sfd}} + 1 \\ (1, 1, \ldots, 1) & \text{otherwise.} \end{cases} \tag{6.17}$$

Further, N_{sfd} is usually small ($N_{\text{sfd}} = 8$ for the mandatory mode of the IEEE 802.15.4a standard). For every block $\mathbf{y_k}$ the DESSERT algorithm calculates the likelihood of each of the $N_{\text{sfd}} + 1$ possible sequences according to (6.15) and declares presence of the SFD if the sequence with maximum likelihood is $\mathbf{s^{2(\text{sfd})}}$. In other words, the SFD is detected if

$$\mathbf{s^{2(\text{sfd})}} = \arg\max_{\mathbf{s_k^2}} f(\mathbf{y_k}|\mathbf{\Theta}, \mathbf{s_k^2}) \tag{6.18}$$

which is equivalent to

$$\ln(f(\mathbf{y_k}|\mathbf{\Theta}, \mathbf{s^{2(\text{sfd})}})) \geq \ln(f(\mathbf{y_k}|\mathbf{\Theta}, \mathbf{s_k^2})), \quad \forall \mathbf{s_k^2} \tag{6.19}$$

Combining (6.15) and (6.19), this can be expressed as the following log-likelihood ratio

$$\begin{aligned} \text{LLR}_{\text{sfd}}(\mathbf{y_k}|\mathbf{\Theta}, \mathbf{s_k^2}) &= \ln(f(\mathbf{y_k}|\mathbf{\Theta}, \mathbf{s^{2(\text{sfd})}})) - \ln(f(\mathbf{y_k}|\mathbf{\Theta}, \mathbf{s_k^2})) \\ &= \sum_{j,m} \sum_{\substack{i \\ s_i^{2(\text{sfd})} \neq s_{k+i}^2}} (2s_i^{2(\text{sfd})} - 1) \cdot \text{LLR}(y_{(k+i)CM+jM+m}|N_0/2, 2BT, c_j^2 q_m) \\ &\geq 0, \quad \forall \mathbf{s_k^2} \end{aligned} \tag{6.20}$$

where

$$\begin{aligned} \text{LLR}(y_m|N_0/2, 2BT, q_m) &= \ln\left[\frac{f_{NC\chi^2}\left(\frac{y_m}{N_0/2}\middle|2BT, \frac{q_m}{N_0/2}\right)}{f_{\chi^2}\left(\frac{y_m}{N_0/2}\middle|2BT\right)}\right] \\ &= \ln\left[{}_0F_1\left(;BT;\frac{q_m y_m}{N_0^2}\right)\right] - \frac{q_m}{N_0} \end{aligned} \tag{6.21}$$

corresponds to the log-likelihood ratio leading to the optimal decision rule for payload demod-

ulation, which we already encountered in Equation (5.6) of Chapter 5. For a fast and efficient evaluation, a tabulated version of the logarithm of the confluent hypergeometric limit function $\ln\left[{}_0F_1\left(;BT;x\right)\right]$ would have to be stored in a practical receiver.

If the SFD is found present, the algorithm stops and payload decoding can begin. Otherwise the algorithm continues by receiving the next preamble symbol. If the SFD is not found during a maximum of N_{sync} preamble symbols, the algorithm can safely assume that it must have missed the SFD. In this case reception of the packet is abandoned and the receiver goes back to packet detection and timing acquisition mode.

Note that due to the independence assumption, the contribution of every sample to the likelihood (6.15) can be computed individually, sample by sample.

The algorithm explained so far uses soft-decision decoding. It compares the soft log-likelihood ratios of the possible observable sequences $\mathbf{s_k^2}$ and tests whether the sequence corresponding to the SFD is the one with the highest log-likelihood. Evidently it is also possible to adapt the algorithm such that is uses hard-decision decoding when testing for the most likely sequence. In the case of hard-decision decoding, every preamble symbol is decoded individually, yielding the sequence of blocks of decoded preamble symbols $\hat{s}_k^2$

$$\hat{\mathbf{s}}_\mathbf{k}^\mathbf{2} = (\hat{s}_k^2, \ldots, \hat{s}_{k+N_{\text{sfd}}-1}^2) \tag{6.22}$$

The preamble symbols are decoded according to the optimal OOK decision rule

$$\sum_{j,m} \text{LLR}(y_{kCM+jM+m}|N_0/2, 2BT, c_j^2 q_m) \underset{\hat{s}_k^2=0}{\overset{\hat{s}_k^2=1}{\gtrless}} 0 \tag{6.23}$$

With hard-decision decoding, the DESSERT algorithm calculates the Hamming distance between the decoded sequence $\hat{\mathbf{s}}_\mathbf{k}^\mathbf{2}$ and each of the possible sequences given by (6.17). If the sequence closest to $\hat{\mathbf{s}}_\mathbf{k}^\mathbf{2}$ is $\mathbf{s}^{\mathbf{2}(\text{sfd})}$ and no other sequence is equally close, detection of the SFD is declared and the algorithm stops.

Offline Algorithm

An offline algorithm to detect the SFD was proposed in [111]. The algorithm is similar in nature to our DESSERT algorithm with hard-decision decoding. The algorithm from [111] also decodes preamble symbols according to the decision rule in (6.23). However, it does not decide on the presence of the SFD in an online fashion. Rather, it stores decoded preamble symbols $\hat{s}_k^2$ in memory. It does so until it has received N_{mem} preamble symbols. N_{mem} is

chosen such that it is ensured that the N_{mem} preamble symbols contain the SFD. The algorithm then goes through the stored decoded preamble symbols and calculates the Hamming distance between every block $\hat{\mathbf{s}}_{\mathbf{k}}^{\mathbf{2}}$, $k \in \{0, \ldots, N_{\text{mem}} - N_{\text{sfd}}\}$, and the SFD sequence $\mathbf{s}^{\mathbf{2}(\text{sfd})}$. Detection of the SFD is declared at the position k where the Hamming distance is minimal. [111] does not specify what happens if we encounter several minima. In our implementation we in this case declare the SFD at the earliest (in time) position with minimal Hamming distance.

The drawback of such an offline algorithm is that the SFD is not detected instantly but only after an important part of the payload has already been received. This is because N_{mem} has to be chosen large enough such that it for sure contains the SFD. This in turn makes the search window within which the algorithm looks for the SFD extend into the payload part of the packet. Consequently, this leaves the receiver with two options, both of which are suboptimal:

1. Received samples are stored in memory until the SFD is detected which is impractical if not infeasible. The portion of stored samples that eventually turn out to belong to the payload can only be processed for payload decoding once the position of the SFD is known

2. Payload decoding is performed in parallel to SFD detection, putting additional stress on the receiver. Decoded symbols that turn out to belong to the preamble once the SFD is detected, are discarded. The receiver thus also performs part of the work in vain.

For the offline algorithm it is also possible to define a detection criterion that is based on a soft metric instead of the Hamming distance. We propose to use the following metric, borrowing from the soft metric (6.20) used in the DESSERT algorithm

$$\text{LLR}_{\text{sfd}}^{\text{off}}(\mathbf{y_k}|\mathbf{\Theta}) = \sum_{i,j,m} (2s_i^{2(\text{sfd})} - 1) \cdot \text{LLR}(y_{(k+i)CM+jM+m}|N_0/2, 2BT, c_j^2 q_m) \tag{6.24}$$

With the soft metric, the offline algorithm decides on the index k for which (6.24) is maximal.

An alternative would be to directly use the absolute log-likelihood function of the SFD sequence, $\ln(f(\mathbf{y_k}|\mathbf{\Theta}, \mathbf{s}^{\mathbf{2}(\text{sfd})}))$, as a metric for the soft version of the offline algorithm. However, this approach does not work well as we will see in Section 6.4.

6.3.3 Improving Robustness to MUI

None of the algorithms presented so far is a priori robust to MUI. However, for algorithms that are based on decoding we may be able to reuse ideas that proved to successfully mitigate

interference during payload decoding. Since SFD detection according to both the DESSERT and the offline algorithm requires the estimation of the same parameters as payload decoding, it seems only natural to adopt the robust parameter estimation using order statistics that we described in Section 5.2.1 of Chapter 5.

In addition, the adaptive thresholding scheme to reject strong interference terms proposed in Section 5.2 can also be employed. The receiver can calculate the adaptive threshold

$$\eta_m^{\text{sfd}} = \frac{\hat{N}_0}{2} F^{-1}_{\text{NC}\chi^2}(1 - P_{\text{MUI}}^{\text{FA,SFD}} | 2BT, \frac{\hat{q}_m}{\hat{N}_0/2}) \tag{6.25}$$

that governs the non-linearity

$$g(y_m^{\text{pre}} | \hat{q}_m, \hat{N}_0/2) = \begin{cases} y_m^{\text{pre}} & \forall m : y_m^{\text{pre}} \leq \eta_m^{\text{sfd}} \\ \hat{q}_m & \forall m : y_m^{\text{pre}} > \eta_m^{\text{sfd}} \end{cases} \tag{6.26}$$

In the DESSERT and in the offline algorithm, $g(y_m^{\text{pre}} | \hat{q}_m, \hat{N}_0/2)$ can then be used in place of y_m^{pre}. Other non-linear operations are of course also possible. We also considered erasing samples that exceed the threshold such that they do not contribute to the decision. However, since we could not find any noticeable performance difference between different schemes, we adopted the same scheme that we already used for decoding and that is given in (6.26).

6.4 Performance Evaluation

Like in the chapters before, we use packet-based IEEE 802.15.4a system simulations to evaluate the performance. We simulate one receiver and N_{u} transmitters. Each of the transmitters generates packets according to a Poisson process with rate $R = 100$ packets/s. At the maximum allowed packet size with a payload of 1016 information bits, this corresponds to roughly half the peak rate (see Section 2.5.2). Two types of interfering users are simulated: Near-far interferers with a power level 10 dB higher than the UOI and equal power interferers with a power level equal to the UOI.

We consider the mandatory LPRF mode of IEEE 802.15.4a with channel number 3. The corresponding preamble codes 5 and 6 have a length of $C = 31$ and the cross-correlation pattern shown in Figure 6.1. The preamble has the default length of $N_{\text{sync}} = 64$ code repetitions during the SYNC part, followed by the $N_{\text{sfd}} = 8$ symbols forming the SFD. Our receiver has an integration time of $T = T_{\text{c}}$, but all algorithms would also allow for longer integration times.

The propagation channel is modeled according to the IEEE 802.15.4a residential NLOS (CM2) and the office LOS (CM3) channel models [14]. However, the results of NLOS and LOS being very similar, we mostly only show the results for NLOS.

The signal-to-noise ratio (SNR) is defined as $\mathrm{SNR} = \frac{E_p}{N_0}$ where E_p is the received energy per pulse (after the convolution of the pulse with the impulse response of the channel). Confidence intervals shown are at the 95% level.

6.4.1 Performance of Packet Detection and Timing Acquisition

For packet detection and timing acquisition, the receiver uses a template with $N_{\mathrm{G}} = 10$ repetitions of the code and performs $N_{\mathrm{V}} = 16$ verification steps. This set of parameters, found through simulations, keeps timing acquisition fast enough such that other tasks, e.g., channel estimation, can still be performed on the preamble. Channel estimation is performed according to the methods explained in Chapter 5 and lasts for $N_{\mathrm{CH}} = 16$ preamble symbols.

The probabilities that set the sensitivity of the various thresholds during timing acquisition correspond to $P_{\mathrm{AWGN}}^{\mathrm{FA}} = 1e-3$, $P_{\mathrm{AWGN}}^{\mathrm{FA,PID}} = 0.2$, $P_{\mathrm{AWGN}}^{\mathrm{FA,PID,fine}} = 1e-4$. Extensive simulations showed that a wide range of values gives similar performance, as long as $P_{\mathrm{AWGN}}^{\mathrm{FA,PID}}$ is not set too low.

Our main performance metric is the packet acquisition error rate (ACQER) which includes FAs and MDs. A packet is correctly acquired if the receiver synchronizes on a multipath component of the correct preamble code symbol.

Coarse Timing Acquisition

Figure 6.3 shows the performance of the different algorithms in a near-far scenario with two near interferers ($N_{\mathrm{u}} = 3$). The two interferers use preamble code 6, the UOI uses code 5. The baseline method is not robust: More than 10% of the packets are lost due to interference. On the other hand, the PID method is able to reduce the impact of large interference terms generated by interfering preambles or payloads of the near interferers. The ACQER is improved by about one order of magnitude at high SNR. Still, even with the PID and the use of different preamble codes, FAs occur due to the imperfect cross-correlation. The PICNIC algorithm reduces this type of interference, we gain another order of magnitude. For reference, we show the single user performance of both the baseline and the PICNIC method. Single user performance of the PID method is identical to PICNIC. In this case, the PICNIC algorithm performs slightly worse than the baseline method, which is due to the fact that the threshold η_{pid} here also removes some

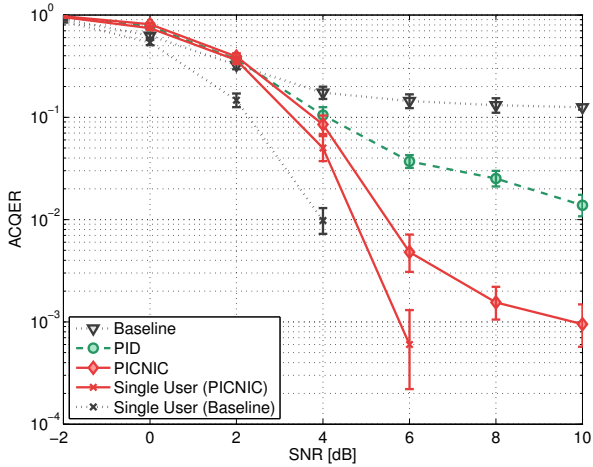

Figure 6.3: ACQER for the different algorithms with two interferers in a near-far configuration. Preamble codes of interferes differ from the one of the UOI. The baseline method is not robust. PID is able to reduce strong interference. Interference due to imperfect cross-correlation is only reduced by the PICNIC method, yielding a gain of up to two orders of magnitude.

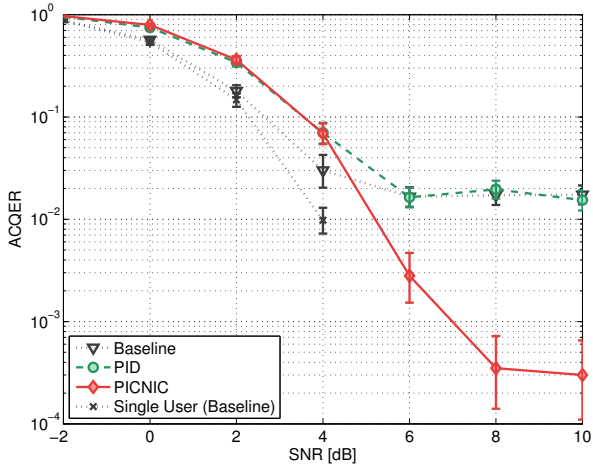

Figure 6.4: ACQER for the different algorithms with two interferers in an equal power configuration. Preamble codes of interferes differ from the one of the UOI. PID and baseline method perform the same since all of the errors are caused by interference due to imperfect cross-correlation. This type of interference is only reduced by the PICNIC method, yielding again a significant improvement.

useful signal information.

Figure 6.4 shows performance for two interferers with different codes, but with power levels equal to the UOI. In the interference limited SNR regions, the PID and the baseline methods have equal performance and show an error floor due to the imperfect cross-correlation. In this case interference is not high enough to dominate the output of the correlation. Consequently, the PID does not improve the performance since all of the acquisition errors are due to the imperfect cross-correlation. The PICNIC method again significantly reduces this type of interference.

If all transmitters use the same preamble codes, lots of FAs occur because the receiver cannot distinguishing an interfering signal from the signal of the UOI (see also Section 4.3.5). Independently of the algorithms used, the ACQER is consequently very high (around 15%). A more meaningful metric is needed that allows for the quantification of the capture effect capabilities of the receiver. We define the *capture error rate* (CER) as the probability that a UOI packet is lost and the receiver does not correctly acquire an interfering packet either. Our simulation results show that if two packets arrive at the receiver at about the same time, the one

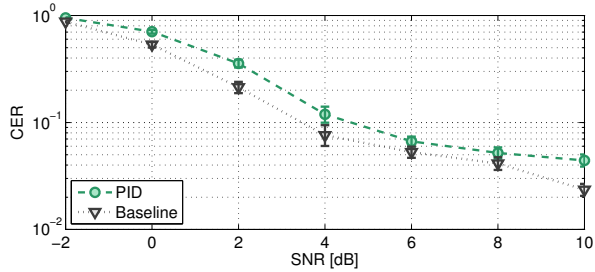

Figure 6.5: Worst case scenario to assess capture effect: One equal power interferer using the same code as the UOI. Further, the interferer is always present and starts its transmission at about the same time as the UOI. Still, the probability that we acquire neither of the two is below 5%.

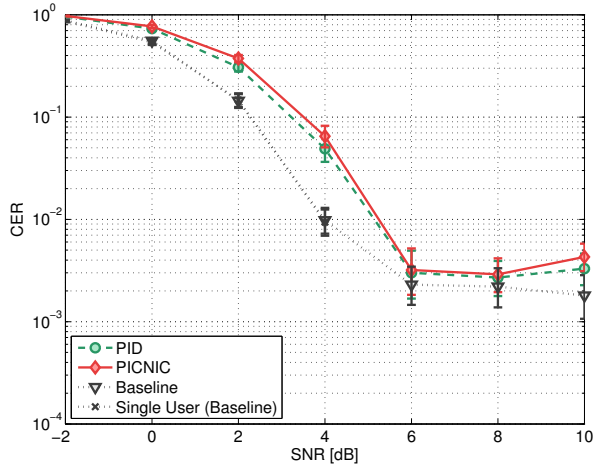

Figure 6.6: Probability of neither acquiring the UOI nor an interferer for two interferers with same code and same power level as the UOI. All algorithms show a good capture effect.

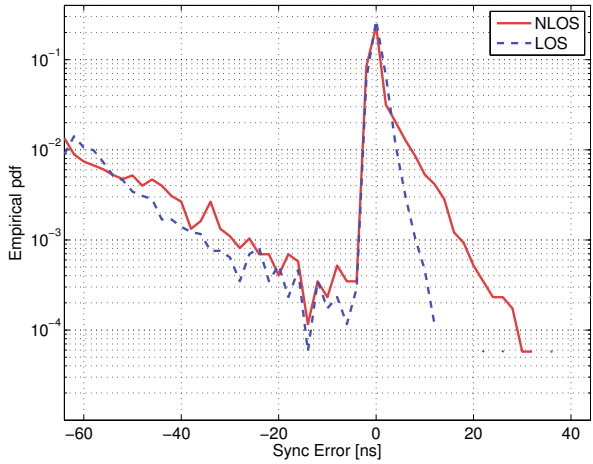

Figure 6.7: Empirical PDF of synchronization error after fine timing acquisition with the baseline algorithm. Results shown are for the near-far configuration at $SNR = 10$ dB. FAs occur at the beginning of the search-back window of size $W = M/2 = 64$ ns.

with higher power is usually acquired. The hardest case is when these two packets have similar power levels. The verification phase may then never succeeds because the receiver switches back and forth between the two packets. To evaluate this scenario we simulate an equal-power interferer that is always present and always starts at about the same time as the UOI. The results are shown in Figure 6.5. For both baseline and PID (PICNIC is not shown because it coincides with PID if identical preamble codes are used), capture is above 95% at high SNR. One effect that helps here is that even though the two users have the same power level, the received energies are distributed differently because of the different propagation channels. Further, we see that baseline performs even a bit better than PID. We attribute this to the fact that the PID, to some extent, levels out different power levels through the thresholding operation on the correlation input. Figure 6.6 also shows the equal power scenario with identical preamble codes but here again with three users that generate packets according to a Poisson process and use the IEEE 802.15.4a Aloha back-off procedure.

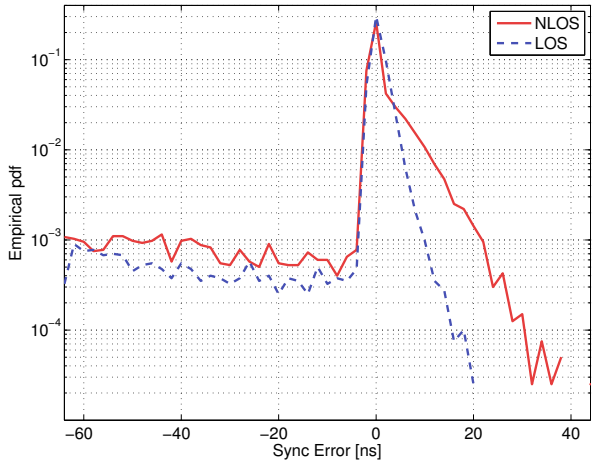

Figure 6.8: Empirical PDF of synchronization error after fine timing acquisition with the PICNIC algorithm (near-far configuration at $SNR = 10$ dB). PICNIC prevents FAs, resulting in more accurate synchronization.

Fine Timing Acquisition

Figures 6.7 and 6.8 show the empirical PDF of the fine timing acquisition error after the jump-back-and-search-forward algorithm for the baseline and PICNIC algorithm, respectively. The PDFs shown are for packets that were correctly acquired during coarse acquisition in the near-far scenario with different codes and at an SNR of 10 dB. The search-back window that we fixed to $W = M/2 = 64$ ns in our simulations, results in the tail at the left of the distribution. We can clearly see in Figure 6.7 that the baseline algorithm is vulnerable to FAs during fine timing acquisition. These FAs are caused by high interference terms and lead the receiver to synchronize to early. Using the PICNIC algorithm clearly limits these FAs, which can be observed in Figure 6.8. We can also see in both cases that synchronization accuracy is better with the LOS channel, which of course is to be expected.

Please note that our algorithm and especially the choice of W is optimized for communication rather than ranging. Still, we get a rather good synchronization accuracy even under severe interference as can be seen from the values summarized in Table 6.1. The values shown in the table are again for the near-far case with different codes at 10 dB. Further, they were obtained with the NLOS channel model. Here we see again that the baseline method performs

Algorithm	RMSE	Mean	50%	75%	90%
Baseline	22.8ns	11.0ns	< 1.1 ns	< 5.5 ns	< 52.6 ns
PID	9.2ns	3.7ns	< 0.9 ns	< 2.9 ns	< 9.1 ns
PICNIC	10.0ns	4.0ns	< 0.9 ns	< 2.9 ns	< 9.5 ns

Table 6.1: Precision of the synchronization for NLOS and two near interferers.

worse because a lot of FAs occur inside window W due to large interfering terms that exceed the threshold. Remember that all the values are calculated over all correctly acquired packets. In the case of PICNIC this includes more packets with interference than in the case of the PID, which explains why PID here seemingly has a slightly better accuracy.

6.4.2 SFD Detection and Overall Synchronization Performance

In the following, we will evaluate the performance of the different SFD detection algorithms. All of the following results include packet detection and timing acquisition, which is done according to either the baseline or the PICNIC algorithm. Therefore, the following results show the overall synchronization performance, which we measure in terms of the synchronization error rate (SER). The SER is the percentage of packets that were missed because of synchronization errors. It includes both, FAs and MDs.

Single User Performance

Figure 6.9 shows the performance of the different SFD detection algorithms if no MUI is present. Timing acquisition is performed according to the baseline algorithm which has the best single user performance.

Interestingly, the performance of the online DESSERT algorithm is almost undistinguishable from the offline algorithm. For both the DESSERT and the offline algorithm, the versions using the soft metric perform slightly better than the hard ones. However, this difference is negligible.

As already pointed out shortly, the offline algorithm using the log-likelihood directly as a metric does not work very well. This is the case because if only the absolute log-likelihood function is considered, a few samples that fit, e.g., the noise-only hypothesis very well are able to dominate the likelihood function leading to detection of the SFD at the wrong position.

Finally, we observe that the algorithms based on sequential decoding outperform the cor-

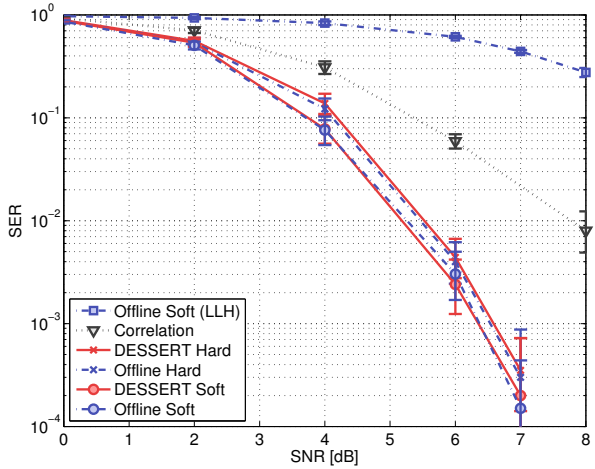

Figure 6.9: Single user performance of different SFD detection algorithms. The online DESSERT algorithms perform very close to their more complicated offline counterparts.

relation based algorithm from Chapter 4. Due to the fact that there is already a considerable performance difference of roughly 2.5 dB in the single user case, we will in the following no longer consider the correlation based algorithm.

Performance under Multi-User Interference

To assess the performance with MUI, we again consider a near-far and an equal power setting with $N_{\mathrm{u}} = 3$ users. We here only consider the four algorithms based on decoding that had similar single user performance. As a reference we also show the performance of a receiver that does not perform any form of interference mitigation. The reference receiver uses the baseline algorithm for timing acquisition and estimates the channel parameters according to Section 5.1.2 which is not robust to MUI. Finally, it performs SFD detection with the DESSERT algorithm using soft decoding but no thresholding.

Figure 6.10 shows the SER for the four algorithms in the near-far setting if no thresholding is used to reject interference during SFD detection. Channel estimation, however, is performed in a robust fashion and so is timing acquisition which is performed according to the PICNIC algorithm. The online algorithms show similar performance and they are not robust to MUI. Robust channel estimation alone gives only a minor performance improvement compared to

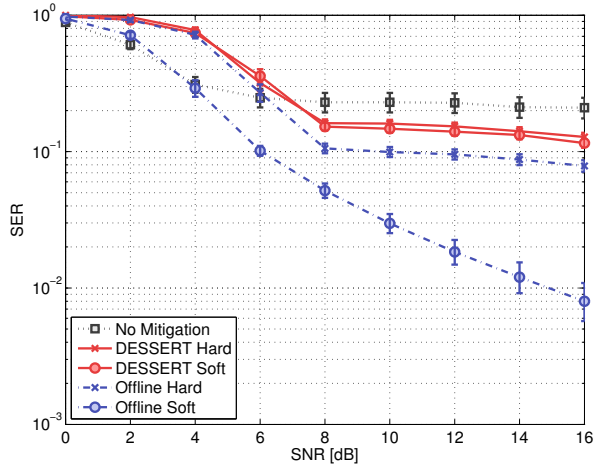

Figure 6.10: Performance of online DESSERT algorithms and their offline counterparts in a near-far setting with 3 users. None of the algorithms uses thresholding to mitigate high interference terms. The only algorithm that is robust to MUI is the offline algorithm using the soft metric.

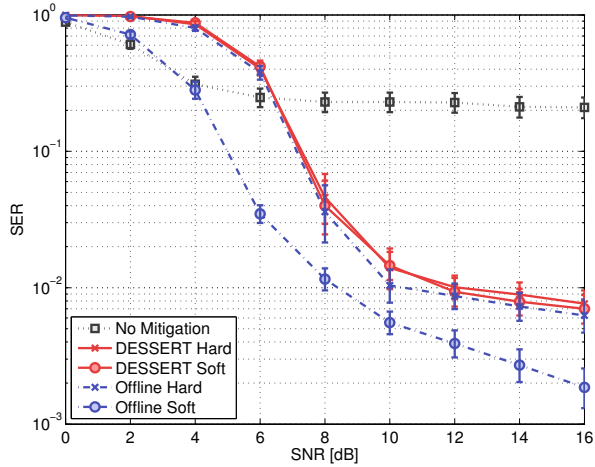

Figure 6.11: Scenario of Figure 6.10 but now all of the algorithms apply a threshold to mitigate interference. They all show a good robustness against MUI but the soft offline algorithm again has a performance advantage throughout the whole SNR range.

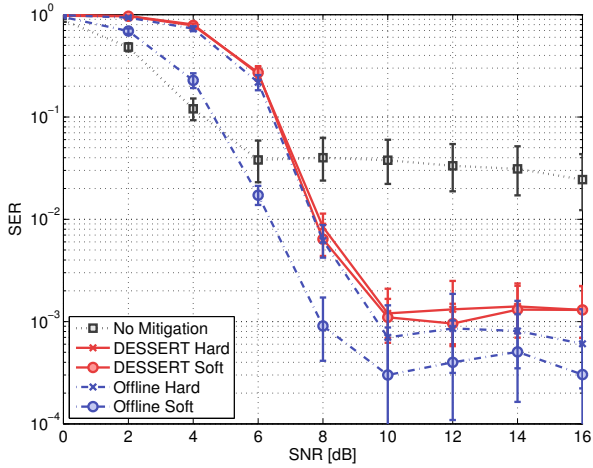

Figure 6.12: Performance of online DESSERT algorithms and their offline counterparts in an equal power setting with 3 users. None of the algorithms uses thresholding to mitigate high interference terms. Still, interference is not high enough to significantly affect SFD detection with any of the four algorithms considered.

the reference receiver, making any gain achieved during timing acquisition void. The offline algorithm based on hard decoding shows a slight gain but is not very robust to MUI either. The offline algorithm using soft decoding on the other hand clearly outperforms all of the other algorithms. It already shows a decent robustness against MUI. At high SNR, it achieves an SER improvement of more than one order of magnitude with respect to the reference receiver. At low SNRs it outperforms the other algorithms by 2 dB.

Figure 6.11 shows again the same near-far settings but this time all of the receivers except for the reference receiver employ a threshold to reject high interference terms during SFD detection. The probability that governs the adaptive threshold is set to $P_{\text{MUI}}^{\text{FA,SFD}} = 0.01$. This value was found through simulations and shows a good performance over various interference scenarios.

The online algorithms profit the most from the use of the threshold. They now also perform over one order of magnitude better than the reference receiver. The offline algorithms also improve. The one using hard decoding performs similar to the online algorithms. The one using soft decoding is now even two orders of magnitude better than the reference receiver in the interference limited SNR region.

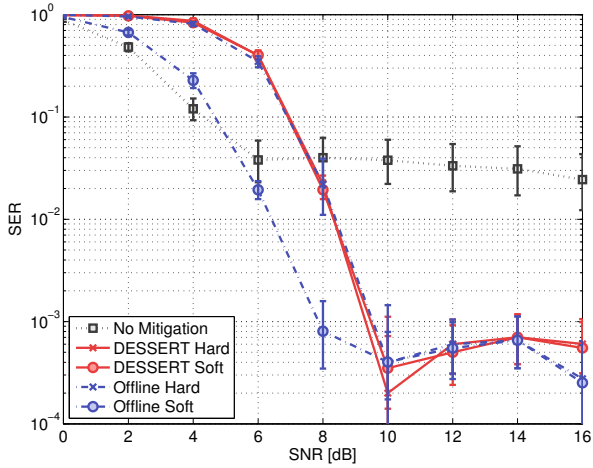

Figure 6.13: Scenario of Figure 6.12 but now all of the algorithms apply a threshold to mitigate interference. With respect to Figure 6.12, a minor improvement for the online algorithms can be observed.

None of the algorithms is severely affected by the weaker interference in the equal power setting. This can be seen in Figure 6.12 where all the algorithms have a decent performance despite the fact that no threshold is used during SFD detection. Further they all perform similarly with the exception of the soft online algorithm that again shows a 2 dB gain at low SNR. Also note that at high SNR at least the offline algorithms already have an SER that is close to the corresponding ACQER of PICNIC shown in Figure 6.4. There are thus hardly any errors due to SFD detection, but most of the remaining errors are due to timing acquisition. Consequently, thresholding does not really improve the performance as can be seen in Figure 6.13. The only slight improvement we see is for the online algorithms. All algorithms achieve a SER improvement of just under two orders of magnitude with respect to the reference receiver that does not use any form of interference mitigation.

As a conclusion, all of the four algorithms considered for robust SFD detection achieve their goal of providing robustness against MUI. However, using the offline algorithm based on hard decoding is not advisable since it performs the same as the online algorithms but at a higher complexity. The two online algorithms are very similar, both in terms of complexity and performance. The reason for similar complexity lies in the fact that both algorithms act on the (soft) receiver output samples and require evaluation of the log-likelihood given by (6.21).

The soft version is in this case preferable since it has a slightly better single user performance. Finally, the soft version of the offline algorithm gives the best performance in every scenario considered. We recommend using this algorithm if the 2 dB gain at low SNR is needed and if the additional complexity with respect to the online algorithms is not an issue.

6.5 Conclusion and Possible Directions for Future Work

We presented PICNIC, a low-complexity algorithm for timing acquisition with an IR-UWB energy-detection receiver in the presence of MUI. PICNIC uses a mixture of thresholding and interference cancelation and outperforms classic timing acquisition algorithms by up to two orders of magnitude if MUI is present. Furthermore, PICNIC exhibits a near perfect capture property: if several transmitters compete for timing acquisition at the receiver, one signal will be acquired with practically no missed detections. Such a property is very desirable in a network where devices transmit concurrently without coordination.

We have also compared different SFD detection algorithms in both single and multi-user settings. SFD detection is often neglected in the related work and we are not aware of any other work comparing different strategies. We proposed a robust online algorithm based on sequential decoding and showed that it can perform close to a more complicated offline version.

The only other interference robust synchronization algorithms for energy-detection receivers that we are aware of are the non-linear filtering techniques used for TOA estimation in [83, 84]. It is easy to see that PICNIC could be combined with these techniques by filtering the signal at the input or output of the correlation. Extending PICNIC in this fashion could be interesting to further increase its robustness against MUI and is a possible direction for future work.

Chapter 7

Interference Mitigation by Statistical Interference Modeling in an Impulse Radio UWB Receiver

IR-UWB networks may be subject to concurrent transmissions without power control, e.g., due to MAC protocols that do not use power control, or coexisiting, non-coordinated piconets. If uncontrolled, impulsive interference stemming from these concurrent transmissions is not properly handled on the physical layer, receiver performance is severely hurt. We have seen this in the case of the energy-detection receiver in Chapter 4. One option to prevent such a performance loss is to mitigate interference through an adaptive thresholding mechanism. With thresholding, the receiver defines a threshold that is based on the estimated power of the signal of interest and the estimated level of the background noise. It then considers samples that lie above the threshold as interference and limits their contribution. Chapter 5 shows that such a mechanism is an effective means to mitigate MUI, even in the case of an energy-detection receiver. Thresholding has also been shown to mitigate interference in coherent receivers [52, 54].

In this chapter we push a little bit further and explore whether interference mitigation can be improved by assuming that MUI follows a given statistical model. It is known that due to its impulsive nature, MUI in an IR-UWB system is not accurately modeled through a Gaussian approximation [39, 41]. A Gaussian model is thus not well suited for our purpose. A popular non-Gaussian model for MUI is the Gaussian mixture model (GMM), see, e.g., [58]. The GMM assumes that the interference has an underlying probability distribution formed by a mixture

of Gaussian distributions with different variances. Each interference term is then assumed to be generated by one of these mixture components. The GMM thus seeks to classify each sample and typically attributes samples with high interference to mixture components with high variances. In [59], the GMM has been proposed as MUI model for IR-UWB and it has been shown how to do channel and interference statistics estimation based on this model. We take the approach taken in [59] one step further and propose a receiver that employs a GMM to mitigate MUI in the decoding phase.

The GMM is a non-Gaussian model without memory. It assumes that the mixture components are independently chosen. However, due to the multipath nature of the channel this is not necessarily true because samples with a high interference level are likely to occur in bursts. Therefore, we propose to introduce correlation by modeling the sequence of mixture components with a homogeneous Markov chain. The resulting MUI model is a hidden Markov model (HMM) where each state is associated with a Gaussian output distribution. The GMM is just a special case of the more general HMM where the choice of the next state is independent of the current state.

We derive the optimal receiver under both models and show that it has the structure of a Rake receiver that penalizes interference terms through an appropriate weighting function. We also show that in packet based systems, mitigation through interference modeling alone is not always sufficient and that some kind of thresholding may still be needed. The resulting receiver employs a combination of statistical interference modeling and thresholding and outperforms solutions that do not mitigate interference or that only rely on a simple threshold. Finally, we find that the HMM is indeed better suited than the GMM to model interference in IR-UWB. However, the resulting performance difference is not huge and comes at the cost of increased receiver complexity.

This chapter is organized as follows: In Section 7.1 we shortly review the system model of the classical IR-UWB PHY that is used in this chapter. In Section 7.2 we introduce the different non-Gaussian MUI models and derive the optimal receiver. Section 7.3 contains the performance evaluation where the different methods and models are compared and their ability to mitigate MUI in a realistic multipath environment is assessed. A conclusion is reached in Section 7.4.

7.1 System Model and Assumptions

In this chapter we are studying a classical IR-UWB PHY with time-hopping and BPSK modulation. We have seen in Section 2.2 that for such a system comprising N_u users, the received signal can be written as

$$r(t) = \sum_{i=0}^{N_\mathrm{s}-1} a_i \tilde{h}(t - c_{\mathrm{THS},i} T_\mathrm{c} - i T_\mathrm{f} - \nu_0) + v(t) \tag{7.1}$$

where $\tilde{h}(t)$ is the compound channel impulse response of the UOI and $v(t)$ accounts for both, MUI from the $N_\mathrm{u} - 1$ interferers and AWGN background noise. PHY transmissions occur in packets of N_s data symbols and $a_i \in \{\pm 1\}$ is the ith symbol of the BPSK modulated data sequence of the UOI. $c_{\mathrm{THS},i}$ is the time-hopping sequence (THS) of the UOI. We assume that each user has its own pseudo-random THS that is known at the receiver. We further require the THS to be constrained such that no inter-symbol interference (ISI) occurs. The THS is integer valued and uniformly distributed on $[0, N_\mathrm{h} - 1]$, where N_h is the number of chips of duration T_c available for time-hopping. The length of a frame equals $T_\mathrm{f} = N_\mathrm{h} \cdot T_\mathrm{c} + T_\mathrm{g} = (N_\mathrm{h} + N_\mathrm{g}) \cdot T_\mathrm{c}$ where N_g is the number of guard slots preventing ISI and T_g is the length of the guard interval. Further, we will in the following assume that synchronization has been achieved by some means and we therefore set $\nu_0 = 0$.

The coherent receiver filters the received signal with a bandlimiting filter of bandwidth B and samples the resulting signal at the Nyquist frequency $1/T = 2B$, yielding the discrete signal

$$y_n = \sum_{i=0}^{N_\mathrm{s}-1} a_i \tilde{h}(nT - c_{\mathrm{THS},i} T_\mathrm{c} - i T_\mathrm{f}) + v_n \tag{7.2}$$

where $v_n = w_n + n_n$ with $n_n \sim \mathcal{N}(0, N_0 B)$ and w_n denoting the contribution of the interferers. Assuming that the guard interval is properly designed to prevent ISI and that T_c is an integer multiple of T, with $T_\mathrm{c} = K \cdot T$, (7.2) can be written as

$$y_n = \sum_{i=0}^{N_\mathrm{s}-1} a_i \sum_{m=0}^{M-1} q_m \cdot \delta_{n - K(c_{\mathrm{THS},i} + i(N_\mathrm{h}+N_\mathrm{g})) - m} + v_n \tag{7.3}$$

where δ_n is the Kronecker delta and $q_m = \tilde{h}(mT)$ is the discrete channel impulse response. With this definition, $M = \left\lceil \frac{T_\mathrm{ch}}{T_\mathrm{c}} \right\rceil$ where T_ch denotes the spread of the compound channel impulse response. Equation (7.3) is the discrete time representation of our system and will be used

throughout the rest of this paper.

7.2 Interference Mitigation

We are now going to introduce our solution to interference mitigation at the receiver. We are following a data-aided approach, meaning that part of the data sequence, the training sequence, is known to the receiver. The N_s data symbols of a physical layer packet are thus divided into two parts: The first N_CH symbols constitute the above mentioned training sequence (which is typically short in comparison to the length of the packet), the remaining N_pay samples contain the information to be transmitted, the payload [1]. Accordingly, our receiver proceeds in two phases: A training phase and a data reception phase. In the training phase, the channel coefficients q_m as well as the statistics of v_n are estimated based on the known training sequence. In the data reception phase, these estimates are then used to mitigate the effect of interference and to recover the unknown data sequence.

7.2.1 Taxonomy of interference types

Before going into details of our receiver design, we analyze the different interference scenarios we are facing. This gives us a better understanding of what can happen and where the challenges are.

If interference occurs during packet reception, it must fall into one of the following three categories:

1. Interference is present during both training and data reception (called interference of type 1 from here on)
2. Interference is present during training only (type 2)
3. Interference present during data reception only (type 3)

If the system we are going to design works perfectly, interference of type 1 should not pose too big of a problem. Ideally we would estimate the interference during the training phase and then deal with it during data reception. Interference of type 2 should do even less harm: we

1. We have seen in the previous chapters that in a complete system, there is at least a third part preceding the training sequence. This part, the synchronization preamble, is used for signal acquisition and synchronization. For simplicity, we here assume that synchronization has already been achieved perfectly and we therefore neglect this part of the packet in the following. However, we still account for the presence of a synchronization preamble in our performance evaluation, see section 7.3.

have estimated it and are thus prepared to face it, but finally it is not even present during data reception. Interference of type 3 however is more difficult to tackle. It is not present during the training phase and we have thus no means to gather any knowledge about it whatsoever. We will thus need some additional mechanism to take care of type 3 interference. We will address estimation and mitigation of interference types 1 and 2 in subsections 7.2.3 and 7.2.4. A possible solution to mitigate interference of type 3 is presented in Section 7.2.4.

7.2.2 Interference Models

We consider two ways of increasing complexity to model interference: the Gaussian mixture model (GMM) and the hidden Markov model (HMM).

In the case of the GMM we assume that the interference and noise samples v_n are i.i.d random variables and that the vector $\mathbf{v} = (v_0, \ldots, v_{N-1})$ has underlying probability distribution

$$f_{\text{GMM}}(\mathbf{v}|\mathbf{\Theta_v}) = \prod_{n=0}^{N-1} \sum_{p=0}^{P-1} \lambda_p \cdot f_{\mathcal{N}}(v_n|\sigma_p^2) \tag{7.4}$$

where $\mathbf{\Theta_v}$ is the vector of model parameters, $\mathbf{\Theta_v} = (\mathbf{\Lambda}, \mathbf{\Sigma}) = (\lambda_0, \ldots, \lambda_{P-1}, \sigma_0^2, \ldots, \sigma_{P-1}^2)$, P is the model order specifying the number of mixture components, λ_p is the prior probability of component p and $f_{\mathcal{N}}(v_n|\sigma_p^2)$ is the pth component density assumed to be zero-mean Gaussian with variance σ_p^2, i.e. $f_{\mathcal{N}}(v_n|\sigma_p^2) = \frac{1}{\sqrt{2\pi\sigma_p^2}} \exp(-v_n^2/2\sigma_p^2)$. The choice of a Gaussian with zero mean is motivated by the fact the we are considering BPSK modulation. However even with a modulation of nonzero mean, a random phase of the channel coefficients leads to samples of zero-mean.

The HMM introduces correlation, the samples v_n are no longer required to be independent. In addition to the sample vector $\mathbf{v}$, we now also have a hidden vector of states, $\mathbf{x} = (x_0, \ldots, x_{N-1})$, with $x_n \in \{0, \ldots, P-1\}$. To each state x_n is associated a Gaussian component density $f_{\mathcal{N}}(v_n|\sigma_{x_n}^2)$ defined as in the GMM case. Each state x_n determines which of the P component densities generates the sample v_n. The initial state is x_0. It is described by the vector of initial state probabilities $\mathbf{\Pi}$ with entries of the form $\pi_j = \Pr(x_0 = j), j \in \{0, \ldots, P-1\}$. Transitions among the states occur according to a matrix of transition probabilities $\mathbf{\Psi}$. An entry $\psi_{j,k} = \Pr(x_n = k|x_{n-1} = j)$ of $\mathbf{\Psi}$ with $j, k \in \{0, \ldots, P-1\}$ is the probability that sample v_n is generated by component density k, knowing that v_{n-1} was generated by component density

j. The vector $\mathbf{v}$ then has the following probability distribution

$$f_{\text{HMM}}(\mathbf{v}|\mathbf{\Theta_v}) = \sum_{\mathbf{x}\in\mathcal{X}} \pi_{x_0} f_{\mathcal{N}}(v_0|\sigma^2_{x_0}) \prod_{n=1}^{N-1} \psi_{x_{n-1}x_n} f_{\mathcal{N}}(v_n|\sigma^2_{x_n}) \tag{7.5}$$

where $\mathcal{X}$ is the space of all possible state vectors and the vector of model parameters is now $\mathbf{\Theta_v} = (\mathbf{\Pi}, \mathbf{\Psi}, \mathbf{\Sigma})$. Note that the GMM is only a special case of the HMM with $\pi_j = \lambda_j$ and $\psi_{j,k} = \lambda_k$.

7.2.3 Training phase

Here we show how to estimate the statistics of interference that is present during the training phase (interference of types 1 and 2). In addition, we also estimate the channel. We thus want to find the maximum-likelihood estimate of $\mathbf{\Theta} = (\mathbf{\Theta_v}, \mathbf{\Theta_c})$ from the training sequence, where $\mathbf{\Theta_v}$ stands for the parameters of the Interference model and $\mathbf{\Theta_c} = (q_0, \ldots, q_{M-1})$ the parameters of the channel [2]. The discrete time received signal is given by (7.3) and we find the sequence v_n during the training phase as

$$v_n = y_n - \sum_{i=0}^{N_{\text{CH}}-1} a_i \sum_{m=0}^{M-1} q_m \cdot \delta_{n-K(c_{\text{THS},i}+i(N_\text{h}+N_\text{g}))-m} \tag{7.6}$$

with $n = 0, \ldots, N_{\text{train}} - 1$ and $N_{\text{train}} = N_{\text{CH}} \cdot K \cdot (N_\text{h} + N_\text{g})$ the number of samples during the training phase. Note that v_n depends on the channel parameters. However, to ease notation, we will write v_n instead of $v_n(\mathbf{\Theta_c})$ whenever possible. We can now formulate the maximum-likelihood estimation problem as

$$\hat{\mathbf{\Theta}} = \arg\max_{\mathbf{\Theta}} \ln(f(\mathbf{v}|\mathbf{\Theta})) \tag{7.7}$$

where we chose to maximize the log-likelihood rather than the likelihood because it simplifies expressions and where f is replaced by (7.4) or (7.5) depending on the interference model. In general, direct maximization of (7.7) is difficult and the classical method of choice is the EM-algorithm [172]. The EM-algorithm is an iterative algorithm used to find the maximum-likelihood parameter estimate in situations where optimization of the likelihood function is simplified by assuming the existence of hidden data $\mathbf{x}$ in addition to the observation $\mathbf{v}$. The

2. In practice, the number of channel parameters can be significant. We will therefore not be able to estimate all of them, and accept to only estimate the first few ones.

complete-data log-likelihood is then defined as $\ln(f(\mathbf{v}, \mathbf{x}|\boldsymbol{\Theta}))$.

In what follows we will denote current parameter estimates with a prime, e.g., $\hat{\lambda}'_p$ denotes the currently available estimate of $\hat{\lambda}_p$.

The EM-algorithm loops over the following steps:

1. E-Step: calculate the conditional expectation of the complete-data log-likelihood with respect to the hidden data $\mathbf{x}$ given the observed data $\mathbf{v}$ and the current parameter estimate $\hat{\boldsymbol{\Theta}}'$

$$Q(\boldsymbol{\Theta}, \hat{\boldsymbol{\Theta}}') = E\big[\ln(f(\mathbf{v}, \mathbf{x}|\boldsymbol{\Theta}))\big|\mathbf{v}, \hat{\boldsymbol{\Theta}}'\big] \tag{7.8}$$

2. M-Step: find the new parameter estimate as

$$\hat{\boldsymbol{\Theta}} = \arg\max_{\boldsymbol{\Theta}} Q(\boldsymbol{\Theta}, \hat{\boldsymbol{\Theta}}') \tag{7.9}$$

The parameter estimate found with the EM-algorithm is the solution of (7.7). Parameter Estimation for GMM and HMM by EM is a well-known and widely used procedure (see e.g. [173]). We now give the resulting algorithms for the two models under consideration.

EM-algorithm for the GMM

The hidden data sequence is $\mathbf{x} = (x_0, \ldots, x_{N_{\text{train}}-1})$ where $x_n \in \{0, \ldots, P-1\}$ indicates which component density generated generated sample v_n. The random variables x_n are thus i.i.d. with $\Pr(x_n = p) = \lambda_p$. Following similar procedures as the ones described in [59, 173], the algorithm looping over the following steps results

1. E-Step 1: calculate $\gamma_p(n) = \Pr(x_n = p|v_n, \hat{\boldsymbol{\Theta}}'_{\mathbf{v}}, \hat{\boldsymbol{\Theta}}'_{\mathbf{c}})$

2. M-Step 1: find the new parameter estimate $\hat{\boldsymbol{\Theta}}_{\mathbf{v}}$ as

$$\hat{\boldsymbol{\Theta}}_{\mathbf{v}} = \arg\max_{\boldsymbol{\Theta}_{\mathbf{v}}} Q((\boldsymbol{\Theta}_{\mathbf{v}}, \hat{\boldsymbol{\Theta}}'_{\mathbf{c}}), (\hat{\boldsymbol{\Theta}}'_{\mathbf{v}}, \hat{\boldsymbol{\Theta}}'_{\mathbf{c}}))$$

   resulting in

$$\hat{\lambda}_p = \frac{1}{N}\sum_{n=0}^{N_{\text{train}}-1} \gamma_p(n), \quad \hat{\sigma}_p^2 = \frac{\sum_{n=0}^{N_{\text{train}}-1} v_n(\hat{\boldsymbol{\Theta}}_{\mathbf{c}})^2 \cdot \gamma_p(n)}{\sum_{n=0}^{N_{\text{train}}-1} \gamma_p(n)} \tag{7.10}$$

3. E-Step 2: calculate $\gamma_p(n) = \Pr(x_n = p|v_n, \hat{\boldsymbol{\Theta}}_{\mathbf{v}}, \hat{\boldsymbol{\Theta}}'_{\mathbf{c}})$

4. M-Step 2: compute $\hat{q}_0, \ldots, \hat{q}_{M-1}$ sequentially by

$$\begin{aligned}\hat{q}_0 &= \arg\max_{q_0} Q((\hat{\boldsymbol{\Theta}}_{\mathbf{v}}, (q_0, \hat{q}'_1, \ldots, \hat{q}'_{M-1})), (\hat{\boldsymbol{\Theta}}_{\mathbf{v}}, \hat{\boldsymbol{\Theta}}'_{\mathbf{c}})) \\ \hat{q}_1 &= \arg\max_{q_1} Q((\hat{\boldsymbol{\Theta}}_{\mathbf{v}}, (\hat{q}_0, q_1, \ldots, \hat{q}'_{M-1})), (\hat{\boldsymbol{\Theta}}_{\mathbf{v}}, \hat{\boldsymbol{\Theta}}'_{\mathbf{c}}))\end{aligned}$$

and so on, resulting in

$$\hat{q}_m = \frac{\sum_{n=0}^{N_{\text{train}}-1} r_{n,m} s_{n,m} \sum_{p=0}^{P-1} \frac{\gamma_p(n)}{\hat{\sigma}_p^2}}{\sum_{n=0}^{N_{\text{train}}-1} s_{n,m}^2 \sum_{p=0}^{P-1} \frac{\gamma_p(n)}{\hat{\sigma}_p^2}} \tag{7.11}$$

In the above equations, $\gamma_p(n)$ can be interpreted as the posterior probability that the interference and noise sample v_n was drawn from the Gaussian mixture with variance σ_p^2. It is given by

$$\gamma_p(n) = \Pr(x_n = p | v_n, \hat{\boldsymbol{\Theta}}') = \frac{\lambda'_p \cdot f_{\mathcal{N}}(v_n(\hat{\boldsymbol{\Theta}}'_{\mathbf{c}}) | \sigma_p^{2\prime})}{\sum_{\tilde{p}=0}^{P} \lambda'_{\tilde{p}} \cdot f_{\mathcal{N}}(v_n(\hat{\boldsymbol{\Theta}}'_{\mathbf{c}}) | \sigma_{\tilde{p}}^{2\prime})} \tag{7.12}$$

Finally, it can be shown that the two remaining quantities are given by

$$r_{n,m} = y_n - \sum_{i=0}^{N_{\text{CH}}-1} a_i \sum_{\substack{\tilde{m}=0 \\ \tilde{m} \neq m}}^{M-1} \hat{q}_{\tilde{m}} \cdot \delta_{n-K(c_{\text{THS},i}+i(N_{\text{h}}+N_{\text{g}}))-\tilde{m}} \tag{7.13}$$

and

$$s_{n,m} = \sum_{i=0}^{N_{\text{CH}}-1} a_i \cdot \delta_{n-K(c_{\text{THS},i}+i(N_{\text{h}}+N_{\text{g}}))-m}. \tag{7.14}$$

Note that this is a modified version of the EM-algorithm that alternates between updating the interference model parameters $\hat{\boldsymbol{\Theta}}_{\mathbf{v}}$ and updating the channel parameters $\hat{\boldsymbol{\Theta}}_{\mathbf{c}}$. This simplifies the joint maximization of the parameters and is known as the space-alternating generalized EM-algorithm (SAGE) [174]. Also note that for the same reason the joint maximization of the channel parameters was replaced by a sequentially updating heuristic.

EM-algorithm for the HMM

In case of the HMM the hidden data sequence $\mathbf{x} = (x_0, \ldots, x_{N_{\text{train}}-1})$ corresponds to the hidden sequence of states. The random variables $x_n \in \{0, \ldots, P-1\}$ thus form a homogeneous Markov chain with initial state probabilities $\boldsymbol{\Pi}$ and transition matrix $\boldsymbol{\Psi}$. The resulting algorithm turns out to have the same structure as in the case of the GMM. For a derivation of the update

equations, the reader is referred to [173]. The update equations are

$$\hat{\pi}_p = \gamma_p(0), \quad \hat{\psi}_{j,k} = \frac{\sum_{n=0}^{N_{\text{train}}-1} \xi_{j,k}(n)}{\sum_{n=0}^{N_{\text{train}}-1} \gamma_p(n)} \tag{7.15}$$

$\hat{\sigma}_p^2$ and $\hat{q}_m$ are found to be given again by (7.10) and (7.11). As opposed to the GMM we now have to calculate two quantities, $\gamma_p(n) = \Pr(x_n = p|v_n, \hat{\mathbf{\Theta}}')$ and $\xi_{j,k}(n) = \Pr(x_{n-1} = j, x_n = k|v_n, \hat{\mathbf{\Theta}}')$, during the E-step. Further there exists no closed form expression for the above quantities. However they can be determined via an iterative method known as the forward-backward or Baum-Welch algorithm [175].

7.2.4 Data Reception Phase

After having estimated the parameters of the interference model and the channel parameters, we can now use these estimates to recover the data sequence. In the following we are first going to show how interference of types 1 and 2 can be mitigated. This is done by using a decoder that takes advantage of the estimated interference model. We will then explain why this interference mitigation procedure is not effective against interference of type 3 and propose a possible solution. At the end of this subsection we will finally give the complete algorithm for the data reception phase.

Decoding using statistical interference modeling

Similar to the training phase, the sequence of the noise and interference terms during data reception is given by

$$v_n = y_{n+N_{\text{train}}} - \sum_{i=N_{\text{CH}}}^{N_{\text{s}}-1} a_i \sum_{m=0}^{M-1} q_m \cdot \delta_{n+N_{\text{train}}-K(c_{\text{THS},i}+i(N_{\text{h}}+N_{\text{g}}))-m} \tag{7.16}$$

with $n = 0, \ldots, N_{\text{data}}-1$ where $N_{\text{data}} = N_{\text{pay}} \cdot K \cdot (N_{\text{h}}+N_{\text{g}})$ is the number of samples during the data reception phase. Assuming that the information symbols are equiprobable, the optimum decoding rule to recover the data sequence $\mathbf{a} = (a_{N_{\text{CH}}}, \ldots, a_{N_{\text{s}}-1})$ is the maximum likelihood criterion. We thus want to find the data sequence $\hat{\mathbf{a}} = (\hat{a}_{N_{\text{CH}}}, \ldots, \hat{a}_{N_{\text{s}}-1})$ that maximizes the likelihood of observing the sequence $\mathbf{v} = (v_0, \ldots, v_{N_{\text{data}}-1})$

$$\hat{\mathbf{a}} = \arg\max_{\mathbf{a}} \ln(f(\mathbf{v}|\mathbf{a})) \tag{7.17}$$

Since the distribution of $\mathbf{v}$ is given by (7.4) or (7.5), depending on the model, this places us again in the framework of the EM-algorithm, which has been shown in [79] for a general context. Combining (7.8) and (7.16) and dropping all the terms not depending on $\mathbf{a}$, we obtain for both models

$$\tilde{Q}(\mathbf{a}, \hat{\mathbf{a}}') = \sum_{i=N_{\mathrm{CH}}}^{N_{\mathrm{s}}-1} \sum_{m=0}^{M-1} \sum_{p=0}^{P-1} q_m \cdot y_{t_{i,m}} \cdot a_i \cdot \frac{\gamma_p(t_{i,m})}{\sigma_p^2} \tag{7.18}$$

where $\gamma_p(n) = \Pr(x_n = p | v_n, \hat{\mathbf{a}}')$ can again be calculated according to (7.12) and $t_{i,m} = K(c_{\mathrm{THS},i} + i(N_{\mathrm{h}} + N_{\mathrm{g}})) + m$. Note that for the GMM the E-step remains exactly the same as in the training phase, only that the parameter to estimate is now $\mathbf{a}$ instead of $\mathbf{\Theta}$. For the case of the HMM the E-step is also similar to the training phase but we no longer have to calculate $\xi_{j,k}(n)$. The M-step is the same for both models, but different from the training phase as we are now maximizing (7.18) with respect to $\mathbf{a}$. We see from (7.18) that the maximization reduces to a classical max-sum problem, that can be solved by the Viterbi algorithm [176] with branch metric

$$m(a_i) = \sum_{m=0}^{M-1} \sum_{p=0}^{P-1} q_m \cdot y_{t_{i,m}} \cdot a_i \cdot \frac{\gamma_p(t_{i,m})}{\sigma_p^2} \tag{7.19}$$

Note that the receiver discussed here implicitly has the structure of a Rake receiver performing maximum-ratio combining. This can be seen from (7.19): All of the estimated components of the multipath channel contribute to the detection of the symbol. Further they are weighted according to their path gain, q_m, and an additional factor, $\gamma_p(t_{i,m})/\sigma_p^2$, accounting for interference. From this observation we can get a good intuition on how the algorithm actually mitigates interference. Recall that $\gamma_p(t_{i,m})$ is the posterior probability that the sample $y_{t_{i,m}}$ has an interference and noise term drawn from a zero-mean Gaussian with variance σ_p^2. In the above metric, samples with noise terms that stem with high probability from a distribution with high variance consequently get penalized through the factor $\gamma_p(t_{i,m})/\sigma_p^2$. This factor plays the role of a weighting function: Samples with low interference level get a larger weight than samples that are likely to be polluted by a high interference and noise term. The functioning of this interference mitigation technique is illustrated in Figure 7.1.

A similar principle applies to the estimation of the channel coefficients during the training phase. The same weighting factor appears in (7.11). Therefore we will also have a more reliable channel estimate when interference mitigation is performed, which has been shown in [59] for the case of the GMM. Interference of type 1 thus gets mitigated through the weighting function in (7.19) as well as indirectly through the better channel estimate. When facing interference of type 2, the weighting factor has less impact. In this case we mostly benefit from the better

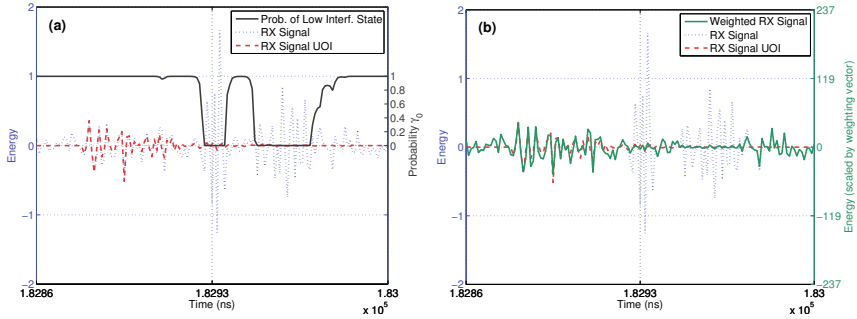

Figure 7.1: Here we show how an algorithm based on interference modeling performs interference mitigation. A two-state hidden Markov model is assumed for the MUI. In Figure (a), one pulse of the received signal and its component corresponding to the user of interest is shown. For each sample, the receiver estimates the probability that it has a low contribution from interfering users (low interference state) or that it is polluted with a high interference term (high interference state). The estimated probability of being in the low interference state is also shown in Figure (a). We can see that the algorithm nicely identifies the part that suffers from a high interference term. Based on this estimation, the receiver designs a weighting vector that is applied to the received signal. Figure (b) additionally shows the received signal after it has been multiplied with the weighting vector, and we can see that the MUI has been successfully removed.

channel estimate.

A final remark is needed for the GMM. Due to the GMM having independent samples, the above iterative decoding procedure is not necessarily required. The reason is that in this case a direct, independent evaluation of the decision rule

$$\ln(f_{\text{GMM}}(\mathbf{v}|\hat{a}_i = +1)) \underset{\hat{a}_i=-1}{\overset{\hat{a}_i=+1}{\gtrless}} \ln(f_{\text{GMM}}(\mathbf{v}|\hat{a}_i = -1)) \tag{7.20}$$

is possible for every symbol a_i. This yields the following decision rule

$$\sum_{m=0}^{M-1} \ln \left[\sum_{p=0}^{P-1} \lambda_p f_{\mathcal{N}}(y_{t_{i,m}} - q_m|\sigma_p^2) \right] \underset{\hat{a}_i=-1}{\overset{\hat{a}_i=+1}{\gtrless}} \sum_{m=0}^{M-1} \ln \left[\sum_{p=0}^{P-1} \lambda_p f_{\mathcal{N}}(y_{t_{i,m}} + q_m|\sigma_p^2) \right] \tag{7.21}$$

that may be used in place of the EM algorithm during data decoding.

Thresholding to mitigate interference type 3

As already mentioned, the situation is different for interference of type 3. It is not present during training and therefore the estimated variances of the interference and noise term will be rather small (in the order of the background noise variance). Samples with a lot of interference will thus still get a relatively high weight. This observation is the key to our solution to mitigate interference of type 3. After the training phase we determine the largest of the estimated variances. Assume this is σ_{p*}^2. We then determine a threshold η_{p*}, such that $\Pr(X \geq \eta_{p*}) \leq P_{\text{MUI}}^{\text{FA}}$, where $X \sim \mathcal{N}(0, \sigma_{p*}^2)$ and $P_{\text{MUI}}^{\text{FA}}$ is some predetermined small probability. After each E-step we set

$$\gamma_p(n) = 0; \quad \forall p, n \text{ such that } v_n > \eta_{p*}$$

This ensures that samples, that with high probability cannot be explained by the estimated interference model, do not contribute to the branch metric. As this is likely to affect predominantly samples polluted by interference of type 3, we have found a way to mitigate this type of interference.

To summarize 1) and 2), this results in the following algorithm for the data reception phase:

1. Initialization: Determine the threshold η_{p*}. Take an initial guess, $\hat{\mathbf{a}}^{(0)}$, for $\mathbf{a}$
2. E-Step: calculate $\gamma_p(n)$ in the same way as in the training phase (i.e. using (7.12) in the case of the GMM, using the forward-backward algorithm in the case of the HMM).
3. Thresholding: set $\gamma_p(n)$ to zero for any sample v_n that lies above the threshold η_{p*}.

4. M-Step: find the new parameter estimate $\hat{\mathbf{a}}$ by means of the Viterbi algorithm with metric given by (7.19).

5. Repeat steps 2 to 4 until convergence

7.3 Performance Evaluation

In this section we evaluate the performance of a receiver using statistical interference modeling to mitigate interference. Through simulations, we compare four different receiver types. Two receivers perform statistical interference modeling according to the GMM or HMM algorithms described in this chapter. Further, we simulate two reference receivers, the first of which uses no interference mitigation. We can view this receiver as a special case of a receiver that uses the GMM, but only with a single state (i.e., P=1). Further, this receiver does not perform the non-linear thresholding operation described in Section 7.2.4. This receiver thus makes a Gaussian approximation, neglecting the impulsive nature of MUI. Finally, we also simulate a receiver that only uses a simple threshold to mitigate interference. This receiver is identical to the receiver that uses no interference mitigation except for the fact that it uses the non-linear thresholding operation.

7.3.1 Simulation Setup and Receiver Parameters

The performance metric we used is the bit error rate (BER) versus signal to noise ratio (SNR, defined as $\frac{E_\mathrm{b}}{N_0}$, where E_b is the received energy per bit) at the receiver. Our simulation setup is the following. $T_\mathrm{c} = 2$ ns, $N_\mathrm{h} = N_\mathrm{g} = 128$, which results in a pulse repetition frequency of 1.95 MHz. Channel coding is performed with a simple pulse repetition code of rate $1/4$, resulting in a physical layer peak data rate of 0.49 Mb/s.

The pulse shape $p(t)$ is chosen to be the second derivative of a Gaussian monopulse given by $p(t) = (1-4\pi(t^2/\tau^2)) \cdot \exp(-2\pi t^2/\tau^2)$ with $\tau = 2.4$ ns, resulting in a signal with a -10 dB bandwidth of about $B = 500$ MHz. Correspondingly we sample with a sampling frequency of 1 GHz, yielding a sampling period $T = 1ns$.

We ran simulations for the IEEE 802.15.4a indoor office LOS and NLOS channel models. Since we did not observe fundamental differences between these two channel models, we only show the results for the NLOS model.

Physical layer data packets are assumed to be generated by a homogeneous Poisson process with rate R. The simulations shown here assume a rate corresponding to one half or one fourth

of the peak data rate. Each packet has a length of 127 bytes or $N_{\text{s}} = 1016$ data symbols. The first $N_{\text{sync}} = 112$ symbols are assumed to be reserved for the synchronization preamble, the next $N_{\text{CH}} = 32$ symbols form the training sequence and the remaining $N_{\text{pay}} = 872$ symbols constitute the payload. Two successive packets are assumed to be at least separated by the duration of a packet of size $N_{\text{sync}} + N_{\text{CH}}$, leaving room for acknowledgements.

To simulate interference, different near-far scenarios with varying numbers of users were chosen. At the receiver, the interferers have power levels of 0 dB, 10 dB and 20 dB with respect to the UOI.

The model order of the interference models is fixed to $P = 2$. We also ran simulations with a model order of $P = 3$. This did not change the results fundamentally, we observed a slight improvement for the HMM and no improvement for the GMM, and hence we do not show these results here.

Initialization of the estimators in the training phase is done as in [59], i.e. $\hat{\lambda}^{(\mathbf{0})} = (0.99, 0.01)$, $\hat{\sigma}_0^{(0)} = \frac{1}{N_{\text{train}}} \sum_{n=0}^{N_{\text{train}}-1} v_n^2$, $\hat{\sigma}_1^{(0)} = 50 \cdot \hat{\sigma}_0^{(0)}$ and $\hat{q}_m^{(0)} = \frac{1}{\sqrt{M}}$. The estimator for the data sequence $\hat{\mathbf{a}}$ is initially simply set to zero. The probability governing the threshold is set to $P_{\text{MUI}}^{\text{FA}} = 10^{-4}$, which was found through simulations. The number of estimated channel parameters is $M = 56$ which at a sampling frequency of 1 GHz roughly corresponds to the channel spread.

7.3.2 Simulation Results

Our simulations show that this 2-state model achieves a significant performance gain over traditional techniques that do not account for MUI or use a simple thresholding mechanism. In Figure 7.2 we show the BER for a near-far scenario with $N_{\text{u}} = 4$ users. The three interferers are received with power levels that exceed the UOI signal by 0 dB, 10 dB and 20 dB, respectively. Rate R is equal to half the peak data rate. Interference modeling achieves a significant gain, independent of the interference model used. The difference between the GMM and the HMM however is much less pronounced. The HMM gets a slight performance gain, indicating that it is better suited to model MUI. However, considering the additional complexity needed by the HMM, the GMM seems to be a more realistic solution.

Further, we find that (1) modeling of interference alone is not sufficient, thresholding is still needed to prevent losses, and (2) these losses are effectively due to type 3 interference. This shows the importance of establishing a classification of different interference types and justifies our approach of mitigating each type accordingly. The corresponding results are shown in Figures 7.3 and 7.4. All of these figures show a near far scenario with three strong interferers,

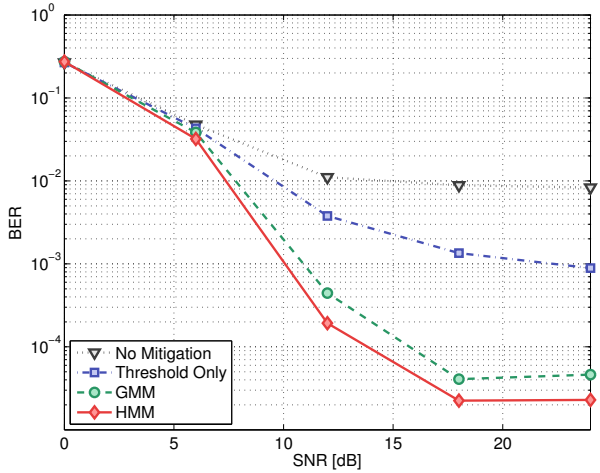

Figure 7.2: We compare our interference mitigation technique with a receiver that neglects multi-user interference (MUI) completely and with a receiver performing only simple thresholding. The curves shown are for a near-far scenario with three interferers. The interferers have received power levels of 0 dB, 10 dB and 20 dB with respect to the UOI. All packets are generated at half the peak data rate. It can be seen, that the performance gain from modeling the interference is significant. Using the more sophisticated HMM to characterize MUI however only gives a small additional gain compared to the GMM model.

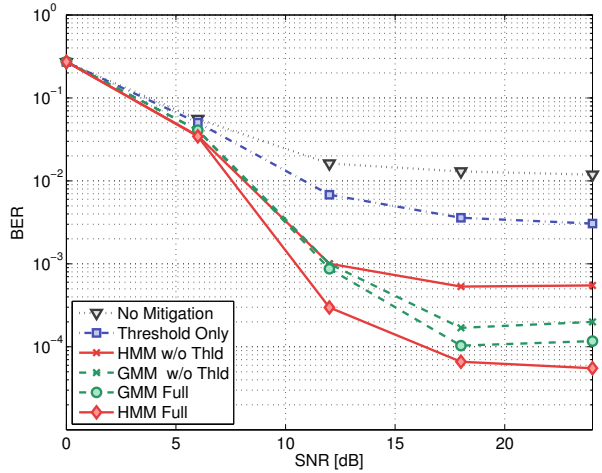

Figure 7.3: Here we show the effect of interference type 3. We use a repetition code of rate $1/4$ and assume packet generation at one fourth of the peak rate. We see that whether or not to perform thresholding greatly affects performance. Especially for the HMM the effect is huge: Without thresholding it barely performs better than the simple thresholding receiver.

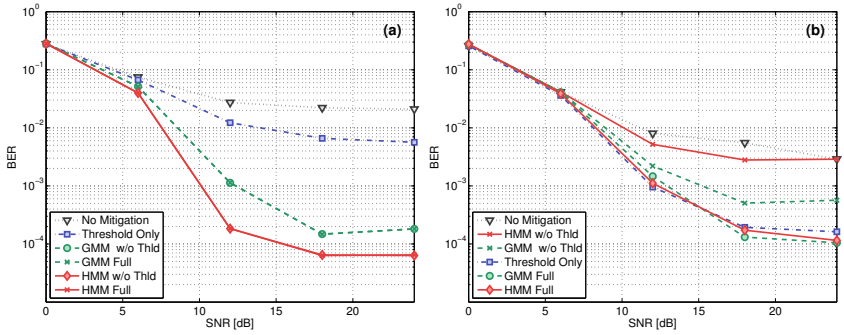

Figure 7.4: Here we have categorized all packets depending on the type of interference they experience and show the performance for types 1 and 3 separately. Results for type 1 are shown in (a), results for type 3 in (b). The simulation setup is the same as in Figure 7.3. The plots confirm our hypothesis that the losses seen in Figure 7.3 when not using thresholding are almost exclusively due to packets that face interference of type 3. Note that for interference type 3 both GMM and HMM without thresholding, perform worse than the simple threshold based receiver, that does not model interference.

each having a power level exceeding the UOI by 20 dB. Packets are generated at one fourth of the peak data rate. Figure 7.3 shows the overall BER. We further show the curves for the GMM and the HMM when thresholding is omitted. We have a clear performance degradation in both cases. The HMM even performs worse than the GMM in this case. We interpret this to be due to the HMM modeling MUI better and therefore being less flexible, when encountering situations where the predicted model does not apply. Figure 7.3 thus shows that the thresholding mechanism is needed even when interference modeling is performed.

To confirm that the performance degradation in Figure 7.3 effectively is due to interference of type 3, we classified all data packets with respect to the interference type they were facing. This allows us to plot the BER curve for each of the types separately, which is shown in Figure 7.4(a) for type 1 and Figure 7.4(b) for type 3. In our classification we considered interference present in a part of the packet, if more than 20% of this part was affected by interference. E.g., interference was considered present during the data reception phase if it was present during at least 20% of data reception. We see that for packets with type 1 interference, the threshold has no impact. Interference of type 1 is thus efficiently reduced by the statistical interference mitigation mechanism. Packets with type 3 interference however suffer from a large loss if the threshold is not applied. Note that in this case both the HMM and the GMM

perform worse than the threshold based receiver. This further confirms that better modeling makes the receiver less flexible in situations where the model no longer applies, which can be overcome by using the proposed thresholding algorithm.

7.4 Conclusion

In this chapter we have presented an effective method to mitigate the effect of MUI in a multipath environment and shown that it outperforms existing techniques significantly. We have introduced a new and more general model for multi-user interference in IR-UWB and have compared with existing models its ability to accurately characterize interference. Finally, we have identified three different types of MUI and the need to deal with each of them in an appropriate manner.

The main goal of this chapter was to understand whether a performance improvement can be achieved by exploiting statistical properties of MUI. There are thus several aspects that have been omitted on purpose. First of all, we did not do a complexity analysis for the choice of the interference model. The HMM requires much more processing than the GMM due to the forward-backward algorithm in the E-step. Also from a viewpoint of energy consumption, the GMM has an advantage. During chips where the signal of interest is not present, the receiver can be turned off. The interference level can later still be estimated as there is no time-dependence between the samples. Despite its complexity, the HMM is still of interest because it is able to accurately model MUI. This makes it suitable for the theoretical analysis of IR-UWB systems with MUI, which has been recognized, e.g., by [76].

Second, we concentrated on coherent receivers and the classical IR-UWB PHY. We will show in the next chapter that interference mitigation techniques can also be applied to the energy-detection receiver and IEEE 802.15.4a. Third, interference of type 3 may be further reduced by re-estimating interference during idle periods. Since our initial work on this topic, this approach has been explored in [60] and we will also use it in Section 8 for the energy-detection receiver. Finally, other non-Gaussian models may be used to model MUI in IR-UWB networks. For a discussion of possible options please refer to Section 3.1.2 in the related work.

Chapter 8

Non-Gaussian Interference Modeling with an Energy-Detection Receiver

In this chapter we revisit the energy-detection receiver and show that interference mitigation techniques based on statistical modeling, similar to those that we have seen in the previous chapter, can also enhance the robustness to MUI of this receiver architecture.

In Chapter 7 we found the Gaussian mixture model (GMM) to be a suitable model for noise and interference in IR-UWB networks. Hence, we also use the GMM in this chapter and show that at the output of the energy-detection receiver this translates to a model with a mixture of gamma distributions. The parameters of this gamma mixture model can further be estimated from the IEEE 802.15.4a preamble without making a lot of sacrifices in terms of receiver complexity. We then derive the optimal energy-detection receiver for the gamma mixture model and show how to perform SFD detection and data demodulation. Furthermore, our design explicitly addresses the time-varying nature of MUI in packet based systems through an algorithm that constantly adapts the MUI model. We have seen in Chapter 7 that if this is not accounted for, the gain from modeling interference can become void. Finally, our performance evaluation demonstrates that this approach yields energy-detection receivers that are very robust to MUI.

Although there is related work on interference mitigation by statistical modeling for coherent receivers (see the related work in Section 3.1.2), this study is the first one for non-coherent energy-detection IR-UWB receivers. Further, our receiver is compliant with the IEEE 802.15.4a standard.

This chapter is organized as follows; in Section 8.1, we review the necessary details about

the IEEE 802.15.4a physical layer needed to understand the results of the performance evaluation. The interference model is introduced in Section 8.2 and estimation of its parameters is addressed. The optimal receiver is derived in Section 8.3. In Section 8.4, we present the results of our performance evaluation. Finally, we conclude this chapter in Section 8.5.

8.1 System Model and Assumptions

In this chapter, we consider an IEEE 802.15.4a IR-UWB physical layer with binary pulse position modulation (BPPM) and non-coherent reception with energy-detection. System model and receiver architecture are thus the same we used in Chapter 5 and in Chapter 6. The details are explained in the system model in Chapter 2. We here merely recall the notation and the most important details.

The received signal during the IEEE 802.15.4a preamble is given by (2.28) and equals

$$r_{\text{pre}}(t) = \underbrace{\sum_{i=0}^{N_{\text{pre}}-1} s_i \sum_{j=0}^{C-1} c_j \cdot \tilde{h}(t-(j+iC)L_{\text{s}}T_{\text{c}}-\nu_0)}_{\tilde{x}_{\text{pre}}(t-\nu_0)} + v(t) \tag{8.1}$$

where $\tilde{x}_{\text{pre}}(t)$ is the signal contribution of the UOI and $v(t)$ accounts for *both* MUI and thermal noise. Thermal noise is modeled as a zero-mean AWGN process (bandlimited to B, the bandwidth of the receiver bandpass filter) with power spectral density (PSD) $N_0/2$. MUI is created by the $N_{\text{u}}-1$ interfering transmitters that use the same PHY.

During the first N_{sync} preamble symbols, the modulation coefficient s_i equals $s_i = 1$, for the remaining $N_{\text{sfd}} = N_{\text{pre}} - N_{\text{sync}}$ symbols it corresponds to the SFD sequence. In this chapter we are not concerned with robust timing acquisition. We can therefore proceed as we did in Section 6.3 of Chapter 6 and assume that the TOA is $\nu_0 = 0$ but that N_{sync} is unknown. For robust timing acquisition, please refer to Chapter 6.

Furthermore, one pulse is sent every L_{s} chips and modulated according to the preamble code $c_j \in \{-1, 0, +1\}, j = 0, 1, \ldots, C-1$. Finally, $\tilde{h}(t)$ denotes the unknown compound channel impulse response.

During the IEEE 802.15.4a payload, we receive the signal

$$r_{\text{pay}}(t) = \underbrace{\sum_{i=0}^{N_{\text{pay}}-1} a_i \sum_{j=0}^{N_{\text{cpb}}-1} b_{ij} \cdot \tilde{h}(t - iT_{\text{f}} - d_i T_{\text{f}}/2 - jT_{\text{c}}) + v(t)}_{\tilde{x}_{\text{pay}}(t)}. \tag{8.2}$$

In contrast to the preamble signal, each of the N_{pay} payload symbols of duration T_{f} is composed of a short, continuous burst of N_{cpb} pulses with pseudo-random polarity, $b_{ij} \in \pm 1$. $d_i \in \{0, 1\}$ denotes the i-th BPPM data bit. Note that with respect to the full model given in (2.30) of Chapter 2 we omitted the time-hopping sequence (THS) since it is perfectly known at the receiver and only complicates the notation. In our performance evaluation, we of course take the THS into account.

The discrete samples at the output of the energy-detection receiver are given by

$$y_m = \int_{mT}^{(m+1)T} [r(t)]^2 \, \mathrm{d}t \tag{8.3}$$

where $r(t)$ corresponds to either (8.1) or (8.2), depending on whether we are receiving the preamble or the payload, and the integration time is of the form $T = \frac{L_{\text{a}}}{M} T_{\text{c}}$.

8.2 Interference Modeling and Parameter Estimation

In this section we first describe the interference model used by our receiver, then how the parameters of the model can be estimated, and finally, a practical low-complexity parameter estimation procedure.

8.2.1 Distribution of Receiver Output without MUI

We recall that when $v(t)$ comprises only AWGN, the probability density function (PDF) of the receiver output y_m can be closely approximated (see Section 2.3.3) with a scaled non-central chi-square distribution

$$f(y_m | N_0/2, BT, p_m) = \frac{1}{N_0/2} f_{\text{NC}\chi^2}\Big(\frac{y_m}{N_0/2}\Big| 2BT, \frac{p_m}{N_0/2}\Big) \tag{8.4}$$

with $2BT$ degrees of freedom and where the non-centrality parameter depends on the coefficients

$$p_m = \int_{mT}^{(m+1)T} [\tilde{x}(t)]^2 \, \mathrm{d}t. \tag{8.5}$$

The received signal contribution of the UOI, $\tilde{x}(t)$, corresponds to either $\tilde{x}_{\text{pre}}(t)$ or $\tilde{x}_{\text{pay}}(t)$. The coefficients p_m depend on the channel energy-delay profile of the UOI signal and they can be estimated during reception of the preamble, when the preamble code elements are nonzero, i.e., $c_j \neq 0$ and $\tilde{x}_{\text{pre}}(t) \neq 0$. We have seen receiver architectures and algorithms that allow for a robust estimation of p_m in the presence of MUI in Chapter 5. We therefore assume in the following that the coefficients p_m are known.

8.2.2 Modeling of Receiver Output Distribution with MUI

If MUI is present, the AWGN assumption for $v(t)$ does no longer hold and Gaussian approximations for $v(t)$ are inaccurate in the case of IR-UWB physical layers [39, 41]. With coherent receivers, a widely used non-Gaussian model for $v(t)$ is the Gaussian-mixture model (GMM) that we already encountered in Chapter 7. In the following we demonstrate that this model can also be successfully used in the case of energy-detection receivers. In the next paragraph, we extend the GMM to energy-detection receivers and derive the associated receiver output.

Let us assume a model with P zero-mean Gaussian component processes

$$v_p(t) \in \{v_0(t), v_1(t), \ldots, v_{P-1}(t)\}.$$

The power spectral density (PSD) of every component process $v_p(t)$ is σ_p^2 and its prior probability is λ_p. Hence, under the Gaussian-mixture model, the noise contribution of every sample y_m is generated with probability λ_p by the component process $v_p(t)$ with PSD σ_p^2. The parameters of the model are

$$\boldsymbol{\Theta} = (\boldsymbol{\Lambda}, \boldsymbol{\Sigma}) = (\lambda_0, \lambda_1, \ldots, \lambda_{P-1}, \sigma_0^2, \sigma_1^2, \ldots, \sigma_{P-1}^2)$$

with $\sum_{p=0}^{P-1} \lambda_p = 1$. Making the same approximations as for equation (8.4), the PDF of the receiver output y_m can be written as

$$f_{NC\chi^2\text{Mix}}(y_m|\boldsymbol{\Theta}, BT, p_m) = \sum_{p=0}^{P-1} \lambda_p \cdot \frac{1}{\sigma_p^2} f_{NC\chi^2}\Big(\frac{y_m}{\sigma_p^2}\Big|2BT, \frac{p_m}{\sigma_p^2}\Big) \tag{8.6}$$

which corresponds to a mixture of scaled non-central chi-square distributions.

8.2.3 Estimation of the Model Parameters

As stated in Section 8.2.1, the energy-delay profile or, equivalently, the coefficients p_m can be estimated during the reception of the preamble when the signal of the UOI $\tilde{x}_{\text{pre}}(t) \neq 0$. We next show that the parameters $\boldsymbol{\Theta}$ can be estimated from samples of the preamble where $\tilde{x}_{\text{pre}}(t) = 0$ (no contribution from the UOI). This corresponds to the set of $\bar{M} = N_{\text{sync}} \cdot (C - \sum_{j=0}^{C-1} c_j^2) \cdot M$ samples

$$\begin{aligned} \bar{\mathbf{y}} &= (\bar{y}_0, \ldots, \bar{y}_{\bar{M}-1}) \\ &= \left\{ y_m |\ m \in \{0, 1, \ldots, N_{\text{sync}} CM - 1\}, \quad c_{\lfloor \frac{m}{M} \rfloor \bmod C} = 0 \right\} \end{aligned} \tag{8.7}$$

For all $\bar{M}$ samples, $p_m = 0$ and it follows from (8.6) that their joint PDF is given by a mixture of gamma distributions (see also Section 2.3.3)

$$f_{\Gamma\text{Mix}}(\bar{\mathbf{y}}|\boldsymbol{\Theta}, BT) = \prod_{m=0}^{\bar{M}-1} \sum_{p=0}^{P-1} \lambda_p \cdot f_\Gamma(\bar{y}_m | BT, 2\sigma_p^2) \tag{8.8}$$

where

$$f_\Gamma(y|\kappa, \theta) = \frac{1}{\theta^\kappa \Gamma(\kappa)} y^{\kappa-1} e^{-y/\theta} \tag{8.9}$$

is the PDF of the gamma distribution with scale parameter k and shape parameter θ.

We have transformed the problem of estimating the parameters of the model in (8.6) to the problem of estimating the parameters of the Gamma Mixture Model given by (8.8). Note that the payload also contains parts where the receiver knows that no contribution of the UOI is present. This is due to the THS and the guard intervals that are introduced in the payload to prevent inter-symbol interference (ISI). An analogous derivation that is omitted here can thus be made for these payload parts.

To find the maximum-likelihood estimate

$$\hat{\boldsymbol{\Theta}} = \arg\max_{\boldsymbol{\Theta}} \ln(f_{\Gamma\text{Mix}}(\bar{\mathbf{y}}|\boldsymbol{\Theta}, BT)) \tag{8.10}$$

of the parameters $\boldsymbol{\Theta}$, we use the classical framework for mixture-density parameter estimation, which is commonly done by using the iterative Expectation-Maximization (EM) algorithm [172]. The EM-algorithm starts with an initial guess $\hat{\boldsymbol{\Theta}}^{(0)} = (\hat{\lambda}_0^{(0)}, \ldots, \hat{\lambda}_{P-1}^{(0)}, \hat{\sigma}_0^{2(0)}, \ldots, \hat{\sigma}_{P-1}^{2(0)})$

of the parameters to be estimated and updates the estimate in every iteration, guaranteeing that the log-likelihood always increases. Following similar derivations as in [173], we find the following update equations for our model:

$$\hat{\lambda}_p^{(k+1)} = \frac{1}{\bar{M}} \sum_{m=0}^{\bar{M}-1} \gamma_p^{(k)}(m), \quad \hat{\sigma}_p^{2(k+1)} = \frac{\sum_{m=0}^{\bar{M}-1} \bar{y}_m \cdot \gamma_p^{(k)}(m)}{2BT \sum_{m=0}^{\bar{M}-1} \gamma_p^{(k)}(m)} \tag{8.11}$$

where

$$\gamma_p^{(k)}(m) = \frac{\hat{\lambda}_p^{(k)} \cdot f_\Gamma(\bar{y}_m | BT, 2\hat{\sigma}_p^{2(k)})}{\sum_{\tilde{p}=0}^{P} \hat{\lambda}_{\tilde{p}}^{(k)} \cdot f_\Gamma(\bar{y}_m | BT, 2\hat{\sigma}_{\tilde{p}}^{2(k)})} \tag{8.12}$$

can be interpreted as the posterior probability that a sample $\bar{y}_m$ was drawn from the gamma component density with scale parameter $2\sigma_p^2$.

8.2.4 Low-complexity Recursive Formulation of EM-algorithm

Due to its iterative nature, the EM-algorithm given by (8.11) and (8.12), is both time and memory consuming. The algorithm must wait for all of the $\bar{M}$ samples to become available and then processes all of them simultaneously. This is impractical for a low-complexity implementation. Therefore, we resort to a simplified recursive version [177] of the EM-algorithm that yields fast convergence at a moderate complexity. A similar approach has been taken in [60] with a coherent receiver. The recursive algorithm treats $\bar{\mathbf{y}}$ in an online fashion by blocks of $\tilde{M} < \bar{M}$ consecutive samples. At the k-th recursion, the k-th block $\bar{\mathbf{y}}^{(k)} = (\bar{y}_{k\tilde{M}}, \ldots, \bar{y}_{(k+1)\tilde{M}-1})$ is processed. In its recursive formulation, the update equation for the parameters becomes

$$\hat{\mathbf{\Theta}}^{(k+1)} = \left(1 - \frac{1}{(k+1)^\epsilon}\right) \hat{\mathbf{\Theta}}^{(k)} + \frac{1}{(k+1)^\epsilon} \tilde{\mathbf{\Theta}}^{(k+1)} \tag{8.13}$$

where $\epsilon \in (0,1)$ is a forgetting factor and $\tilde{\mathbf{\Theta}}^{(k)} = (\tilde{\lambda}_0^{(k)}, \ldots, \tilde{\lambda}_{P-1}^{(k)}, \tilde{\sigma}_0^{2(k)}, \ldots, \tilde{\sigma}_{P-1}^{2(k)})$ with

$$\tilde{\lambda}_p^{(k+1)} = \frac{1}{\tilde{M}} \sum_{m=0}^{\tilde{M}-1} \gamma_p^{(k)}(k\tilde{M} + m) \tag{8.14}$$

$$\tilde{\sigma}_p^{2(k+1)} = \frac{\sum_{m=0}^{\tilde{M}-1} \bar{y}_{k\tilde{M}+m} \cdot \gamma_p^{(k)}(k\tilde{M} + m)}{2BT \sum_{m=0}^{\tilde{M}-1} \gamma_p^{(k)}(k\tilde{M} + m)} \tag{8.15}$$

With respect to the iterative version, complexity is greatly reduced: (8.12) is only computed once for every sample and the only values that need to be kept in memory are the number

of iterations performed, the current parameter estimate, the sum in (8.14) and its weighted counterpart in the numerator of (8.15).

8.2.5 Initialization, Adaptive Model Order

MUI is a time-varying process. It may be present during some parts of a packet and absent during others. Our receiver therefore adapts the model order to these varying conditions according to the following procedure:

Initialize

Start with a single component ($P = 1$) and $\hat{\mathbf{\Theta}}^{(0)} = (\hat{\lambda}_0^{(0)}, \hat{\sigma}_0^{2(0)}) = (1, \frac{1}{\tilde{M}} \sum_{m=0}^{\tilde{M}-1} \bar{y}_m)$

Add

Add a new component to the mixture at iteration k if there are samples in $\bar{\mathbf{y}}^{(k)}$ that are not well explained by the model. We consider a sample to be not well explained if it lies above the threshold $\eta_{p*}^{(k)} = F^{-1}_{\Gamma_{p*}^{(k)}}(0.99)$ where $F_{\Gamma_{p*}^{(k)}}(x)$ is the cumulative distribution function of the mixture component $p*$ which has the highest associated noise variance $\hat{\sigma}_{p*}^{2(k)}$. Let $\tilde{\mathbf{y}}^{(k)} = (\tilde{y}_0, \ldots, \tilde{y}_{L-1})$ denote the samples that fulfill this condition. The new state then has initial parameters $\hat{\lambda}_P^{(0)} = L/\tilde{M}$ and $\hat{\sigma}_P^{2(0)} = \frac{1}{L} \sum_{l=0}^{L-1} \tilde{y}_l$. The existing component probabilities $\hat{\lambda}_0^{(k)}, \ldots, \hat{\lambda}_{P-1}^{(k)}$ are scaled accordingly such that all the probabilities sum to one. Adding new components in this way ensures that interference that was not present before can still be appropriately mitigated and does not lead to a performance degradation like we have observed it for interference type 3 in Chapter 7.

Merge

Two components x_i and x_j are merged if their associated noise variances are too close. They are considered too close if $F^{-1}_{\Gamma_{x_i}^{(k)}}(0.4) < F^{-1}_{\Gamma_{x_j}^{(k)}}(0.6)$ while $\hat{\sigma}_{x_i}^{2(k)} > \hat{\sigma}_{x_j}^{2(k)}$. The state resulting from the merging operation, has parameters $\hat{\lambda}_p^{(k)} = \hat{\lambda}_{x_i}^{(k)} + \hat{\lambda}_{x_j}^{(k)}$ and $\hat{\sigma}_p^{2(k)} = \frac{\hat{\lambda}_{x_i}^{(k)} \hat{\sigma}_{x_i}^{2(k)} + \hat{\lambda}_{x_j}^{(k)} \hat{\sigma}_{x_j}^{2(k)}}{\hat{\lambda}_{x_i}^{(k)} + \hat{\lambda}_{x_j}^{(k)}}$

Remove

Components with associated probability $\hat{\lambda}_p^{(k)} < 10^{-3}$ are deemed irrelevant and removed from the model.

8.3 Decoding With the Estimated Model

8.3.1 Detection of the Start Frame Delimiter

As explained in Section 8.1, the end of the preamble is indicated with the SFD sequence. After synchronization and channel estimation, the receiver starts to look for the SFD. When detected, decoding of the data bits of the payload starts. We saw in Chapter 6 that the SFD detection can be interpreted as a decoding problem where the receiver looks at N_{sfd} consecutive received preamble symbols and tries to determine whether they correspond to the squared SFD sequence $\mathbf{s}^{\mathbf{2}(\text{sfd})} = (s_0^{2(\text{sfd})}, \ldots, s_{N_{\text{sfd}}-1}^{2(\text{sfd})})$.

In Chapter 6 we also presented a suite of algorithms that are able to detect the SFD in an online or offline fashion. All of these algorithms rely on the assumption that the samples at the receiver output y_m are independently distributed according to a non-central or central chi-square distribution if there is no MUI. The type of the distribution depends on whether the signal of the UOI is present or absent due to the ternary preamble and SFD codes. We further showed, that each of the algorithms only depends on calculation of the quantity

$$\text{LLR}(y_m|N_0/2, 2BT, q_m) \quad = \quad \ln\left[\frac{f_{NC\chi^2}\left(\frac{y_m}{N_0/2}\middle|2BT, \frac{q_m}{N_0/2}\right)}{f_{\chi^2}\left(\frac{y_m}{N_0/2}\middle|2BT\right)}\right] \tag{8.16}$$

representing the log-likelihood ratio of a sample rather being drawn from the non-central rather than from the central chi-square distribution. The energy-delay profile

$$q_m = \int_{mT}^{(m+1)T} [\tilde{h}(t)]^2 \,\mathrm{d}t, \quad m \in \{0, 1 \ldots, M-1\} \tag{8.17}$$

can be estimated during the preamble, which we have shown in Chapter 5.

Under the mixture model presented in this chapter, the samples at the receiver output are still assumed independent. Further, they are distributed according to a mixture of scaled non-central chi-square distributions if the signal of the UOI is present, or according to a mixture of Gamma distributions if it is not (see Section 8.2.2 and Section 8.2.3). We can therefore use the exact same algorithms we saw in Chapter 6 to detect the SFD also with the mixture model. The

only change we have to make is to replace the LLR given in (8.16) with

$$\begin{aligned} \mathrm{LLR}_{\mathrm{Mix}}(y_m|\boldsymbol{\Theta}, q_m) &= \ln\left(\frac{f_{NC\chi^2\mathrm{Mix}}(y_m|\boldsymbol{\Theta}, q_m)}{f_{\Gamma\mathrm{Mix}}(y_m|\boldsymbol{\Theta})}\right) \\ &= \ln\left(\sum_{p=0}^{P-1} {}_0F_1\Big(;BT;\frac{q_m y_m}{4\sigma_p^4}\Big) e^{-\frac{q_m}{2\sigma_p^2}} \gamma_p(m)\right) \end{aligned} \tag{8.18}$$

where $\gamma_p(m)$ is given by (8.12) and the confluent hypergeometric limit function is given by

$${}_0F_1\Big(;BT;\frac{q_m y_m}{4\sigma_p^4}\Big) = \quad 2^{BT-1}\Gamma(BT)\Big(\frac{\sigma_p^2}{\sqrt{y_m q_m}}\Big)^{BT-1} I_{BT-1}\big(\frac{\sqrt{q_m y_m}}{\sigma_p^2}\big) \tag{8.19}$$

8.3.2 Demodulation of Payload Data Bits

For demodulation of the data bits, the exact same considerations as for detection of the SFD apply. This time we can reuse the optimal decision rule from equation (5.5) in Chapter 5. Again, we just need to replace the LLR with its mixture counterpart given by (8.18). With $N_\mathrm{f} = T_\mathrm{f}/T$, this yields the following decision rule for the data symbols

$$\sum_{m=0}^{M-1} \mathrm{LLR}_{\mathrm{Mix}}(y_{m+iN_\mathrm{f}}|\boldsymbol{\Theta}, q_m) \underset{\hat{d}_0=1}{\overset{\hat{d}_0=0}{\gtrless}} \sum_{m=0}^{M-1} \mathrm{LLR}_{\mathrm{Mix}}(y_{m+iN_\mathrm{f}+N_\mathrm{f}/2}|\boldsymbol{\Theta}, q_m) \tag{8.20}$$

8.4 Performance Evaluation

For the performance evaluation, we simulate a complete packet-based IEEE 802.15.4a system according to the assumptions in Section 2.5. Different packet arrival rates R were considered for the maximum allowable size of 1208 coded bits per packet.

We simulated both the mandatory low pulse repetition frequency mode (LPRF) and the mandatory high pulse repetition frequency mode (HPRF). In all simulations, the UOI uses preamble code 5 and the interferers use code 6.

Further, the receiver we simulated operates with an integration time of $T = T_\mathrm{c} = 2$ ns. With this choice $BT = 1$ and the interference model to be estimated reduces to a mixture of exponentials. Our receiver estimates the interference model during both the preamble and the payload, during periods where the signal of the UOI is known to be absent. The forgetting factor that yielded best results was determined through simulations and found to be $\kappa = 0.9$. The size of a block $\tilde{M}$ for the recursive estimation was set to $\tilde{M} = 384$, which corresponds to

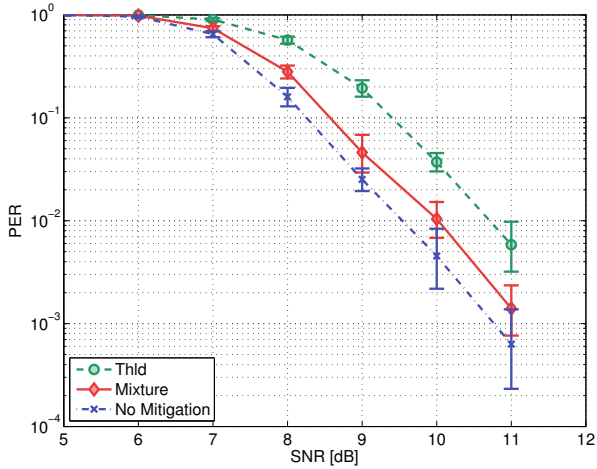

Figure 8.1: Single user scenario. The thresholding receiver shows a loss of about 1 dB because useful information is cut by the threshold.

the number of samples per payload symbol that do not have a contribution of the UOI.

The channel models used in our simulations are the IEEE 802.15.4a residential non-line-of-sight (NLOS) and the office line-of-sight (LOS) models [14]. Our performance metrics are the packet error rate (PER) and the synchronization error rate (SER) that are plotted against the signal to noise ratio (SNR), defined as $\mathrm{SNR} = \frac{E_p}{N_0}$ where E_p is the received energy *per preamble pulse*.

We compare the mixture receiver presented in this chapter ("Mixture") with the optimal receiver for the single user case ("No Mitigation") derived in Chapter 5 and with a robust receiver that uses an adaptive threshold to reject interference ("Thld") also derived in Chapter 5.

Figure 8.1 shows the performance of the three receivers in the single user case with no interference. The slight performance difference between the mixture receiver and the optimal one is due to the robust but slightly suboptimal channel estimation from Chapter 5 used by the former. The thresholding receiver shows a loss of about 1 dB which is due to useful information being cut by the threshold.

Figures 8.2 and 8.3 show near-far scenarios where we have one interferer that is received with a power 10 dB higher than the UOI. The interferer sends at $R = 100$ packets/s which is roughly half of the peak rate. In Figure 8.2, we show results for LPRF and the NLOS

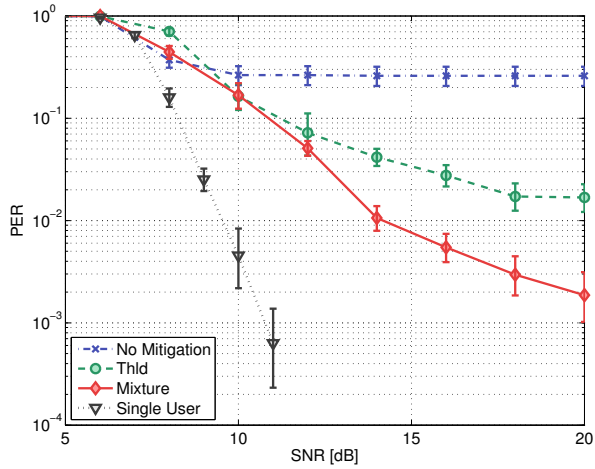

Figure 8.2: Near-far scenario with one interferer received with a power 10 dB higher than the UOI. Results shown are for LPRF with a NLOS channel. The proposed receiver shows a substantially better packet error rate (PER) than receivers that do not mitigate interference at all or use an adaptive threshold.

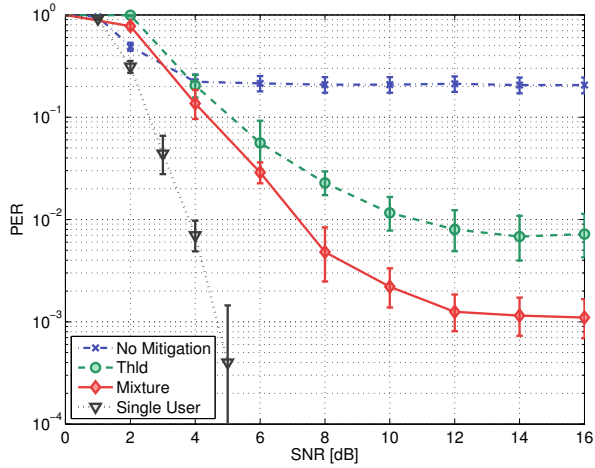

Figure 8.3: Same scenario as in Figure 8.2, but this time for HPRF and the LOS channel. Again, the receiver employing a mixture model shows a substantially better packet error rate (PER) than receivers that do not mitigate interference at all or use an adaptive threshold.

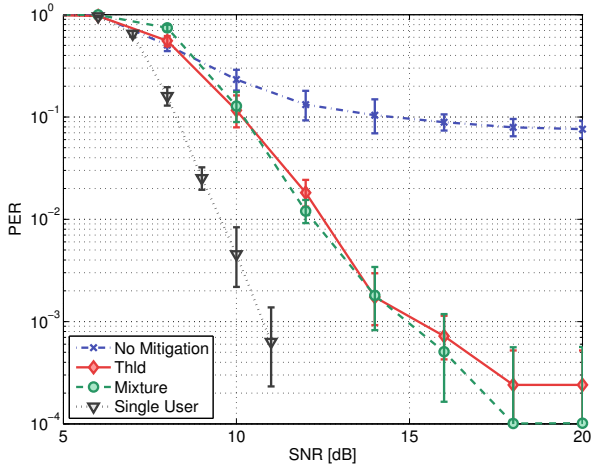

Figure 8.4: Scenario with three interferers that have a received power equal to the UOI. Again, the mixture as well as the thresholding receiver outperform receivers that do not mitigate interference substantially.

channel model, in Figure 8.3 for HPRF and the LOS channel model. In both scenarios, the optimal single user receiver that makes a Gaussian approximation is severely affected by MUI and shows PERs in the order of $20 - 25\%$. Using the mixture model achieves a very good robustness against MUI, yielding a reduction in PER of up to two orders of magnitude. We can also see that in both cases the mixture model outperforms the thresholding receiver, resulting in substantially lower PER floors.

Figure 8.4 shows a scenario where the interferers and the UOI have the same received power. We simulated three interferer at a rate of $R = 200$ packets/s. Again, both the mixture and the thresholding receiver improve substantially over a receiver that does not mitigate interference. The mixture receiver and the thresholding receiver have almost equal performance. Using an interference model no longer gives an advantage, the thresholding receiver already achieves a very good mitigation. Further, the lower interference terms (with respect to the near-far case) are more difficult to identify for the EM algorithm. It can therefore happen that the EM algorithm finds a local minimum that does not characterize the interference sufficiently well. This in turn may result in a mitigation scheme that is closer to the Gaussian approximation, which has a bad performance as can be seen from the “No Mitigation” curve.

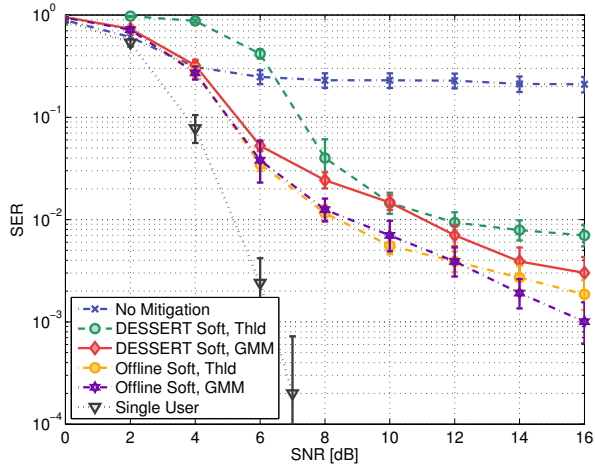

Figure 8.5: Synchronization error rate (SER) versus SNR in a near-far scenario with three interferers. For the DESSERT algorithm, a 2 dB gain is achieved at low SNR if the GMM is used instead of a thresholding scheme. For the offline algorithm, no difference in performance is observed between the GMM and a threshold based implementation.

Finally, we also show the effect of the mixture model on the SFD detection performance. Figure 8.5 shows results for the soft DESSERT and soft offline algorithms that were both introduced in Chapter 6. The results shown are for LPRF and four users with a rate of $R = 100$ packets/s. The three near interferers have 10 dB higher power than the UOI. For the offline algorithm we see only a negligible improvement when the GMM is used, it almost coincides with the offline algorithm that uses only a threshold. For the online DESSERT algorithm, using the mixture model in lieu of the thresholding scheme proposed in Chapter 6 yields a 2 dB improvement at low SNR where the system is noise limited and the threshold therefore cuts useful information as we have seen from the single user results. At high SNR, where the limiting factor is MUI, only a small performance gain is observed for the mixture model and a threshold is sufficient to limit the impact of MUI on SFD detection.

8.5 Conclusion and Possible Directions for Future Work

We presented an energy-detection receiver that mitigates interference from concurrent transmissions by employing statistical interference modeling. We showed that this approach gives a huge performance advantage compared to receivers that do not mitigate interference. Compared to receivers that use an adaptive threshold, the performance gain is smaller but still considerable in near-far scenarios. Another advantage over thresholding schemes is the fact that the mixture model includes the special case where no MUI is present, thus resulting in a smaller performance degradation in the single user case.

Evidently, the performance increase also comes at a certain cost. We see the biggest problem in this respect in the evaluation of the log-likelihood ratio needed for the optimal decision rule in both SFD detection and data decoding. Although the form of the expression lends itself for tabulation, at higher model orders several evaluations are necessary, thus adding to the complexity. On the other hand, the optimal expressions may prove valuable for future work since they allow to get insight into how a suboptimal non-linearity with good interference mitigation characteristics should look like. Another possible direction for future work is an evaluation of a complete network to determine how the performance gain in packet error rate translates into improvement of the overall network throughput.

Part II

Effect of Drifting Clocks on IR-UWB Energy-Detection Receivers

Chapter 9

Clock-Offset Tracking Software Algorithms For IR-UWB Energy-Detection Receivers

The implementation of IEEE 802.15.4a devices faces several challenges. First, the mandatory medium access control layer (MAC) in the standard is based on Aloha and hence, completely uncoordinated; receivers must be robust to occasional interference. In Part I of this thesis we presented various mechanisms that enhance the robustness to such interference of low-complexity receivers based on the energy-detection architecture. Further, with devices mainly operated indoors, receivers should be able to adapt to a time-varying environment and, in particular, a time-varying propagation channel. Finally, with a strong focus on low-priced devices, the underlying hardware can be of average quality. For instance, because of possibly low-quality frequency oscillators used for clock generation, the standard allows for relative clock offsets as large as 40 parts per million (ppm) [35].

In the case of energy-detection receivers, these challenges lead to a trade-off between robustness to the environment and resilience to clock drifts. Extremely low complexity energy-detection receivers are built with a large and constant integration duration, on the order of several tens of nanoseconds [109]; they are quite robust to clock drifts but are sensitive to noise enhancement effects [31] and cannot adapt to channel variations. More sophisticated energy-detection receivers attempt to estimate the power delay profile of the propagation channel. They use a shorter integration duration and combine several weighted outputs of the energy collector according to the estimate of the power delay profile. Examples of such receivers are

[31, 33] and the receiver we used throughout Part I of this thesis. These receivers are robust to noise enhancement effects, can adapt to channel variations and offer a better performance than non-adaptive receivers. Further, they allow for receiver designs that are robust to multi-user interference (see Part I). However, because of the shorter integration duration and consequently higher sampling frequency, they become sensitive to clock offsets.

In packet based systems, such as IEEE 802.15.4a networks, there is no global synchronization. For each received packet, the reception of the payload is preceded by a packet detection and timing acquisition phase: Its purpose is (1) to detect the presence of the packet on the wireless medium, (2) to synchronize the clocks of the transmitter and the receiver and (3) to find out when the payload starts. Now, and especially with cheap hardware, the clocks at the transmitter and the receiver drift. For instance with clock offsets of -4 ppm at the transmitter and 18 ppm at the receiver, the overall clock offset is 22 ppm (clock drifts are measured with respect to a global reference clock.). With a clock frequency of 500 MHz, the clock will be offset by one sample every 90 μs. As we show in Section 9.4 of this chapter, this can severely degrade the performance of even energy-detection receivers that are in general believed to be robust against timing impairments.

Unfortunately, none of the classical solutions to deal with clock offsets that are known from narrowband physical layers are directly applicable to energy-detection receivers, as we already pointed out in Section 3.2.2 when discussing related work. Further, there does not seem to exist any related work studying clock offsets in the context of energy-detection receivers and their impact on this particular receiver architecture is thus not well understood.

Our main contribution in this chapter is the development and performance evaluation for IR-UWB energy-detection receivers of (1) a clock-offset compensation solution and (2) a clock-offset tracking algorithm. By tracking, we imply *both* the estimation and correction of the clock offset. Our tracking algorithm is constructed around the Radon transform [178, 179], an image processing tool traditionally used to detect line features in images. Our solution does not increase the hardware complexity of the receiver and naturally takes advantage of the multipath propagation channel for the estimation of the clock offset. Further, our algorithms are directly applicable to the IEEE 802.15.4a standard and the tracking algorithm reduces the performance loss due to clock offset to less than 0.5 dB.

This chapter is organized as follows: An updated system model, incorporating the effect of clock offsets on the received signal, is given in Section 9.1. We describe our clock offset compensation and tracking algorithms in Section 9.2 and Section 9.3, respectively. The performance of both algorithms is evaluated in Section 9.4. We conclude the chapter in Section 9.5.

9.1 System Model and Assumptions

Without loss of generality, we consider an IEEE 802.15.4a IR-UWB physical layer (introduced in Section 2.4. However, everything that follows equally applies to other IR-UWB PHYs, such as the classical one described in Section 2.2. Further, we focus on non-coherent, energy-detection reception with binary pulse position modulation (BPPM) (see Section 2.3.3).

The received signal model is a slightly modified version of the one we already encountered in the foregoing chapters and that was introduced in Section 2.4 of Chapter 2. A modification is needed to incorporate the effect of imperfect clocks. The new received signal model for the preamble is then the following

$$r_{\text{pre}}(t) = \sum_{i=0}^{N_{\text{pre}}-1} s_i \sum_{j=0}^{C-1} c_j \cdot \tilde{h}(t-(1+\epsilon)(j+iC)L_{\text{s}}T_{\text{c}} - \nu_0) + n(t) \tag{9.1}$$

where ϵ denotes the unknown relative clock offset between the transmitter and the receiver. We assume both ϵ and the compound channel impulse response $\tilde{h}(t)$ to be invariant for the duration of one packet. Further, $n(t)$ is assumed to be a zero-mean AWGN process with power spectral density $N_0/2$ (i.e., we here assume no interference). In this chapter, the pulse repetition period $L_{\text{s}}T_{\text{c}}$ plays an important role as we shall see later. However, already here we simplify the notation by introducing $k \doteq j+iC$. This allows us to combine the N_{pre} preamble symbol modulation coefficients s_i and the ternary preamble code c_j of length C into a compound preamble symbol $\tilde{c}_k = s_{\lfloor \frac{k}{C} \rfloor} \cdot c_{k \mod C}$. With this notation, the received signal is now

$$r_{\text{pre}}(t) = \sum_{k=0}^{N_{\text{pre}}C-1} \tilde{c}_k \cdot \tilde{h}(t-(1+\epsilon)kL_{\text{s}}T_{\text{c}} - \nu_0) + n(t) \tag{9.2}$$

The receiver still uses the preamble for packet detection, timing acquisition, channel estimation and SFD detection. In this chapter we will see that it additionally serves to perform clock-offset estimation and tracking.

During the payload, the received signal model including the effect of clock drifts is

$$r_{\text{pay}}(t) = \sum_{i=0}^{N_{pay}-1} \sum_{j=0}^{N_{\text{cpb}}-1} b_{ij} \cdot \tilde{h}(t-(1+\epsilon)T_{i,j} - \nu_0) + n(t) \tag{9.3}$$

with $T_{i,j} = iT_{\text{f}} + c_{\text{THS},i}T_{\text{burst}} + d_iT_{\text{f}}/2 + jT_{\text{c}}$ and where N_{pay} is the number of symbols in the

payload, T_f is the duration of a symbol, $d_i \in \{0,1\}$ is the i-th symbol of the payload, $c_{\mathrm{THS},i}$ denotes the time-hopping sequence and $b_{ij} \in \pm 1$ is the scrambling sequence.

9.1.1 Receiver Model and Receiver Operations

On the receiver side, the signal is filtered with a bandpass filter of bandwidth B, squared and integrated by the energy-detection receiver. The output of the integrator is sampled at rate $1/T$. The integration time T is chosen such that M samples are taken per pulse repetition period of the preamble i.e., $T = \frac{L}{M} T_\mathrm{c}$.

This results in the following discrete signal

$$y_{m,n} = \int_{mT+\psi(n)}^{(m+1)T+\psi(n)} [r_\mathrm{pre}(t)]^2 \, \mathrm{d}t, \quad m = 0, 1 \ldots, M-1. \tag{9.4}$$

With $\psi(n) = nL_\mathrm{s}T_\mathrm{c}$, $y_{m,n}$ denotes the m-th sample of the n-th block of $L_\mathrm{s}T_\mathrm{c}$ consecutive preamble samples.

Our receiver employs the baseline packet detection and timing acquisition algorithm described in Chapter 4. Upon detection of a signal, the receiver undergoes a verification phase. It is also at this point that we start the clock-offset tracking algorithm proposed in this chapter. If verification is successful, fine synchronization is performed using the jump-back-and-search-forward algorithm that is also explained in Chapter 4. The receiver then performs a period of channel estimation where it estimates the energy-delay profile of the channel according to what we saw in Chapter 5. At the same time it also begins to look for a special signal sequence called start-frame-delimiter (SFD). The SFD is used to designate the end of the preamble and the beginning of the payload and detected according to the soft online algorithm introduced in Chapter 6.

For the demodulation of the n-th data bit d_n of the payload, the receiver may use the optimum decision rule from Chapter 5

$$\sum_{m=0}^{M-1} y_{m,n} \cdot q_m \underset{d_n=1}{\overset{d_n=0}{\gtrless}} \sum_{m=0}^{M-1} y_{m+\frac{N_\mathrm{f}}{2},n} \cdot q_m \tag{9.5}$$

where the weighting coefficients q_m are derived from the energy-delay profile of the channel. Alternatively, a traditional energy-detection receiver with an integration window of fixed dura-

tion $T_{\text{Fix}} = \frac{L}{M_{\text{Fix}}} T_{\text{c}}$ can also be used. In this case we have

$$q_m = 1 \text{ if } m \leq M_{\text{Fix}}, \text{ and } 0 \text{ otherwise.} \tag{9.6}$$

During data reception, the samples $y_{m,n}$ are given by (9.4) where $r_{\text{pre}}(t)$ is replaced with $r_{\text{pay}}(t)$ and $\psi(n)$ now equals $\psi(n) = nT_{\text{f}} + c_{\text{THS},n} T_{\text{burst}} + \hat{\nu}_0$.

9.2 Window Expansion: Clock Drift Compensation

A very simple and natural way of addressing clock drift is to gradually expand the length of the integration window of traditional energy-detection receivers at a fixed rate ϵ_r. The integration window is increased by one sample every $1/\epsilon_r$ samples. Hence, as the signal drifts, the major part of its energy stays within the window. Assuming we know the precision of the oscillators used in our system, ϵ_r is typically chosen to be roughly the expected clock drift. A clock-drift estimation is then not required. We call this method *Window Expansion*. The drawback of this method is noise enhancement due to the increasing integration duration.

Window Expansion can be generalized to receivers applying a weighting function, such as the one in (9.5): To expand the window, we smooth the weighting function employed in the decision rule by convolving it with a time-varying window function that increases over time at rate ϵ_r.

Increasing the integration time of traditional energy-detection receivers is then equivalent to convolving the rectangular weighting function given by (9.6) with a rectangular window of increasing length. For the receiver employing the optimum decision rule in (9.5), we found a convolution with a triangular window of increasing length to yield better results than a convolution with a rectangular one.

Note that we do not need any criterion to stop the expansion of the window. The reason is that within the time it takes to receive an IEEE 802.15.4a packet, the window cannot grow beyond a symbol duration, even at the growth rate corresponding to the maximum expected clock drift.

9.3 Radon Tracking: A Clock Offset Tracking Algorithm Based on the Radon Transform

In this section, we present a clock-offset tracking algorithm, called *Radon Tracking*, which allows for more sophisticated clock-offset compensation than Window Expansion.

9.3.1 Equivalence Between Slope Estimation of a Line in a Gray-Scale Image and Clock Drift Estimation

During the preamble, the samples $y_{m,n}$ can be rearranged in an "energy matrix", $\mathbf{Y} = [y_{m,n}]_{M \times N}$ [83, 84]. The n-th column contains the M consecutive samples corresponding to the n-th pulse of the preamble. This energy matrix is then equivalent to a gray-scale image where $y_{m,n}$ is the intensity of the pixel at coordinate (m, n). Our clock-offset tracking algorithm relies on the following observation: the estimation of the clock drift is equivalent to finding the slope of parallel lines in this gray-scale image. With perfect clock synchronization ($\epsilon = 0$), the signal in (9.2) is $L_s T_c$-periodic. Consequently, the discrete signal given by (9.4) is M-periodic. Accordingly, the energy matrix displays a pattern resembling parallel horizontal lines. These parallel lines correspond to the multi-path components of the signal. Figure 9.1 (a) shows a discrete signal with 40 samples per frame ($M = 40$) and two multi-path components. Figure 9.1 (b) shows the gray-scale image of the corresponding energy matrix. We can observe the two parallel lines corresponding to the two multi-path components.

Now, let's assume for a sample $y_{m,n}$ that the integration window is aligned with the received signal such that it captures the entire energy of a pulse. If the clocks of the transmitter and the receiver exhibit a relative positive (respectively negative) clock drift of ϵ, the alignment of the integration window for the sample $y_{m,n+1}$ is no longer perfect: Some of the energy "leaks" into the sample $y_{m-1,n+1}$ ($y_{m+1,n+1}$, respectively). In consequence, the energy matrix displays now a pattern resembling parallel lines at a given angle ϕ. There is a one-to-one relationship between ϕ and the clock drift ϵ:

$$\phi = \arctan(M\epsilon) \tag{9.7}$$

Figure 9.1 (d) shows our example signal with the two multi-path components but this time subject to a clock offset [1] of $\epsilon = 1e-3$. Figure 9.1 (e) shows the corresponding energy matrix. The signal drifts by one samples every 25 columns ($= 1000$ samples at $M = 40$), leading to

1. This value is illustrative only; relative clock offsets found in oscillators are usually two orders of magnitude lower.

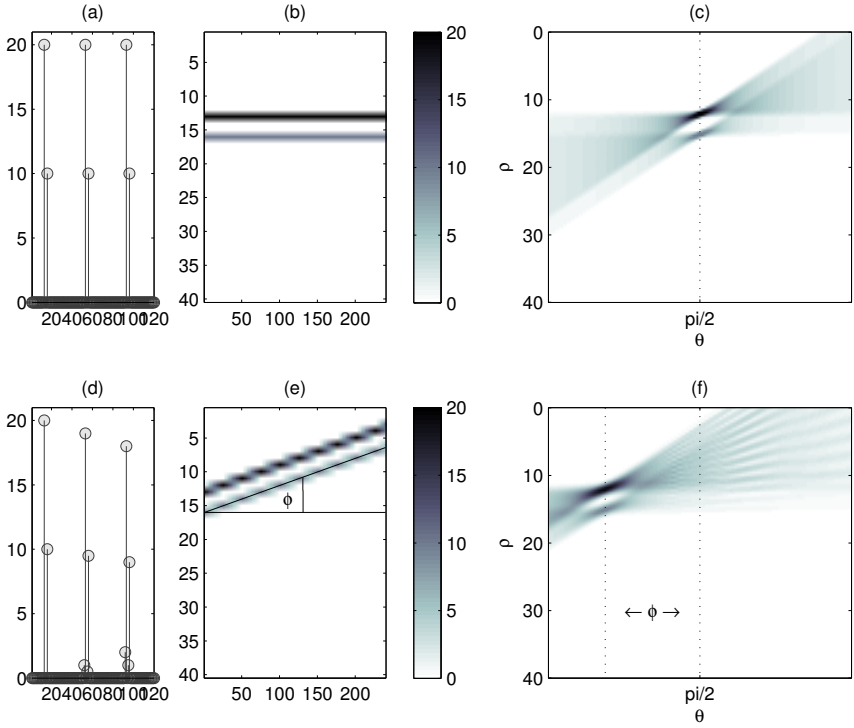

Figure 9.1: A relative clock offset ϵ leaves a distinctive pattern resembling parallel lines at an angle $\phi = \arctan(M\epsilon)$ in the energy matrix of an $M-$periodic signal. In Radon space the maximum is off the right angle by ϕ.

$\phi = \arctan(1/25) = \arctan(40 \cdot 1e-3)$.

9.3.2 The Radon Transform: A Tool for Line Detection

The (two-dimensional) *Radon transform* [2] is widely used in image processing for line feature detection. We use the common ρ, θ parametrization [180] where the Radon transform $R(\rho,\theta)[I(x,y)]$ of the two-dimensional image $I(x,y)$ is

$$R(\rho,\theta)[I(x,y)] = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} I(x,y)\delta(\rho - x\cos\theta - y\sin\theta)\,\mathrm{d}x\,\mathrm{d}y \tag{9.8}$$

where ρ is the distance from a line to the origin and θ is the angle of the vector from the origin to the closest point on the line. We refer to the (ρ,θ)-parameter space as *Radon space*.

Every point (ρ,θ) in the Radon space corresponds to the integral along the line $y = -\frac{\cos\theta}{\sin\theta}x + \frac{\rho}{\sin\theta}$ in the original image $I(x,y)$. Finding lines in a gray-scale image corresponds to finding points with high intensities in Radon space. The basic idea of our algorithm is to apply this principle to our problem of clock-drift estimation. This is illustrated in Figures 9.1 (c) and 9.1 (f) that show the Radon transforms of our example signals.

Note that the Radon transform can be computed iteratively (see the algorithm in Section 9.3.4). Hence, we do not need to accumulate the complete energy matrix. In our simulations, we calculate the Radon transform by blocks of M samples. In principle, it can even be calculated sample by sample.

9.3.3 Pre-Processing to Denoise the Input

The Radon transform is already quite robust to noise. However, we further increase this robustness with a pre-processing on the energy matrix before the calculation of the Radon transform. First, we take into account the preamble code by only considering samples $y_{m,n}$ where the corresponding code symbol $\hat{c}_n \neq 0$. Then, along the rows of $\mathbf{Y}$, we apply a moving average filter over the last G pulses, yielding a matrix $\tilde{\mathbf{Y}}$ with elements $\tilde{y}_{m,n}$. Finally, the elements of $\tilde{\mathbf{Y}}$ below a threshold η_{radon} are set to zero [3]. This yields the matrix $\bar{\mathbf{Y}}$ with elements $\bar{y}_{m,n}$. Because the noise approximatively follows a chi-square distribution, we calculate a threshold η_{radon} to

2. For discrete binary input images it is often referred to as *Hough transform*.
3. This also speeds up the algorithm as zero valued entries of $\bar{\mathbf{Y}}$ do not have to be processed in the subsequent steps.

reject samples with a probability of more than 5% to consist of noise only:

$$\eta_{\text{radon}} = \frac{N_0}{2} F^{-1}_{\chi^2}(1 - 0.05|2BT \cdot G) \tag{9.9}$$

where $N_0/2$ is the (estimated) noise power spectral density and $F^{-1}_{\chi^2}(x|2BT \cdot G)$ is the inverse of the cumulative distribution function of the chi-square distribution with $2BT \cdot G$ degrees of freedom.

9.3.4 Computation of the Discrete Radon Transform

The Radon transform in (9.8) is defined for a continuous input. Therefore, we transform the discrete matrix $\bar{\mathbf{Y}} = [\bar{y}_{m,n}]_{M \times N}$ into a continuous "image" $I(x, y)$ via nearest-neighbor interpolation, i.e., $I(x, y) = \bar{y}_{\lfloor y \rceil, \lfloor x \rceil}$, where $\lfloor \cdot \rceil$ denotes the nearest-integer function. Further, as we cannot store the continuous output of the Radon transform, we discretize the Radon space as follows

$$\rho_i = i \cdot \Delta\rho, \quad \theta_j = j \cdot \Delta\theta. \tag{9.10}$$

where $\Delta\rho$ is defined with respect to the size of a pixel in the energy matrix i.e., $\Delta\rho = 1/8$ means that we have 8 discrete values per pixel. Hence, we calculate a discrete version $\mathbf{R} = [R_{i,j}]$ of the Radon transform according to

$$R_{i,j} = \int_{\rho_i}^{\rho_{i+1}} R(\rho, \theta_j)[I(x, y)] \, \mathrm{d}\rho. \tag{9.11}$$

Because the function $I(x, y)$ is piecewise constant, (9.11) is actually simple to compute. Algorithm 2 shows an outline of the algorithm we use, an illustration is given in Figure 9.2. Furthermore, as the precision of the oscillators is known, the range of interest of both θ and ρ is known. Therefore, the Radon transform can be calculated for only the points of interest, limiting both processing and memory requirements. The Radon matrix that we store has a constant size independent of the length of the observation. The Radon transform is akin to a compression scheme: Instead of the energy matrix, it yields an alternative matrix of a smaller dimension, which still captures all the signal information that is necessary to perform clock-drift estimation.

Algorithm 2: Calculates Radon Transform as in (9.11)

Input: Pixel $\bar{y}_{m,n}$ of $\bar{\mathbf{Y}}$
Output: For each ρ, θ an updated entry $R_{\rho,\theta}$ of discrete Radon transform.
if $\bar{y}_{m,n} \neq 0$ **then**
  foreach θ **do**
    foreach ρ **do**
      $f \leftarrow$ fraction of pixel with center $y_{m,n}$ lying between lines parametrized by (ρ, θ) and $(\rho + \Delta\rho, \theta)$;
      $R_{\rho,\theta} \leftarrow R_{\rho,\theta} + f \cdot y_{m,n}$;
    end
  end
end

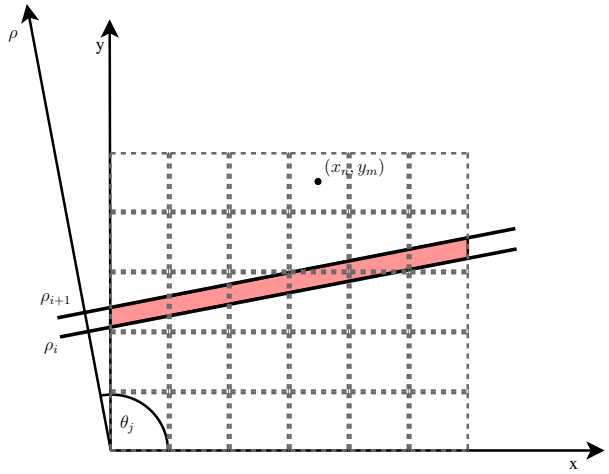

Figure 9.2: Illustration showing how to calculate Radon Transform. The total intensity of the shaded area will be attributed to the entry (ρ_i, θ_j) of the Radon transform matrix.

9.3.5 Angle Estimation by Detection of Maxima in Radon Space

As explained in Sections 9.3.1 and 9.3.2, an estimation of the clock drift is equivalent to finding the estimate $\hat{\phi}$ of the angle ϕ. Thus, we have to look for maxima in the Radon matrix $\mathbf{R}$. Naturally, we should find them around the true angle ϕ. Due to multipath, there will be more than one such maximum at different values of ρ. We want to take advantage of this property by combining the contributions of several multipath components. However, the column sums of the Radon transform matrix are all equal to a constant value i.e. the integral over the whole image $I(x, y)$. Therefore, the naive approach of choosing $\hat{\phi}$ to be the θ_j with the highest intensity averaged over all ρ does not work. Instead, we use the following method to determine $\hat{\phi}$: First, we smooth the matrix $\mathbf{R}$ by convolving each column with a rectangular window of length $W = 2/\Delta\rho$, thus combining the values corresponding to two pixels in the original gray-scale image. Then, we square each entry of the smoothed matrix and compute the column sums. Finally, we set $\hat{\phi}$ to the θ_j corresponding to the column with maximum column sum, i.e.

$$\hat{\phi} = \Delta\theta \cdot \arg\max_j \sum_k ((rect_W * R_{\cdot,j})[k])^2 \tag{9.12}$$

where $rect_W$ is the rectangular window and $R_{\cdot,j}$ denotes the j-th column of $\mathbf{R}$.

9.3.6 Continuous Tracking of the Transmitter Clock

We are not only interested in the estimation of the clock drift. We also want to continuously compensate for it in order to stay aligned with the packet. Tracking of the transmitter clock is commonly done by adjusting the frequency of the oscillator at the receiver. However, new observations are obtained with an updated sampling frequency after the adjustment of the frequency of the oscillator. This implies a change of the pattern in the energy matrix and a modified Radon transform. A priori, this makes a block by block operation necessary where after each update of the receiver clock: (1) the Radon transform is discarded, (2) a large block of new samples is collected in order to obtain a new Radon transform and (3) the current clock drift is re-estimated on the new Radon matrix.

However, such a costly approach can be avoided. It is possible to maintain a single Radon matrix by applying a coordinate transform to the Radon space before an update of the receiver clock. This allows for the conservation of the entire signal history and for a continuous estimation and correction of the clock drift.

An illustration of the tracking process and the involved coordinate transform is given in

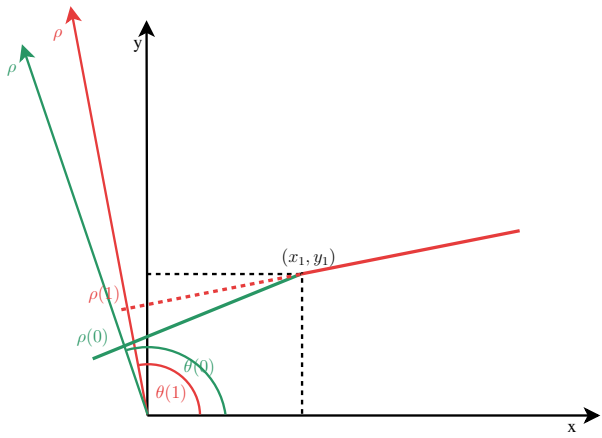

Figure 9.3: Illustration of the coordinate transform needed to cope with changing clock drifts that are due to adjustments of the oscillator frequency. Here the drift changes from $\theta(0)$ to $\theta(1)$ at time index x_1.

Figure 9.3 and derived in the following. The example shows the simplified case of a single line, but extension to the situation in Figure 9.2 with a small band is straightforward. In Figure 9.3 we have a line at an angle $\theta(0)$ that gets projected to $\rho(0)$ according to equation (9.8). At column (or "time") index x_1 we adjust the clock of the receiver resulting in a change of the angle by Δ_1 and consequently in a line at an angle

$$\theta(1) = \theta(0) + \Delta_1 \tag{9.13}$$

For the point (x_1, y_1) we get the following set of equations according to (9.8):

$$\rho(0) = x_1 \cos\theta(0) + y_1 \sin\theta(0) \tag{9.14}$$

$$\rho(1) = x_1 \cos\theta(1) + y_1 \sin\theta(1) \tag{9.15}$$

From (9.13)-(9.15) we get after some basic algebra

$$\rho(1) = \frac{\rho(0)\sin(\theta(0) + \Delta_1) - x_1 \sin\Delta_1}{\sin\theta(0)} \tag{9.16}$$

Generalizing this to the coordinate transform after k clock frequency adjustments, we find

$$\theta_j(k) = \theta_j(0) + \sum_{l=1}^{k} \Delta_l \tag{9.17}$$

$$\rho_i(k) = \frac{\rho_i(0)\sin(\theta_j(k)) - \sum_{l=1}^{k} x_l \sin\Delta_l}{\sin\theta_j(0)}$$

$$\approx \rho_i(0) - \sum_{l=1}^{k} x_l \sin\Delta_l \tag{9.18}$$

where Δ_l is the angle adjustment corresponding to the l-th clock frequency adjustment, $\rho_i(k)$ and $\theta_j(k)$ are the transformed coordinates after k adjustments, and x_l is the column index of the matrix $\bar{\mathbf{Y}}$ corresponding to the l-th point in time where the adjustment occurs. The approximation is possible because for practical clock drifts we have that $\theta_j(m) \approx \pi/2$.

Equations (9.17) and (9.18) define the transformation we have to apply on the coordinate system. To be able to do this transformation we just have to keep the two values $\sum_{l=1}^{k} \Delta_l$ and $\sum_{l=1}^{k} x_l \sin\Delta_l$ up to date.

In our simulations we start filling the Radon matrix once a coarse estimate of the signal arrival time is available. During the verification phase and fine synchronization, we do not enable tracking but attempt to get a better first estimate of the clock offset. After fine synchronization, we track the transmitter clock until the SFD is found.

9.3.7 Handling the Residual Clock Drift

The signaling format change between the preamble and the payload in IEEE 802.15.4a makes it extremely difficult to maintain a consistent Radon matrix. Hence, we perform clock-drift estimation and tracking only during the preamble of a packet. However, the preamble is long enough such that the residual drift is small. Nevertheless, we find a performance increases if we compensate for this residual drift by employing Window Expansion over the payload (Section 9.2). Here again, the growth rate of the window is small enough, such that we do not need a criterion to stop the window expansion (see also Section 9.2).

9.4 Performance Evaluation

To evaluate the effect of clock drift on energy-detection receivers and the performance of our algorithms, we simulate a full IEEE 802.15.4a system with coarse and fine synchronization,

estimation of the energy-delay profile of the channel, SFD detection, and data decoding with the $(63, 55)$ Reed-Solomon code.

The main assumptions for the performance evaluation are again those given in Section 2.5 of Chapter 2. In particular, we focus on the mandatory LPRF mode and use the residential NLOS channel model. Our main performance metric is the packet error rate (PER) and we simulate the maximum allowable packet length of 1016 bits per packet. We use the default length of the preamble of 72 code symbols (including 8 code symbols for the SFD). The signal-to-noise ratio (SNR) is defined as $\mathrm{SNR} = \frac{E_p}{N_0}$ where E_p is the received energy *per pulse* (after the convolution of the pulse with the impulse response of the channel). The physical layer is simulated with an accuracy of 100 ps, corresponding to a *simulation* sampling frequency of 10 GHz. We assume oscillators with a drift uniformly distributed in the range of ± 20 ppm, resulting in relative clock offsets ϵ between transmitters and receivers of up to ± 40 ppm. For the receiver, we mainly focus on the optimal receiver (ED_{OPT}) detailed in Chapter 5 with $T = T_{\text{c}} = 2$ ns resulting in a 500 MHz sampling frequency. The bandpass filter is adapted accordingly. For comparison purposes, we also simulate two reference receivers: one with a fixed integration time $T_{\text{Fix}} = 128$ ns (ED_{FIX}), and one with an integration time that was optimized for our scenario (ED_{VAR}). Without drift, this optimized integration time was found to be 56 ns, which is roughly the channel spread of the residential NLOS model. Further, ED_{VAR} may use Window Expansion to increase the integration time at a constant rate ϵ_r (but ED_{FIX} may not). The duration of the integration time for ED_{FIX} is motivated by the fact that with the IEEE 802.15.4a LPRF mode the minimum inter-pulse spacing during the preamble corresponds to 128 ns as well. Note, that for both ED_{VAR} and ED_{FIX}, we do not simulate the preamble. Instead, we assume that a perfect synchronization puts the integration window such that it captures a maximal amount of energy.

9.4.1 Trade-off Between Noise Enhancement and Robustness to Clock-Drift

Figure 9.4 shows the performance degradation due to a relative clock offset up to ± 40 ppm when no clock-offset compensation is used. We can clearly see the trade-off between noise enhancement and robustness to clock drift. With perfect clocks, ED_{OPT} outperforms the other receivers, even though the curves for ED_{FIX} and ED_{VAR} were obtained with a perfect synchronization. On the other hand, ED_{OPT} is very sensitive to clock drifts: Indeed they cause misalignments of the weights p_m (see (9.5)) that strongly degrade the performance. Interestingly, ED_{VAR} is also severely affected by clock drift, although to a lesser extent than ED_{OPT}.

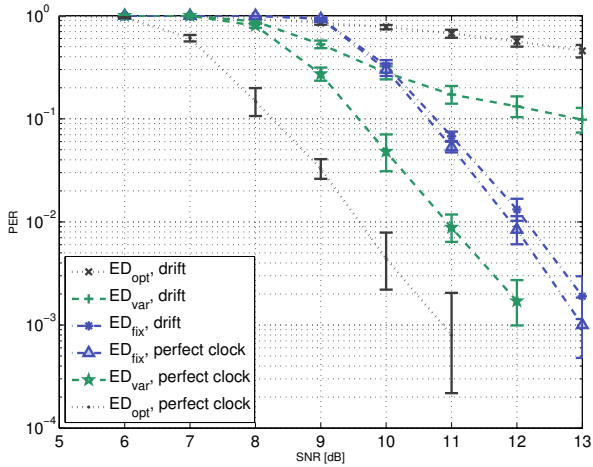

Figure 9.4: PER with perfectly synchronized clocks or with clocks differing by up to ± 40ppm. No receiver compensates for drift. ED_{OPT} is very sensitive to clock drift but shows optimal performance with perfect clocks. For ED_{FIX} the opposite is true. ED_{VAR} is also severely affected by clock drift showing the necessity to compensate clock offsets even in suboptimal receivers. ED_{FIX} and ED_{VAR} suffer from noise enhancement.

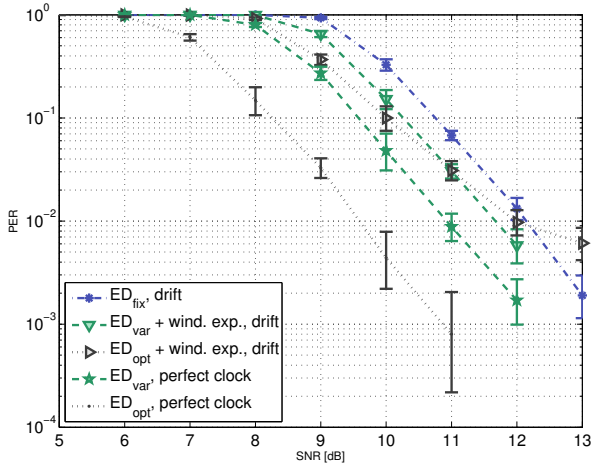

Figure 9.5: ED_{VAR} and ED_{OPT} show much better robustness to clock drift if simple Window Expansion is used. However, with drift the performance of ED_{VAR} is close to ED_{FIX} and $2dB$ worse than the optimum receiver with perfectly synchronized clocks. Even with Window Expansion, ED_{OPT} shows an error floor. More sophisticated mechanisms to cope with drift are necessary.

ED_{FIX}, on the other hand, is barely affected due to its long integration time.

9.4.2 Window Expansion Helps to Some Extent

Figure 9.5 shows the improvement achievable through Window Expansion. We found $\epsilon_r = 32$ ppm to give best performance for ED_{VAR} and $\epsilon_r = 40$ ppm for ED_{OPT}. With drifting clocks, ED_{VAR} with Window Expansion performs hardly better than ED_{FIX}, which amounts to a loss of about 2 dB with respect to the optimum curve for perfectly synchronized clocks. ED_{OPT} with Window Expansion, on the other hand, crosses the other curves and still shows an error floor. Its performance is thus still limited by packets that are lost due to clock offsets. Although Window Expansion gives some improvement, it is definitely not sufficient, specially in the case of ED_{OPT}.

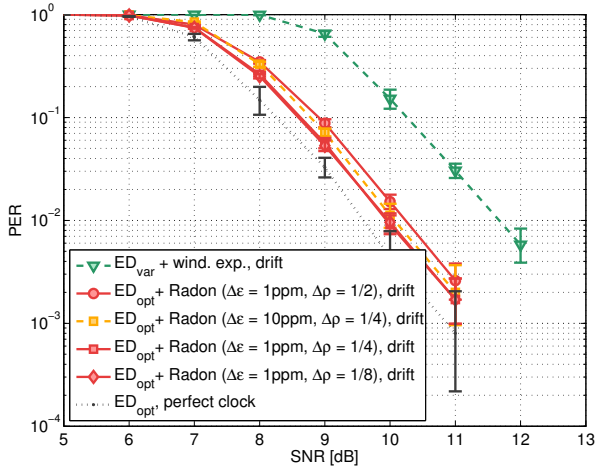

Figure 9.6: PER if Radon Tracking is used. Curves are for ED_{OPT} with different discretizations of $\Delta\epsilon$ and $\Delta\rho$. We also show the two curves from Figure 9.5 giving best performance with and without drifting clocks. Our algorithm performs well, with $\Delta\rho$ equal to $1/8$ and $1/4$. We loose less than 0.5 dB with respect to the optimal curve.

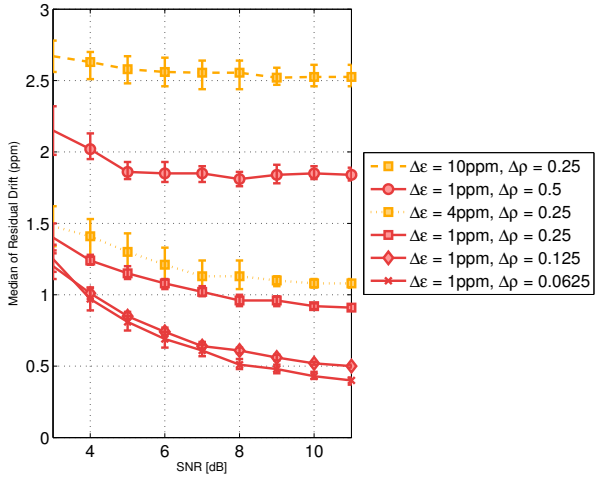

Figure 9.7: (a) Median and 95% confidence intervals of the (absolute) residual drift of Radon Tracking at the end of the preamble. Finer discretization leads to a better clock drift estimate. For all resolutions $\Delta\epsilon$ of the drift estimate, the algorithm is able to approach the optimum of $\frac{\Delta\epsilon}{4}$ at high SNR and provided that the resolution of ρ, $\Delta\rho$, is chosen small enough.

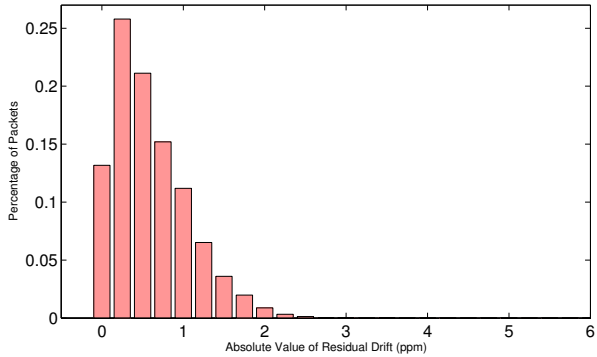

Figure 9.8: Distribution of absolute residual drifts of packets after the preamble with Radon Tracking. Distribution shown is for $\Delta\epsilon = 1$ ppm, $\Delta\rho = 1/4$ and SNR $= 11$ dB.

9.4.3 Radon Tracking Achieves Near Optimal Performance

Figure 9.6 shows the performance of Radon Tracking. We show results for different values of the discretization step $\Delta\rho$. The angular resolution $\Delta\theta$ either corresponds to a resolution of the drift of $\Delta\epsilon = 1$ ppm or a resolution of $\Delta\epsilon = 10$ ppm. To increase robustness against noise we combine $G = 32$ preamble pulses and to account for residual drift we set $\epsilon_r = 2$ ppm. We also show reference curves for ED_{OPT} with perfect clocks, and ED_{VAR} with Window Expansion and drifting clocks. The results show that our algorithm performs extremely well with a loss of less than 0.5 dB with $\Delta\rho = 1/8$ with respect to the optimal curve for perfect clocks. For $\Delta\rho = 1/4$ the performance is virtually identical. If we further increase $\Delta\rho$ we get an additional loss of about 0.2 dB at $\Delta\rho = 1/2$. Even for a resolution of the drift estimator of $\Delta\epsilon = 10$ ppm, we get excellent results at $\Delta\rho = 1/4$. In general, a finer discretization yields a better estimate of the clock drift. This is confirmed by Figure 9.7, where we show the median of the absolute value of the residual drift at the end of the preamble. For comparison, we calculated the mean absolute estimation error of a perfect algorithm that always finds the best estimate but also has the same constraint on the resolution $\Delta\epsilon$. The derivation is given in Appendix A.1 and proves that no algorithm with resolution $\Delta\epsilon$ can do better than $\frac{\Delta\epsilon}{4}$. We see that for all resolutions our algorithm quickly approaches this bound as the SNR increases, provided that the resolution of ρ is sufficient. This indicates that our algorithm has a performance that is close to the optimal algorithm.

Figure 9.8 shows the corresponding distribution of the absolute values of the residual clock drift for $\Delta\epsilon = 1$ ppm and $\Delta\rho = 1/4$ at 11 dB. Only a few packets have a residual drift of more than 2 ppm. This justifies our choice for $\epsilon_r = 2$ ppm. We found this to be similar for other angular resolutions. E.g., for $\Delta\epsilon = 10$ ppm and $\Delta\rho = 1/4$ only few packets have a residual drift above 4 ppm. In the case of $\Delta\epsilon = 10$ ppm we therefore set $\epsilon_r = 4$ ppm.

9.4.4 Effect on Synchronization

For both ranging and communication, synchronization is the most crucial part for the reception of a packet. Figure 9.9 shows the effect of clock drift on the synchronization error rate (SER) for ED_{OPT}. The SER is measured after SFD detection and includes false alarms and missed detections. If the receiver does not use tracking, we can see an error floor at about 5% packets lost due to synchronization errors, mainly because the estimated channel energy-delay profile used in SFD detection is now misaligned due to drift. With Radon Tracking (here shown for $\Delta\rho = 1/4$, results for other values of $\Delta\rho$ were practically identical), the error floor disappears

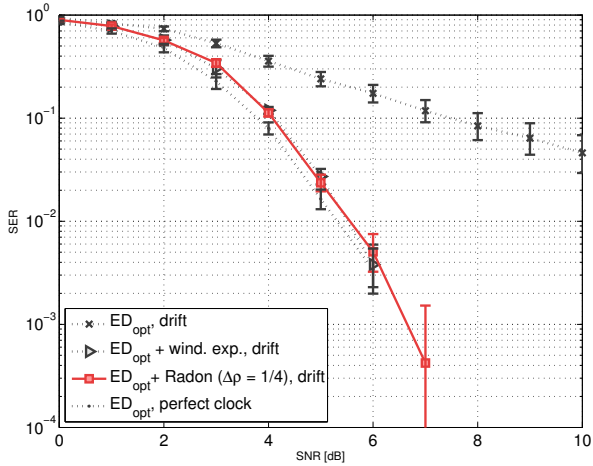

Figure 9.9: Synchronization error rate accounting for false alarms and missed detections for ED_{OPT}. With clock-offset tracking enabled, the error floor disappears and performance is within 0.3 dB of the one with perfect clocks. With only Window Expansion we interestingly have similar performance. In contrast to data decoding, clock-offset tracking is thus not strictly required for synchronization alone.

and the performance is less than 0.3 dB worse than with perfectly synchronized clocks. Interestingly, the same is true for a receiver that only employs Window Expansion. We attribute this mostly to the fact that we continuously update the estimate of the channel energy-delay profile during the preamble. Large misalignments of the weights q_m are therefore not possible during the preamble and the small ones are recovered through Window Expansion. In contrast to data decoding, clock-offset tracking is thus not strictly required for synchronization alone. Even in ranging applications, however, this is only of limited interest because one might need to determine the clock offset in order to account for it in the range calculations.

9.5 Conclusion and Possible Directions for Future Work

We have shown that to prevent a considerable performance loss, clock drifts need to be addressed for energy-detection receivers. We analyzed two solutions: Window Expansion and Radon Tracking. The latter is a clock-offset tracking algorithm based on the Radon transform that yields excellent performance. Further, it can be implemented without requiring any changes to the hardware structure of an energy-detection receiver. The output of the algorithm can also be directly used for estimation of the clock drift between transmitter and receiver, which may be needed to improve the accuracy of ranging measurements.

We believe that this algorithm is even more versatile and could be used for channel estimation and time-of-arrival estimation purposes. An intuition of this claim can be gained by observing that the Radon matrix (see Figure 9.1) retains all the necessary information to perform these tasks: the column of the matrix corresponding to the current clock offset estimate corresponds exactly to the channel energy-delay profile. The Radon Tracking algorithm could thus provide a compact method to perform several receiver operations in parallel. Further investigations on this are required and could form a possible future research direction.

Another point that future work should address is the performance of the algorithm under multi-user interference. However, there is no reason to believe that techniques similar to those that we developed in Part I of this thesis could not be applied to the samples at the input of this algorithm.

Part III

Security of IEEE 802.15.4a Ranging

Chapter 10

Effectiveness of Distance-decreasing Attacks Against Impulse Radio Ranging

Its large bandwidth and the resulting fine temporal resolution, give IR-UWB the unmatched capability of high (sub-meter) precision indoor ranging, even in dense multi-path environments [181]. This ability to accurately measure the distance between two devices is crucial for many location aware applications and services, and was thus one of the main motivations to include IR-UWB as an alternative PHY in IEEE 802.15.4.

On the other hand, many applications that rely on information that depends on a device's physical location, are security-sensitive and require ranging to be executed in a secure manner. By secure we mean that some trusted location information needs to be obtained, even in the presence of an adversary that interferes with the ranging process. One example of such a security-sensitive application is the tracking of goods and people [182]. Consider a valuable item, such as a Swiss watch, equipped with a wireless-enabled (RFID) tag. The watch is displayed in a store equipped with a monitoring system that measures the distance to the tag every few seconds. If anyone tries to move the watch beyond some distance from the monitoring system, an alarm is triggered. The system must prevent an adversary from decreasing the measured distance between the monitoring system and the tag. Otherwise, the adversary can remove the watch from the store and make the monitoring system believe that the watch is still within the store premises. Other examples for applications with similar requirements include RFID access control [183], secure neighbor discovery [184], secure time synchronization [185], and secure localization [152].

Significant research efforts have led to various *secure ranging protocols* (see [147] for the

original description and Section 3.3 of the related work for references to alternate proposals) that guarantee the measured distance between two devices to be an *upper-bound* on their actual distance. [1] With secure ranging protocols, an adversary is prevented from decreasing the estimated distance by injecting bogus messages because the messages exchanged during the ranging process are cryptographically protected.

While secure ranging protocols prevent distance-decreasing attacks on the protocol level, they abstract away from the lower communication layers. Previous work [162] points out that the physical layer (PHY) leaves space for attacks against secure ranging. Such attacks on the PHY make it possible for an adversary to decrease the estimated distance *without* breaking any cryptographic primitives or protocols. However, such attacks are, by nature, PHY-specific, a separate evaluation of their effectiveness on different PHY technologies is necessary.

In this chapter we evaluate the effectiveness of distance-decreasing attacks against the physical layer of IR-UWB. Understanding the impact of PHY attacks on IR-UWB is of particular interest because IR-UWB is a natural candidate for the PHY of a secure ranging protocol: IR-UWB not only has the potential for high precision ranging, but it is to date also the only wireless technology with a standard specifically designed for ranging applications. Further, the IEEE 802.15.4a standard even includes an optional *private ranging mode* meant to enable ranging in the presence of an adversary.

We show that the *de facto* standard for IR-UWB, IEEE 802.15.4a, does not automatically provide security against such attacks. We find that with the mandatory modes of the standard an external attacker can decrease the measured distance by more than 100 meters with a high probability (above 99%). Further, we evaluate possible countermeasures and we also analyze the impact of different receiver structures on the attack performance.

This chapter is organized as follows: In Section 10.1 we give our assumptions on honest and adversarial devices as well as on the secure ranging protocol. In Section 10.2 we propose a set of attacks that compromise the security of secure ranging with IEEE 802.15.4a. These attacks are introduced under the assumption that both honest and adversarial devices are energy-detection receivers. In Section 10.3 we evaluate the effectiveness of the proposed attacks with detailed physical layer simulations. In Section 10.4 we discuss possible countermeasures and explain why the private ranging mode of the standard is not resilient against the proposed attack. In Section 10.5 we relax the assumption on the receiver architecture and reevaluate the effectiveness of distance-decreasing attacks in different scenarios where Rake receivers may be used in addition to energy-detection receivers. Finally, we conclude the chapter in Section 10.6.

1. For this reason, such protocols are often also referred to as *distance-bounding protocols.*

10.1 System Model and Assumptions

10.1.1 Assumptions on Honest Wireless Devices

We assume that devices engaging in a secure ranging protocol share the necessary cryptographic material, and that they are equipped with an IEEE 802.15.4a compliant receiver and transmitter. A description of IEEE 802.15.4a can be found in Section 2.4 of Chapter 2 and the reader is referred there for details. Characteristics of IEEE 802.15.4a that are essential to understand the physical layer attacks will, however, also be highlighted when we introduce the attacks. As usual, we restrict ourselves to the mandatory modes of the standard. Nonetheless, we still discuss some of the optional features, such as the *private ranging mode*, which is discussed in Section 10.4.1.

The choice of transmitter is of little consequence to our investigation, any standard-compliant transmitter is acceptable [124, 186].

In the main part of this chapter, the architecture of the receiver used by honest devices is assumed to be the non-coherent energy-detection receiver from Chapter 5. In a second part, we relax this assumption and also consider the implications of using a coherent Rake receiver. Still, we deem the energy-detection receiver to be a more realistic solution for RFID in terms of cost and complexity and easier to implement on the small and cheap active tags that we consider. It has also been shown to have optimal performance in this class of receivers (see, e.g., Chapter 5 or [33]) and it can be made quite robust to multi-user interference (see Part I of this thesis).

We assume that the receiver performs the usual operations, such as packet detection and timing acquisition, channel estimation, SFD detection and data decoding. Throughout this chapter, we assume no multi-user interference (MUI) and the receiver considered here uses algorithms that, as we have seen in Part I, are not necessarily robust to MUI. In particular, this means that packet detection and timing acquisition are done according to the baseline algorithm in Section 6.2 and channel estimation according to Section 5.1.2. For SFD detection, the receiver employs the DESSERT algorithm using soft-decision decoding, described in Section 6.3, and for data decoding the received samples are weighted with the optimal weights for burst transmissions according to Section 5.1.1. However, all of the considered attacks equally apply to the robust algorithms and we have no reason to believe that any of our findings would fundamentally change in this case.

10.1.2 Assumptions on Secure Ranging Protocol

In general, secure ranging protocols are developed without a particular physical layer in mind. They rather act on the protocol level by enhancing traditional ranging with messages that are cryptographically protected.

A secure ranging protocol is executed between two wireless devices, that are traditionally called the *verifier* and the *prover*. The verifier initiates the ranging process by sending a range request to the prover. The prover then responds with another message and the verifier estimates their mutual distance based on the round trip time. The verifier then also verifies that the obtained result is valid, i.e., that no one tampered with the result.

We distinguish two classes of secure ranging protocols: protocols that attempt to thwart internal attacks and protocols that only try to prevent external attacks. An internal attack occurs if one of the devices that execute the protocol is misbehaving. Protocols that only prevent external attacks, assume both the verifier and the prover to be honest.

We only consider external attacks, for several reasons. First of all protocols preventing internal attacks include, for the most part [2], a *rapid-bit-exchange* phase (RBE): The verifier sends a number of single bit challenges, to which the prover must respond instantly. Such unusual requirements make these protocols difficult to implement. [3] In particular, an IEEE 802.15.4a implementation of RBE would not only be extremely inefficient, as every bit would have to be prefixed with a (relatively long) preamble – it would also be open to packet-level attacks considered in [162]. Furthermore, security against external attacks is sufficient in many applications (e.g., the theft prevention system mentioned earlier). Finally, considering only the PHY, an external attack is more challenging to mount than an internal one. [4] Moreover, the individual components of the attack that we will devise in the following can also be used by a malicious prover to mount an internal attack.

On the other hand, protocols that merely try to prevent external attacks [152, 153] only require the exchange of a small number (typically 2) of ranging messages that are several bits long. This makes them easily implementable on IEEE 802.15.4a compliant devices. Figure 10.1 shows a typical example of such a protocol. Both ranging messages consist of nonces that are unpredictable by an external adversary. The nonces are bound to the shared, secret key via a message authentication code, which allows the verifier to check the validity of the

2. A recent proposal [156] shows that the RBE can be replaced by a full-duplex transmission in protocols that provide security against internal attacks

3. To this date, no implementation of RBE for wireless networks exists.

4. The same is not necessarily true in a more global context, since one could argue that it might be difficult for an adversary to compromise a honest device, which is required to mount an internal attack.

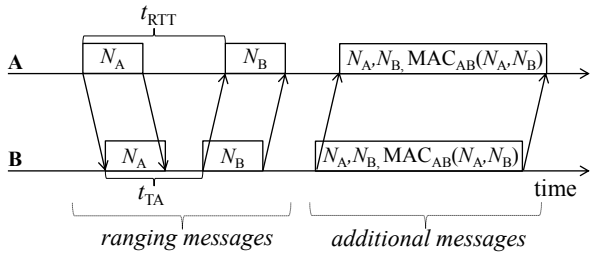

Figure 10.1: Example secure ranging protocol: Device **A** (verifier) estimates the distance to device **B** (prover) with the formula $d_{AB} = c(t_{\mathrm{RTT}} - t_{\mathrm{TA}})/2$, where c is the channel propagation speed. $\mathrm{MAC}_{\mathrm{AB}}$ stands for Message Authentication Code with a symmetric key shared between **A** and **B**, N_A and N_B are freshly generated nonces, t_{TA} is a constant turn-around time that **A** and **B** know, and t_{RTT} is the round-trip-time measured by **A**.

protocol run and prevents an external adversary from injecting bogus messages and thus from interfering with the ranging process.

Two Strategies to Check Validity of a Nonce at the Verifier

We have already said that nonces need to be unpredictable by the adversary. This is usually ensured by a constraint on their length, which makes correct guessing of the nonce hard. If this were not the case, we would open the door for guessing attacks that make it possible for an external adversary to decrease the measured distance. An example of a guessing attack, where the adversary correctly guesses the nonce N_B, is shown in Figure 10.2.

To prevent an adversary from guessing a nonce, we fix a security level P_{guess}, corresponding to the maximum allowable success probability of a guessing attack. If a nonce of length N_{nonce} has equiprobable bits, we have that

$$P_{\text{guess}} = 2^{-N_{\text{nonce}}} \tag{10.1}$$

For a given security level, a first strategy (nonce verification strategy I) consists in fixing the minimum required length of a nonce according to (10.1). The verifier then accepts a nonce as legitimate if it contains no errors.

An alternative strategy (nonce verification strategy II) is to generalize the above decision

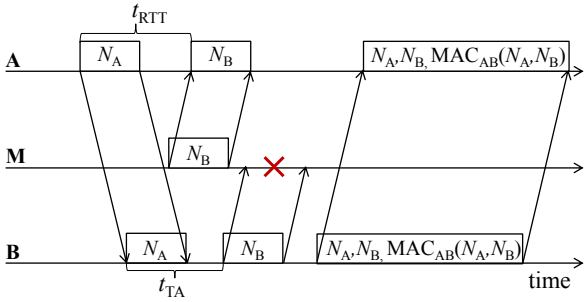

Figure 10.2: Example of a guessing attack: The external adversary **M** guesses the value of the nonce N_B and starts its transmission early enough, such that the measured distance $d_{AB} = c(t_{\text{RTT}} - t_{\text{TA}})/2$ at **A** is decreased. At the same time, **M** ensures that the attack is not detected at **A** by blocking the legitimate response of **B** through, e.g., jamming the signal.

rule at the verifier and to reject a nonce if it contains more than N_{err} bits that are in error, with

$$N_{\text{err}} = F^{-1}_{\text{BIN}}(P_{\text{guess}}|N_{\text{nonce}}, 1/2) \tag{10.2}$$

where $F^{-1}_{\text{BIN}}(x|n,p)$ is the inverse of the CDF of a binomial distribution with parameters n and p.

Although this strategy increases the length of a nonce for $N_{\text{err}} > 0$, it allows the protocol to operate at virtually any bit error rate by simply increasing the length of the nonce (we will see in Section 10.4 why this is an important property for potential countermeasures). Indeed, we have that

$$\lim_{N_{\text{nonce}}\to\infty} N_{\text{err}} = \lim_{N_{\text{nonce}}\to\infty} F^{-1}_{\text{BIN}}(P_{\text{guess}}|N_{\text{nonce}}, 1/2) = \lim_{N_{\text{nonce}}\to\infty} F^{-1}_{\mathcal{N}}(P_{\text{guess}}|N_{\text{nonce}}/2, N_{\text{nonce}}/4) \tag{10.3}$$

where $F^{-1}_{\mathcal{N}}(x|\mu,\sigma^2)$ is the inverse of the CDF of a normal distribution with mean μ and variance σ^2 and the second equality follows from the central limit theorem. The protocol succeeds as long as there are no more than N_{err} errors in a nonce of length N_{nonce}, resulting in a bit error rate of $\text{BER}^{\text{max}} = N_{\text{err}}/N_{\text{nonce}}$. From (10.3) it then follows that, as the length of the nonce increases, the maximum sustainable bit error rate BER^{max} tends to the worst case of $1/2$, i.e.,

$$\lim_{N_{\text{nonce}}\to\infty} \text{BER}^{\text{max}} = \lim_{N_{\text{nonce}}\to\infty} \frac{N_{\text{nonce}}/2 + \sqrt{N_{\text{nonce}}/4}\Phi^{-1}(P_{\text{guess}})}{N_{\text{nonce}}} = \frac{1}{2} \tag{10.4}$$

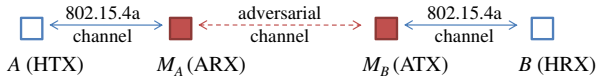

Figure 10.3: Distance-decreasing relay attack setup.

where $\Phi^{-1}(x)$ denotes the inverse CDF of a standard normal distribution.

The applicability of either of these two strategies depends on assumptions made about how the secure ranging protocol is implemented on the honest IEEE 802.15.4a compliant devices.

Consider the likely scenario where the implementor of a secure ranging protocol buys off-the-shelf IEEE 802.15.4a devices and implements the protocol at the application layer. In many cases, it is unlikely that he has access to the lower communication layers and especially to the received bits *before* error correction with the mandatory Reed-Solomon (RS) code is performed. This in turn makes it impossible for him to use nonce verification strategy II because once the RS code is applied, it is no longer possible to determine how many of the received bits were in error. Having said that, he can still use nonce verification strategy I with N_{nonce} equal to the number of information bits at the application layer. This is possible, since the maximum number of errors that the RS code can mask corresponds to the number of parity bits added during RS encoding.

Consequently, nonce verification strategy II can only be applied if information about the number of erroneous systematic bits is available at the protocol level. In this case we assume that the length of the nonce is chosen such that both ranging and communication packets experience similar packet error rates (PER). The length of a nonce is then the solution to the system of equations formed by (10.2), that ensures a certain security level, and

$$N_{\text{err}} = F^{-1}_{\text{BIN}}(1 - \text{PER}_{\text{comm}} | N_{\text{nonce}}, \text{BER}_{\text{comm}}) \tag{10.5}$$

that ensures a PER for ranging PER_{rang} that is in the order of the PER for communication packets PER_{comm}. Equation (10.5) is found by measuring the average BER of communication messages BER_{comm} that leads to the target PER of PER_{comm}. Further, it is assumed that the number of bit errors in a packet is binomially distributed. This is an approximation as individual bits are not identically distributed because of the IEEE 802.15.4a burst transmissions.

10.1.3 Threat Model, Assumptions on Adversary

We consider an external adversary mounting a *distance-decreasing relay attack* between two honest devices, A and B, that execute a secure ranging protocol. The general setup of such an attack is shown in Figure 10.3, the details of the attack are explained in Section 10.2. The adversary acts as a relay and uses two devices M_A and M_B, where A can communicate directly only with M_A and B can communicate directly only with M_B. In some scenarios this is inherent due to the limited radio range of A and B. In other scenarios, this can be achieved by shielding one of the victim devices (e.g., with a "booster bag" that is coated with aluminium foil [187] and acts as a Faraday cage).

Additionally, M_A and M_B exchange information using an out-of-band adversarial channel. The IEEE 802.15.4a channel propagation speed is c, the speed of light. The same speed is assumed for the adversarial channel.

We assume that the adversary is not able to break any cryptographic primitives used by the secure ranging protocol that is executed between A and B. We further assume that the length of nonces is such that guessing attacks are infeasible (see Section 10.1.2).

We focus on the exchange of a single ranging message. In this case one of the honest devices acts as a transmitter (HTX) and the other one as a receiver (HRX). Accordingly, the adversarial devices act as a receiver (ARX), and as a transmitter (ATX). It is easy to extend this attack to an entire secure ranging protocol. The adversary simply mounts the distance-decreasing relay attack on all ranging messages. Any non-ranging messages of the protocol, which are not time critical, can be relayed in an arbitrary fashion.

Assumptions on Adversarial Wireless Devices

The adversarial devices are equipped with transmitters similar to the honest devices, but able to send non-standard-compliant pulse sequences, and to ignore regulatory transmission limits.

Analog to the honest receivers, we for now assume that the receivers of the adversary follow the architecture of the optimal non-coherent energy-detection receiver for IEEE 802.15.4a, introduced in Chapter 5. Later we relax this assumption and consider the case where the adversary may also use a coherent all-Rake (ARake, introduced in Section 2.3.1 of Chapter 2) receiver which gives him the best possibly achievable performance.

The adversary may further equip its devices with high gain antennas, allowing him to increase the SNR observed by both adversarial and honest devices. Such an increase in received SNR can also be achieved by the adversary moving its devices closer to the honest devices.

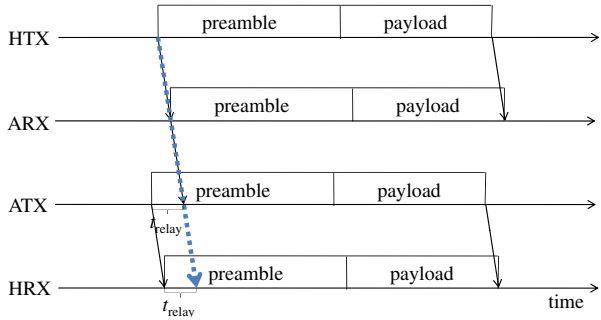

Figure 10.4: Overview of the distance-decreasing relay attack. ARX and ATX are assumed to lie on a line between HRX and HTX. The thick dotted arrow indicates time-of-arrival corresponding to the actual distance between HTX and HRX.

10.2 Distance-decreasing Attack

The main principle of a distance-decreasing relay attack, is for an adversary to relay messages between HTX and HRX in such a way that they seem "shifted back in time" by some positive offset t_{relay} that we call the *relay time-gain*. This is illustrated in Figure 10.4. The measured distance is then decreased by $c \cdot t_{\text{relay}}$.[5]

The difficulty in mounting this attack is twofold: (1) ATX needs to begin the transmission of the preamble before it learns from ARX *when* HTX started the transmission. (2) ATX needs to transmit the payload before it learns *what* bits the payload carries. Existing work has focused exclusively on the second problem [162, 163], but the first one is equally important: Without shifting the time-of-arrival at HRX, attacks on the payload are in vain.

Naturally, the preamble and the payload must be relayed with the same time-gain. This implies that the upper-bound on the achievable time-gain is the minimum of (1) the upper-bound on the time-gain for the preamble and (2) the upper-bound of the time-gain for the payload. As we will see shortly, the payload upper-bound is more strict and determines the achievable relay-gain.

5. Assuming an optimal configuration from the adversary's perspective where ARX and ATX lie on a line between HTX and HRX. In other configurations the distance decrease will be smaller. Note, however, that the choice of the configuration rests with the adversary.

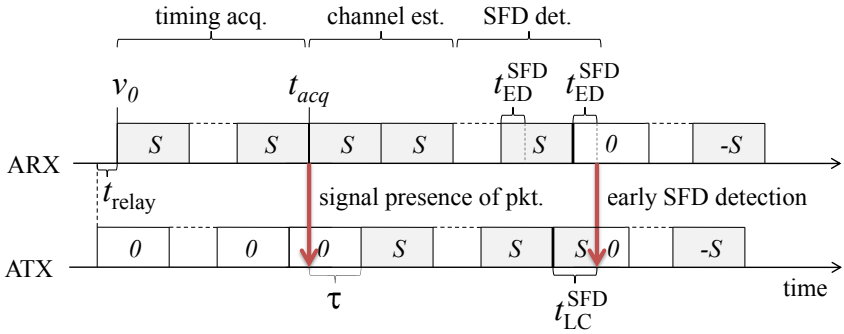

Figure 10.5: Distance-decreasing relay attack on the preamble.

10.2.1 Attack on the Preamble

In general, the adversary cannot know when HTX will begin the transmission of a packet. Although ATX can always choose to begin transmitting a preamble at a random time, there is a good chance it either does this too late, resulting in a distance increase, or too early, such that the distance-decreasing attack on the payload (which, as we already mentioned, will turn out to be the bottleneck) fails. Even if neither is the case, the achieved time-gain is random, which might lead to undesired results, e.g., negative distance estimates.

To circumvent these problems, we propose that the adversary mounts a distance decreasing relay attack on the preamble according to what is depicted in Figure 10.5 and explained in the following. For clarity of presentation, we here assume the distance between ARX and ATX to be 0.

The received signal during the IEEE 802.15.4a preamble is given by (see Section 2.4.2 of Chapter 2)

$$r_{\text{pre}}(t) = \sum_{i=0}^{N_{\text{pre}}-1} s_i \sum_{j=0}^{C-1} c_j \cdot \tilde{h}(t - (j + iC)L_{\text{s}}T_{\text{c}} - \nu_0) + n(t) \tag{10.6}$$

The preamble consists of a succession of N_{pre} preamble symbols of length $T_{\text{psym}} = CL_{\text{s}}T_{\text{c}}$ that are modulated by the sequence s_i. The first N_{sync} preamble symbols form the SYNC part of the preamble and we have that $s_i = 1$, for $i = 0, \ldots, N_{\text{sync}} - 1$. For the remaining N_{sfd} preamble symbols, s_i, $i = N_{\text{sync}}, \ldots, N_{\text{pre}} - 1$, corresponds to the ternary SFD sequence, the mandatory one being $0\ 1\ 0\ {-1}\ 1\ 0\ 0\ {-1}$. In Figure 10.5, we denote preamble symbols for which $s_i = 1$ with S, preamble symbols for which $s_i = -1$ by $-S$ and preamble symbols for which $s_i = 0$

with 0. Every preamble symbol is formed by C code symbols that are modulated according to the ternary preamble code c_j.

Since the preamble codes used in the mandatory IEEE 802.15.4a modes are publicly known, ARX can use the same packet detection and timing acquisition algorithm that is used by honest receivers. Assume that a packet from HTX arrives at ARX at time ν_0 and ARX acquires timing at time t_{acq}. ARX then signals the fact that it acquired timing to ATX. At this point, ARX is synchronized to the boundaries of a preamble symbol. However, the exact time it takes to synchronize depends on the channel conditions and is therefore not known to ARX, thus requiring detection of the SFD as it is also the case for honest receivers.

After timing acquisition, ARX performs channel estimation, again exactly according to the procedures used by honest receivers.

Deviating from honest receivers, ARX then however continues with *early SFD detection*: it chooses an early SFD detection delay $t_{\text{ED}}^{\text{SFD}}$ and tries to detect the presence of the SFD by deliberately only considering the first $M_{\text{ED}}^{\text{SFD}} = t_{\text{ED}}^{\text{SFD}}/T$ samples of every received preamble symbol. T here denotes the sampling period of the receiver. This is in contrast to the honest receiver that will take the entire SFD sequence into account for optimal performance. Early SFD detection can be performed according to the DESSERT algorithm used by honest receivers, with the only difference that every block at its input $\mathbf{y_k}$ now only consists of the first $M_{\text{ED}}^{\text{SFD}}$ samples and is given by

$$\mathbf{y_k} = (y_{kCM}^{\text{pre}}, y_{kCM+1}^{\text{pre}} \ldots, y_{k+M_{\text{ED}}^{\text{SFD}}-1}^{\text{pre}}) \tag{10.7}$$

instead of equation (6.11) of the DESSERT algorithm in Chapter 6. Since preamble symbols during the SYNC part are modulated with a coefficient of 1 and the first preamble symbol of the SFD sequence is modulated with a coefficient of 0, this boils down to testing the two hypotheses "$\mathbf{y_k}$ consists of signal plus noise" versus "$\mathbf{y_k}$ consists of noise only". The SFD is detected if the latter is more likely.

At the other end of the relay, ATX chooses a late SFD commit delay $t_{\text{LC}}^{\text{SFD}}$ and remains silent until ARX signals that timing acquisition was successful. Then, after an appropriately chosen (we explain how shortly) delay $\tau < T_{\text{psym}}$, ATX begins transmitting a sequence of preamble symbols with $s_i = 1$. This is repeated until ARX signals that the SFD was detected. Immediately afterwards, ATX switches to the transmission of a standard compliant SFD, beginning from $t_{\text{LC}}^{\text{SFD}}$ into the SFD. This completes our description of the distance-decreasing attack on the preamble.

In contrast to a standard-compliant preamble, the SYNC part of the preamble generated by

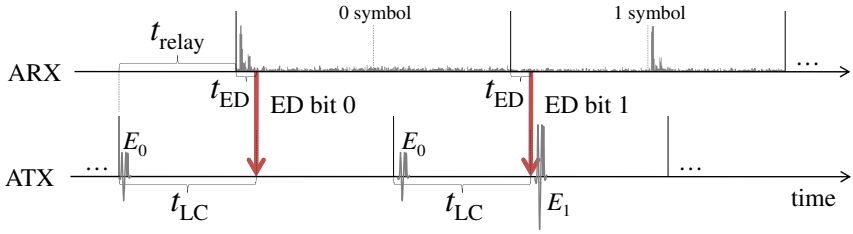

Figure 10.6: Distance decreasing relay attack on the payload symbol.

ATX thus begins with a number of 0 modulated preamble symbols. Further, the beginning of the SFD corresponds to a preamble symbol of the SYNC part for a duration of $t_{\text{LC}}^{\text{SFD}}$, instead of having no signal contribution as it is the case of the standard-compliant SFD. As we show in Section 10.3, the preamble is long enough to accommodate for the missing symbols in the beginning and the modified SFD sequence is close enough to the real one for HRX to synchronize correctly.

The relay time-gain achieved by this attack is $t_{\text{relay}} = t_{\text{LC}}^{\text{SFD}} - t_{\text{ED}}^{\text{SFD}}$. This determines the choice of τ, since $T_{\text{psym}} - \tau = t_{\text{relay}} \mod T_{\text{psym}}$. Note that nothing prevents the adversary from choosing $t_{\text{LC}}^{\text{SFD}} > T_{\text{psym}}$. Rather, the hard upper-bound on the time-gain achievable by the preamble attack is the duration of the SFD T_{sfd}.

10.2.2 Attack on the Payload

The attack on the payload is performed on a frame by frame basis and it again consists of an *early detection* (ED) and a *late commit* (LC) part. An illustration of the attack is shown in Figure 10.6.

The received signal during the i-th IEEE 802.15.4a payload symbol is given by (see Section 2.4.3 of Chapter 2)

$$r_{\text{pay},i}(t) = a_i \sum_{j=0}^{N_{\text{cpb}}-1} b_{ij} \cdot \tilde{h}(t - iT_{\text{f}} - d_i T_{\text{f}}/2 - c_{\text{THS},i} T_{\text{burst}} - jT_{\text{c}} - \nu_0) + n(t) \tag{10.8}$$

Since non-coherent energy-detection receivers are unable to distinguish the BPSK symbols a_i, we can for now assume that $a_i = 1$. Further, the receivers knows the TOA from synchronization allowing us to assume $\nu_0 = 0$. Note that both the scrambling sequence b_{ij}, as well as the time-hopping sequence (THS) $c_{\text{THS},i}$ are publicly known. This leaves us with the BPPM symbol d_i

as the only quantity that is not known to the attacker. For simplicity, we for now assume a time-hopping index of $c_{\text{THS},i} = 0$.

ARX performs an *early detection* attack by deciding on the value of d_i after an *early detection delay* $t_{\text{ED}} < T_{\text{f}}/2$. Instead of taking into account the whole BPPM frame like a honest receiver would, ARX thus deliberately ignores most of the frame. By doing so, ARX essentially replaces BPPM demodulation on-off keying (OOK) demodulation. The maximum likelihood decision rule is then

$$\sum_{m=0}^{M_{\text{ED}}-1} \text{LLR}(y_m | N_0/2, 2BT, q_m) \underset{d_0=1}{\overset{d_0=0}{\gtrless}} 0 \tag{10.9}$$

instead of the optimal BPPM decision rule from Section 5.1.1 of Chapter 5. $M_{\text{ED}} = t_{\text{ED}}/T$ denotes the number of samples that go into the decision and where the log-likelihood ratio (LLR) is given by equation (5.6) of Chapter 5. The time t_{ED} can be made arbitrarily small, it determines the attack's performance. The value for optimal performance is dictated by the channel delay spread, as we show in Section 10.3.

After demodulation, ARX signals the result to ATX. On the other end of the relay, ATX performs a *late commit* attack. ATX begins the transmission of a frame $T_{\text{f}}/2$ before the frame's bit value is received from ARX. At this point in time, ATX does, however, not yet know what bit value the relayed frame should carry. It therefore begins the transmission of the frame, like it would if a 0-bit were transmitted: since it knows both the THS and the spreading sequence it can send an appropriate burst of pulses with energy E_0 in the 0-block of the frame. Once the bit value to be relayed is signalled by ARX, ATX acts accordingly: If it is a 0-bit, it transmits nothing in the 1-block of the frame; if it is a 1-bit, it transmits a burst of pulses with energy $E_1 > E_0$. The resulting transmitted signal has the form

$$x^{LC}_{\text{pay},i}(t) = \sum_{j=0}^{N_{\text{cpb}}-1} b_{ij} \cdot p^{LC}_{d_i}(t - iT_{\text{f}} - c_{\text{THS},i}T_{\text{burst}} - jT_{\text{c}} + t_{\text{relay}}) \tag{10.10}$$

with $p^{LC}_{d_i}(t)$ given by

$$p^{LC}_{d_i}(t) = \sqrt{E_0}p(t) + d_i \cdot \sqrt{E_1}p(t - T_{\text{f}}/2) \tag{10.11}$$

where $p(t)$ is the pulse shape used by the transmitter. The optimal ratio between E_0 and E_1, $\gamma = E_0/E_1$ is determined in Section 10.3. This attack exploits the fact that HRX performs a simple energy comparison to demodulate. The *late commit delay* is $t_{\text{LC}} = T_{\text{f}}/2$. The resulting relay time-gain due to this attack is $t_{\text{relay}} = t_{\text{LC}} - t_{\text{ED}} \leq T_{\text{f}}/2$, which is considerably less than

the upper-bound due to the preamble part of the attack.

Note that the THS does not affect the time-gain of the attack. Indeed, assume that the early detection and late commit delays with time-hopping index $c_{\text{THS},i} = 0$ are denoted by t^0_{ED} and t^0_{LC}, respectively. The relay time-gain under this assumption is given by $t^0_{\text{relay}} = t^0_{\text{LC}} - t^0_{\text{ED}}$. Next, consider mounting these attacks in the case where $c_{\text{THS},i} > 0$. Because $c_{\text{THS},i}$ is publicly known, the adversary simply shifts early detection and late commit in time. Hence, $t_{\text{ED}} = t^0_{\text{ED}} + c_{\text{THS},i} T_{\text{burst}}$ and $t_{\text{LC}} = t^0_{\text{LC}} + c_{\text{THS},i} T_{\text{burst}}$, and consequently $t_{\text{relay}} = t^0_{\text{LC}} + c_{\text{THS},i} T_{\text{burst}} - t^0_{\text{ED}} - c_{\text{THS},i} T_{\text{burst}} = t^0_{\text{LC}} - t^0_{\text{ED}} = t^0_{\text{relay}}$.

10.2.3 Processing Delays

An additional factor that reduces the relay time-gain, and hence the amount by which the distance can be decreased, are the processing delays at ARX and ATX for the IEEE 802.15.4a and adversarial channels. We discuss these delays here, and argue that it is feasible to keep them in the order of $10 - 30$ nanoseconds (or below 10 meters). We focus on the payload, as it is the bottleneck in terms of the achieved delay (the adversary has much more time flexibility during the preamble). We distinguish two cases: (i) ARX and ATX integrated into one device, with appropriate shielding and directional antennas, and (ii) remote ARX and ATX. The latter case can lead to a broader scope of attacks, as the adversary has the flexibility of placing its devices close to the corresponding victim devices. On the downside, remote ARX and ATX are subject to an additional processing delay, due to communication over the adversarial channel.

We first consider the processing delay related to the communication with the honest devices, which applies in both (i) and (ii). At ARX the delay consists of the processing due to demodulation, *after* the necessary signal has been received. Most of the quantities that are needed in the decision rule given by (10.9) can be pre-computed or tabulated such that a fast evaluation is possible. For minimum latency, the adversary further has the option of increasing the integration time of the energy-detection receiver to t_{ED}. In this case, the adversary looses about 1 dB in received power with respect to a receiver that samples at the chip level. However, since $M_{\text{ED}} = 1$, the decision rule can be reduced to a simple check against a pre-computed threshold [105, 106]. Overall, demodulation should be doable in a few clock cycles, leading to processing delays in the order of only a few nanoseconds. At ATX, the delay is of the same order: after the bit value is received from ARX, the transmitter only needs to proceed with or abort the transmission of a previously known burst of pulses (Figure 10.6).

In case (ii), there is an additional delay due to communication over the adversarial channel:

more precisely, the delay of putting the bit value on the adversarial channel at ARX, and demodulating it at ATX. The exact numbers depend heavily on the technology ARX and ATX use to communicate. The adversary is most likely to choose a wireless communication medium, due to its faster propagation speed, but even more so because of the ease of attack deployment compared to a wired channel.

We emphasize that the adversarial channel has unusual requirements. It does not require a high bit-rate, as the adversary only needs to transmit a single bit every 1μs. However, the bit has to be transmitted as fast as possible. Many wireless technologies, even those with very high bit-rates, such as 802.11n, are not suitable: They achieve these high bit-rates through large modulation constellation sizes, rather than a short symbol duration. One valid option is IR-UWB with on-off keying, and a receiver similar to the ED receiver described in Section 10.2.2. Naturally, the adversary will ignore the regulations and transmit with a power high enough to achieve a negligible error rate. To mitigate the multipath delay spread, a highly directive antenna can be used, as proposed for a narrow-band communication system in [188]. The coherent two-level PSK scheme proposed in [188] can also be used as the adversarial channel: It reports bit duration of only 1.6 ns. Overall, in case (ii), a processing delay in the order of $10 - 30$ ns ($3 - 9$ m) seems feasible.

10.3 Performance Evaluation

In this section, we evaluate the effectiveness of the distance-decreasing relay attacks with packet-based system simulations. We simulate a full IEEE 802.15.4a system including all the operations necessary to receive a packet: timing acquisition, estimation of the channel energy-delay profile, SFD detection, and data decoding. The physical layer is simulated with an accuracy of 100 ps.

As explained in Section 10.1.1, we confine ourselves to the two mandatory IEEE 802.15.4a modes (LPRF and HPRF). The standard suggests using the LPRF mode with energy-detection receivers operating in environments with a high multipath delay spread. For energy-detection receivers operating in environments with low delay spread, using the HPRF mode is preferable. Following these suggestions, we therefore use two different channel models to evaluate the LPRF and HPRF modes: The IEEE 802.15.4a residential non-line-of-sight (NLOS) model for LPRF and the the office line-of-sight (LOS) model for HPRF [14].

In all our simulations, we use the ternary preamble code number 5 of length $N_{\mathrm{pcode}} = 31$ given by the standard. The values chosen for $t_{\mathrm{ED}}^{\mathrm{SFD}}$ and $t_{\mathrm{LC}}^{\mathrm{SFD}}$ are chosen with respect to the

structure of this code. Further, the integration time of the energy-detection receiver is set to $T = 2$ ns and all confidence intervals shown are at the 95% level.

Our main performance metrics are the packet error rate (PER) and the synchronization error rate (SER). We consider a packet to be in error if it was not acquired during synchronization or if it does not pass the verification at the verifier. A packet is not acquired if it is not detected (missed detection) or if the synchronization is off by too much for data decoding to be performed correctly (false alarm). We have seen in Section 10.1.2 that depending on the implementation of a secure ranging protocol, different verification strategies that in turn lead to different definitions of the PER are possible. For nonce verification strategy I, based on the RS decoded bits at the application layer, verification fails if more bits are received in error than the RS code can correct. In this case we assume a payload of 128 bits, which we consider a conservative upper-bound on the length of a ranging message in a secure ranging protocol. For nonce verification strategy II, that is based on the systematic bits and does not use the RS code, the number of tolerated bit errors N_{err} depends on the length of the nonce N_{nonce}. Having nonces of $N_{\text{nonce}} = 42$ bits and tolerating up to $N_{\text{err}} = 3$ errors results in a security level of $P_{\text{guess}} = 2^{-32}$. Further, the PERs of ranging packets in this case match those of a communication packet with a payload of 32 bits, reaching a PER of $\text{PER}_{\text{comm}} = 10^{-2}$ at an SNR of about 6.5 dB.

The signal to noise ratio (SNR) is defined as $\text{SNR} = \frac{E_p}{N_0}$ where E_p is the received energy *per pulse* (after the convolution of the pulse with the impulse response of the channel). To evaluate the cost of the attack, we compare the benign case performance (honest receiver and transmitter) with the performance under attack. We then express the cost as the difference in SNR (between the two cases) necessary for the same performance (SER, PER). This tells us by what factor the adversary needs to boost the received signal level to obtain the same performance as in the case of an honest execution of the protocol. He can achieve this by using a high-gain antenna, by transmitting with a higher power, or by moving closer to the victim transceivers.

In a first part of the performance evaluation, we determine the performance of attacks on the preamble and on the payload individually. In a second part, we look at the whole system, putting all the components together, thus allowing us to assess the overall performance of the distance-decreasing relay attack.

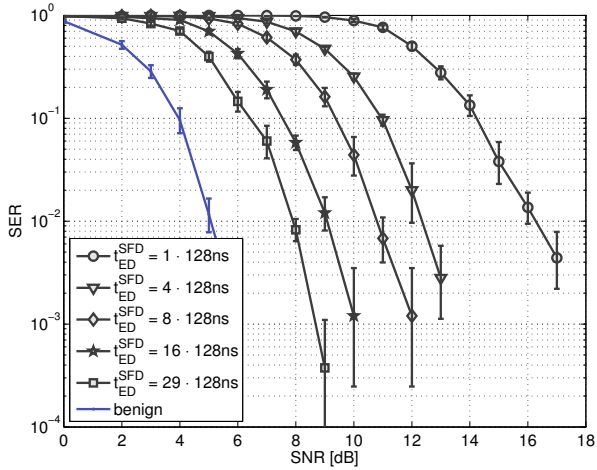

Figure 10.7: SER versus SNR for LPRF comparing benign performance to ED with varying ED delays $t_{\mathrm{ED}}^{\mathrm{SFD}}$.

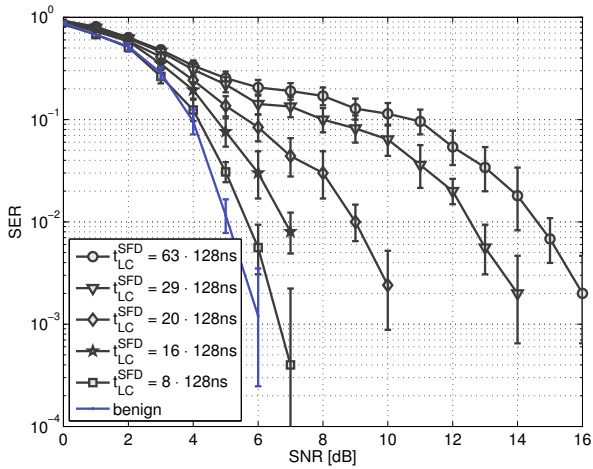

Figure 10.8: SER versus SNR for LPRF comparing benign performance to LC with varying LC delays $t_{\mathrm{LC}}^{\mathrm{SFD}}$.

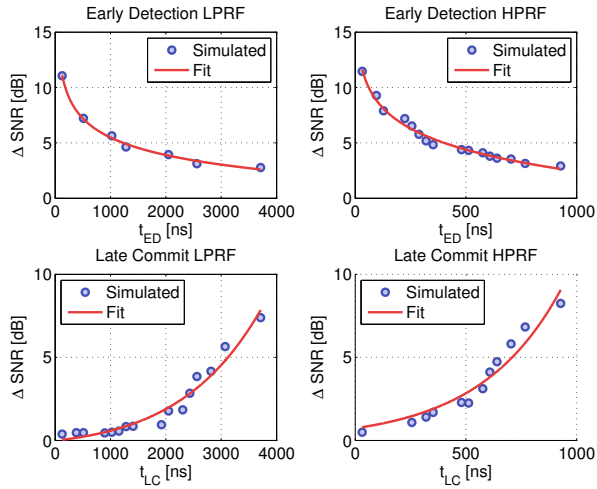

Figure 10.9: Performance loss ΔSNR with respect to benign case at fixed $\mathrm{SER} = 10^{-2}$ versus $t_{\mathrm{ED}}^{\mathrm{SFD}}$ and $t_{\mathrm{LC}}^{\mathrm{SFD}}$ for LPRF and HPRF.

10.3.1 Attack on the Preamble

An honest receiver performing SFD detection takes the entire length T_{sfd} of the SFD into account. For LPRF this equals $T_{\mathrm{sfd}} = 31.8\,\mu\mathrm{s}$, for HPRF $T_{\mathrm{sfd}} = 7.95\,\mu\mathrm{s}$.

Figure 10.7 shows the SER for an honest receiver, as well as for an adversary that performs early SFD detection with different early SFD detection delays $t_{\mathrm{ED}}^{\mathrm{SFD}}$. The curves shown here are for LPRF. Not surprisingly, the earlier an adversary performs SFD detection, the more additional received power with respect to an honest receiver it is going to cost him to reach a given level of SER. If we fix $\mathrm{SER} = 10^{-2}$, detecting the SFD at $t_{\mathrm{ED}}^{\mathrm{SFD}} = 3.712\,\mu\mathrm{s}$ costs the adversary $\Delta\mathrm{SNR} = 2.8$ dB in additional received power, detecting at $t_{\mathrm{ED}}^{\mathrm{SFD}} = 0.128\,\mu\mathrm{s}$ entails a cost of $\Delta\mathrm{SNR} = 11.2$ dB.

For $t_{\mathrm{ED}}^{\mathrm{SFD}}$, we only consider values shorter than the length of the first SFD symbol. Larger values for $t_{\mathrm{ED}}^{\mathrm{SFD}}$ do not make much sense for the adversary because they also force him to commit after the first SFD symbol, which is only possible at a considerable additional cost. This can be seen in Figure 10.8, which shows the SER of an adversary that commits late, at time $t_{\mathrm{LC}}^{\mathrm{SFD}}$ into the SFD. Again, the results shown are for LPRF. Committing at $t_{\mathrm{LC}}^{\mathrm{SFD}} = 8 \cdot 128\mathrm{ns} = 1.02\,\mu\mathrm{s}$, or earlier is within 0.6 dB of the benign case and thus comes at practically no additional cost

at a target SER of 10^{-2}. Committing later comes at an ever increasing cost: Committing at $t_{\mathrm{LC}}^{\mathrm{SFD}} = 29 \cdot 128\mathrm{ns} = 3.712\,\mu\mathrm{s}$, already costs $\Delta\mathrm{SNR} = 7.5$ dB. According to the preamble and SFD codes, no pulse is sent between the 29th code symbol of the first SFD symbol and the first code symbol of the third SFD symbol. So committing anywhere between $t_{\mathrm{LC}}^{\mathrm{SFD}} = 3.712\,\mu\mathrm{s}$ and $t_{\mathrm{LC}}^{\mathrm{SFD}} = 63 \cdot 128\mathrm{ns} = 8.064\,\mu\mathrm{s}$ is equivalent to committing at $t_{\mathrm{LC}}^{\mathrm{SFD}} = 8.064\,\mu\mathrm{s}$, which costs more than $\Delta\mathrm{SNR} = 9$ dB.

Results for HPRF are generally close to those of LPRF shown so far. Performing ED at $t_{\mathrm{ED}}^{\mathrm{SFD}} = 928$ ns, for example, costs the adversary about $\Delta\mathrm{SNR} = 2.9$ dB, compared to 2.8 dB for LPRF. This can be seen in Figure 10.9 where we show the additional cost ΔSNR with respect to an honest receiver versus $t_{\mathrm{ED}}^{\mathrm{SFD}}$ and $t_{\mathrm{LC}}^{\mathrm{SFD}}$ for both LPRF and HPRF and a fixed SER of 10^{-2}. The corresponding SNR values were found via interpolation of curves such as those shown in Figures 10.7 and 10.8. Results for ED are close and within 0.5 dB. Late commit generally costs about 1 dB more in the case of HPRF. Note the different time scales that are due to the fact that a preamble symbol in HPRF is four times shorter (see Section 2.5 of Chapter 2).

An important observation is that none of the curves showing the performance under attack exhibits an error floor. This indicates that by increasing the SNR, the attack success rate can be made arbitrarily large. The same holds for the payload, as we will see shortly.

Alternative SFD detection. For completeness, we also evaluated the effectiveness of the attack against a receiver that performs SFD detection using a correlation-based SFD detection method (which entails a 2 dB performance loss compared to the DESSERT algorithm, see Chapter 6). Such a receiver is also vulnerable to the attack, and the cost of the attack in terms of ΔSNR is close to the previously shown results: for $t_{\mathrm{LC}}^{\mathrm{SFD}}$ in the order of T_{psym}, we find a cost increase of 1 dB.

10.3.2 Attack on the Payload

We now look at the effect of ED and LC on the payload. The following results do not contain effects of synchronization: We assume here that the receiving party, ARX in the case of ED and HRX in the case of LC, is able to perfectly synchronize to each packet. Perfect synchronization here means that an oracle returns the exact packet time-of-arrival (Hence, there are no false alarms or missed detections.). The channel energy-delay profile is still estimated, but the estimation is performed under the assumption that the packet boundaries are perfectly aligned. In the case of LC, we further assume that the packet sent by ATX does not contain any errors due to a preceding ED.

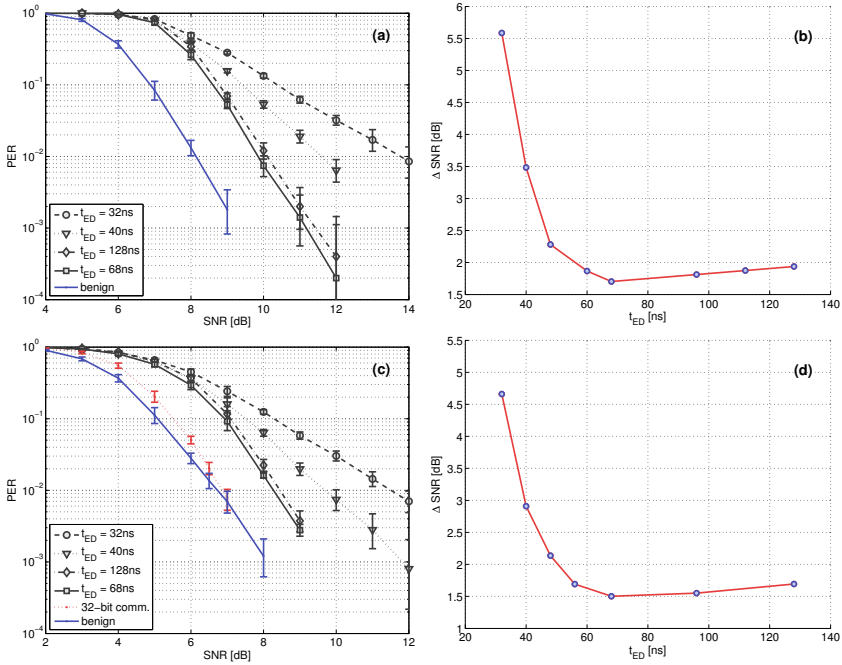

Figure 10.10: (a): PER versus SNR for the payload assuming nonce verification strategy I. We compare benign performance to ED with varying ED delays t_{ED}. The optimal t_{ED} is in the order of the channel delay spread and gives a loss of about 1.7 dB.
(b): More compact representation of the data in (a), showing the loss ΔSNR with respect to the benign case versus t_{ED} for a fixed PER of 10^{-2}.
(c) and (d): Corresponding results for nonce verification strategy II. The cost of the attack stays practically the same. For reference, we show the performance of a 32-bit communication packet in addition to the performance of ranging packets.

Figure 10.10(a) shows the PER at different SNRs for the LPRF mode and nonce verification strategy I. We show the performance curves for a benign receiver and an adversary performing ED at different ED delays t_{ED}. The optimal ED delay for the adversary is in the order of the channel delay spread and found to be $t_{\mathrm{ED}}^{\mathrm{OPT}} = 68$ ns in the present example that uses the NLOS channel model. Deciding on the symbol at $t_{\mathrm{ED}}^{\mathrm{OPT}}$ introduces a loss of about 1.7 dB with respect to the benign curve at a packet error rate of $\mathrm{PER} = 10^{-2}$. This can also be seen in Figure 10.10(b). Here we show the loss in SNR, ΔSNR, with respect to the benign case versus the ED delay t_{ED} for a target packet error rate of $\mathrm{PER} = 10^{-2}$. The curve has been obtained from curves such as those shown in Figure 10.10(a) via interpolation. Detecting after $t_{\mathrm{ED}}^{\mathrm{OPT}}$ gives a slightly worse performance because the adversary then merely integrates more noise instead of useful signal. Performing ED much earlier than $t_{\mathrm{ED}}^{\mathrm{OPT}}$ results in substantially larger loss because a large part of the useful signal energy is lost: Deciding at $t_{\mathrm{ED}} = 32$ ns, for example, introduces a loss of 5.6 dB.

Figures 10.10(c) and (d) show the corresponding results for nonce verification strategy II. In Figure 10.10(c) we additionally show the performance of a 32-bit communication packet that achieves a performance similar to the 42-bit ranging packet in the benign case. Nonce verification strategy II allows the honest devices to operate at lower SNRs. However, this has only little effect on the performance of the attack: the cost of the attack at $t_{\mathrm{ED}}^{\mathrm{OPT}}$ is now 0.2 dB lower than in case of nonce verification strategy I.

Figure 10.11 shows the performance of LC on the payload in the case of LPRF and nonce verification strategy I. As explained in Section 10.2.2, the LC delay t_{LC} is fixed to $t_{\mathrm{LC}} = T_{sym}/2 = 512$ ns. We show the PER for different ratios γ of the energies $E0$ and $E1$ corresponding to the signal energies transmitted by the adversary during the 0-block and 1-block, respectively. $E1$ here corresponds to the energy a benign receiver would transmit and $E0$ is typically smaller (see also Section 10.2.2). A ratio of $\gamma^{\mathrm{OPT}} = 0.35$ gives optimal performance throughout the whole operating range, thus this is the energy ratio we will use in all subsequent simulations. The optimal ratio gives a loss of about 4 dB with respect to the benign case. For nonce verification strategy II, the optimal ratio as well as the associated cost of the attack is identical and we therefore do not show any curves.

For HPRF, we do not show any curves either because the results are very similar to LPRF. With HPRF and the LOS channel, the optimal ED delay is $t_{\mathrm{ED}}^{\mathrm{OPT}} = 48$ ns. Note that this is significantly larger than the channel delay spread. The reason is that in the HPRF mode, a burst of 16 pulses is sent during the payload, spreading the received signal wider in time. The difference in SNR, with respect to the benign case, is 2 dB versus 1.7 dB with LPRF. For LC,

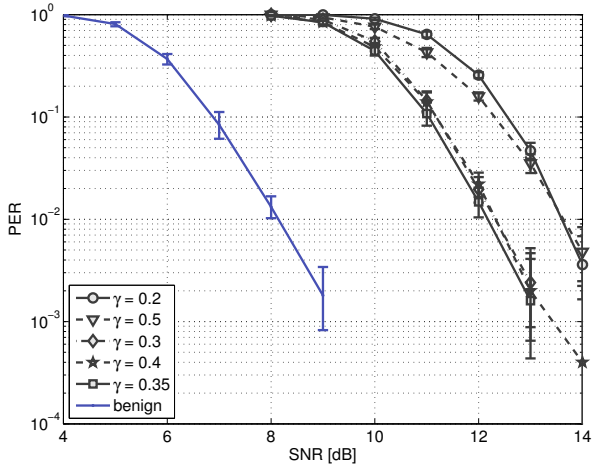

Figure 10.11: PER for LC on the payload with varying energy ratios γ. The optimal ratio at $\gamma^{\mathrm{OPT}} = 0.35$ gives a loss of about 4 dB with respect to the benign setting. The curves shown are for nonce verification strategy I.

we find the optimal energy ratio to be $\gamma^{\mathrm{OPT}} = 0.35$ as well, and the corresponding loss of 3.9 dB is close to the 4 dB found for LPRF.

Alternative Demodulation. We also evaluated a simplistic receiver that demodulates without weighting with the estimated energy-delay profile. Such a receiver is vulnerable to the attack as well, and the attack's cost in terms of ΔSNR is within 0.5 dB of the cost for the baseline receiver.

10.3.3 Overall Performance of the Attack

We now establish the overall performance of the distance-decreasing relay attack. As the relay attack involves two transmissions, ARX and HRX potentially have different received SNRs, which we will denote by $\mathrm{SNR}_{\mathrm{ED}}$ and $\mathrm{SNR}_{\mathrm{LC}}$. This difference can be a result of the topology, but it can also be introduced by the adversary. Depending on his abilities, an adversary can, for example, send with a higher power in order to increase $\mathrm{SNR}_{\mathrm{LC}}$, or move closer to HTX, or use a directive antenna to increase $\mathrm{SNR}_{\mathrm{ED}}$. Combined with the observation that the same relay time-gain, t_{relay}, can be obtained with different combinations of ED and LC delays, this gives the adversary room for a trade-off: Depending on the SNR values achievable for $\mathrm{SNR}_{\mathrm{ED}}$

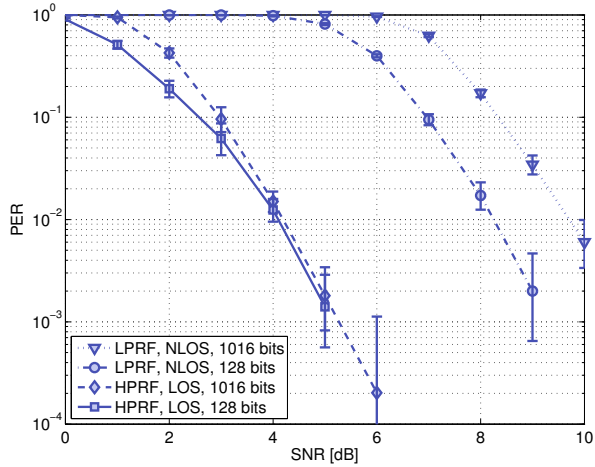

Figure 10.12: Reference curves showing PER in the benign case for LPRF and HPRF with different packet sizes and nonce verification strategy I.

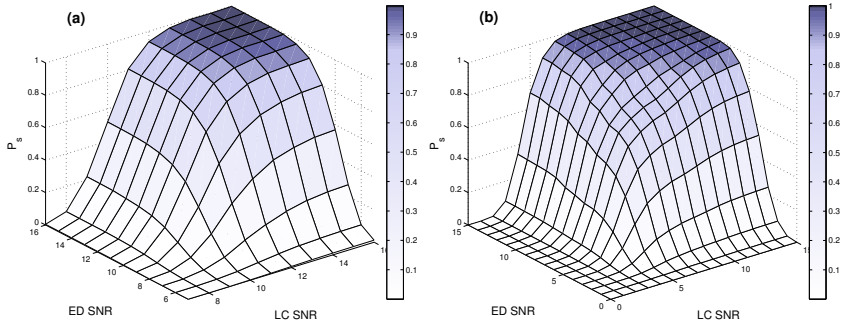

Figure 10.13: Probability of success, P_s, for an attack trying to achieve a distance decrease of 144 m. Packet length is 128bits. (a): With LPRF, $P_s > 99\%$ is reached at a cost of $(\Delta\text{SNR}_{\text{ED}} = 4\text{dB}, \Delta\text{SNR}_{\text{LC}} = 6\text{dB})$. (b): For HPRF $(7\text{dB}, 7\text{dB})$ gives $P_s > 99\%$.

($\mathrm{SNR}_{\mathrm{LC}}$, respectively) the adversary can choose to perform ED earlier or later (commit earlier or later, respectively). If $\mathrm{SNR}_{\mathrm{LC}}$ is high with respect to $\mathrm{SNR}_{\mathrm{ED}}$, the adversary will prefer to commit late in order to be able to detect late as well. If $\mathrm{SNR}_{\mathrm{LC}}$ is low with respect to $\mathrm{SNR}_{\mathrm{ED}}$, the adversary will prefer to detect early in order to be able to commit early.

In our analysis, Figure 10.12 will serve as a benchmark. It shows the PER in the benign case for both LPRF and HPRF with nonce verification strategy I. Packet sizes of 128 and 1016 data bits are shown. A packet size of 1016 bits is the maximum packet size allowed by the standard; as stated earlier, 128 bits correspond to a conservative length of a ranging message. In LPRF the factor limiting performance is the payload. This can be seen by observing that the LPRF curve for the shorter packet size is almost identical to the benign curve in Figure 10.10 (which assumes perfect synchronization). For HPRF the opposite is true, the limiting factor is the synchronization. This can be seen in Figure 10.12 where the size of the packet hardly influences the PER. The reason is that in HPRF 16 times more energy is sent in a payload symbol compared to a preamble pulse, whereas in LPRF it is only 4 times more.

Figure 10.13(a) shows the probability of success for LPRF and an attack that tries to gain 480 ns when relaying a 128bit packet between HTX and HRX. This relay time-gain is equivalent to a 144 m distance decrease between HTX and HRX.[6] The results shown are for different combinations of $\mathrm{SNR}_{\mathrm{ED}}$ and $\mathrm{SNR}_{\mathrm{LC}}$. For every SNR combination, the probability of success that is reported corresponds to the tuple of $(t_{\mathrm{ED}}^{\mathrm{SFD}}, t_{\mathrm{LC}}^{\mathrm{SFD}}, t_{\mathrm{ED}})$[7] yielding best performance among all the tuples that achieve the given relay time-gain of 480 ns. In the benign case we achieve a PER of approximately 10^{-2} at an SNR of around 8 dB, see Figure 10.12. In Figure 10.13(a), a probability of success of $P_s = 0.9869$ is achieved for the pair $(\mathrm{SNR}_{\mathrm{ED}} = 12\mathrm{dB}, \mathrm{SNR}_{\mathrm{LC}} = 14\mathrm{dB})$. For all pairs above $(12\mathrm{dB}, 14\mathrm{dB})$ the probability of success is above 99%. With respect to an honest transmitter-receiver pair, an adversary thus needs an additional 4 dB in SNR for ED and an additional 6 dB for LC, in order to reduce the distance by 144 m with a probability of success in the order of 99%. Attaining SNR values in this range would not pose much of a challenge to the adversary.

The corresponding HPRF results, decreasing the distance by 144 m, are shown in Figure 10.13(b). A probability of success of $P_s = 0.9875$ is reached at $(11\mathrm{dB}, 11\mathrm{dB})$. Compared with Figure 10.12, the additional cost is 7 dB for both ED and LC. Compared with LPRF, we thus see that decreasing the distance by the same amount costs a bit more in HPRF. This was

6. Here we assume for simplicity that the processing delays at the adversarial transceivers are zero. Processing delays are discussed in Section 10.2.3.

7. Recall that $t_{\mathrm{LC}} = 512$ ns is fixed, thus limiting to some extent the degrees of freedom on the payload part.

to be expected for several reasons. First of all we have seen that, contrary to LPRF, the performance is not limited by the payload but by the synchronization. We can thus not hope to achieve the ED/LC performance of the payload-only attacks shown in Figures 10.10 and 10.11. Second, to obtain a given relay time-gain on the preamble is more costly for HPRF because of the closer spacing of the pulses. We have seen in Figure 10.9 that detecting early at the ith code symbol or committing late at the jth code symbol costs roughly the same for both LPRF and HPRF. The distance decrease achieved corresponds to $(j-i)\cdot L\cdot T_c$ which depends on the length of a code symbol $L\cdot T_c$. At the same cost, the distance decrease achieved by HPRF is thus four times shorter than for LPRF.

Increasing the packet length to its allowable maximum of 1016 bits decreases slightly the probability of success: To reach $P_s > 99\%$ we see virtually no cost increase for HPRF. For LPRF the cost in $\mathrm{SNR}_{\mathrm{ED}}$ and $\mathrm{SNR}_{\mathrm{LC}}$ increases by about 1.5 dB each. A smaller distance decrease obviously comes at a lower cost: e.g., for an attack decreasing the distance by 100 m for HPRF with 128bit packets, we found the additional cost for $P_s \approx 99\%$ to be 5 dB for ED and 4 dB for LC. Compared to the corresponding attack achieving a decrease of 144 m, this signifies a 2 dB smaller cost in $\mathrm{SNR}_{\mathrm{ED}}$ and a 3 dB smaller cost in $\mathrm{SNR}_{\mathrm{LC}}$.

We have also conducted the same experiments for verification strategy II. Compared with nonce verification strategy I using 128 bit packets, we found the exact same costs of $(\Delta\mathrm{SNR}_{\mathrm{ED}} = 4\mathrm{dB}, \Delta\mathrm{SNR}_{\mathrm{LC}} = 6\mathrm{dB})$ and $(7\mathrm{dB}, 7\mathrm{dB})$ for LPRF and HPRF, respectively. We therefore do not show any additional figures.

10.4 Countermeasures

The goal of this section is to investigate possible countermeasures against the attacks that we introduced in the preceding sections.

Two main factors determine the quality of a countermeasure. The first is its effectiveness: the maximum distance by which the adversary can decrease the distance with the countermeasure in place. The second factor is its cost: how much does it cost for a receiver using the countermeasure to reach a similar benign case performance (if no attack is taking place) as a system without countermeasures deployed. As an additional third factor, we could consider whether or not the countermeasure is implementable within the limits imposed by the IEEE 802.15.4a standard.

10.4.1 Private Ranging Mode Achieves Only Weak Security

The IEEE 802.15.4a standard includes an optional private ranging mode that allows the legitimate participants to secretly agree on the preamble codes used in the ranging packets. Hence, the adversary does not know the exact structure of the preamble, which makes the attack on the preamble harder. Nevertheless, the honest devices can only choose among 8 allowable preamble codes, which offers little security. First, the adversary could simply guess the codes with a decent success probability. Second, the adversary could detect a packet using, in parallel, all 8 allowable codes. This can be done entirely in the digital domain by correlating the received signal with each of the 8 codes and choosing the one with the highest correlation output. What additionally helps the adversary is the fact that these codes were designed to have minimum cross-correlation. In summary, the private ranging mode only moderately increases the complexity of the distance-decreasing relay attack, and cannot be considered a valid countermeasure. Furthermore, it seems that the private ranging mode was principally designed for more complex coherent receivers: The preamble parameters that the private ranging mode employs imply strong inter-symbol interference (ISI), which a non-coherent receiver, such as the one used in our investigation, cannot cope with well.

A similar but more promising direction than the private ranging mode might be to use secret preamble codes, known only to the communicating honest nodes. This could make (early) preamble detection infeasible, at least within the constrained time budget available to the adversary to mount the relay attack. It is uncertain, but worth investigating, how such random codes without nice auto-correlation properties would affect the benign case performance. Alternatively, secret time-hopping sequences could be used to make early detection of payload symbols more difficult. This would also require further investigation. Both of these alternative approaches are obviously not compliant with the current IEEE 802.15.4a standard.

10.4.2 Decrease Payload Symbol Duration

One of the most straightforward countermeasures is to decrease payload symbol duration [162], as the distance-decreasing attack cannot decrease the distance by more than one symbol duration. This applies to the BPPM modulation used in IEEE 802.15.4a: if the symbol duration is T_{f}, the time-gain of the the attack we proposed in this paper is at most $T_{\mathrm{f}}/2$. [8]

Using Optional Modes of IEEE 802.15.4a

Decreasing the payload symbol duration as a countermeasure can be implemented even within the IEEE 802.15.4a standard. Some non-mandatory modes have symbols as short as 32 ns. However, reducing T_{f} to a value where the attack is not a threat (i.e.,, the maximum achievable distance-decrease is only a few meters), is not without effect on the benign performance. The first problem is ISI, which manifests itself if the symbol duration is close or below the channel delay spread. Low-complexity non-coherent receivers cannot cope well with ISI and even if some solutions exist, they entail a loss of $5 - 10$ dB in the benign case [189]. Furthermore, shorter symbols have less resilience to multi-user interference.

Switching to OOK Demodulation Through Early Detection

Alternatively, the symbol duration can be preserved, but the honest receiver can choose to only take into account the beginning of the symbol [163], essentially performing early detection with OOK demodulation at an offset t_{OOK} from the beginning of the symbol. This is particularly attractive in our case, as switching from BPPM demodulation to OOK demodulation significantly reduces the achievable time-gain: Indeed, $t_{\mathrm{LC}} \geq t_{\mathrm{relay}}$ can be reduced from 512 ns to a value in the order of the channel delay spread $T_{\mathrm{spread}} \approx 60$ ns (for optimal OOK performance). This corresponds to *at most* 18 m distance-decrease (assuming an unrealistic instant ED and no processing delays). Further, this solution does not induce any additional ISI and is compliant with the mandatory modes of the standard. Our simulations evaluating payload early detection (see Section 10.3.2) show that such a countermeasure based on OOK decreases the benign case performance by roughly 1.5 dB because half of the available information is discarded. Additional coding could potentially compensated for this degradation.

In the case of nonce verification strategy II, additional coding can essentially be achieved by increasing the length of the nonces (see Section 10.1.2). A performance evaluation of this countermeasure for LPRF is shown in Figure 10.14. Analog to Section 10.3, we assume that the performance goal is given by a required PER of 10^{-2} at an SNR of 6.5 dB. We have seen in Section 10.3 that without a countermeasure, this can be achieved with nonces of length $N_{\mathrm{nonce}} = 42$, yielding a security level of $P_{\mathrm{guess}} = 2^{-32}$. Figure 10.14 shows the required length of a nonce if the same performance is to be achieved with the countermeasure using different

8. We note that the attack proposed in this paper can be further improved (in terms of the achieved distance-decrease), by employing late commit techniques in the fashion of [163]. However, the additional challenge is the weighting by the channel mask performed by the baseline receiver. Evaluation of such attacks might thus be worth further investigation.

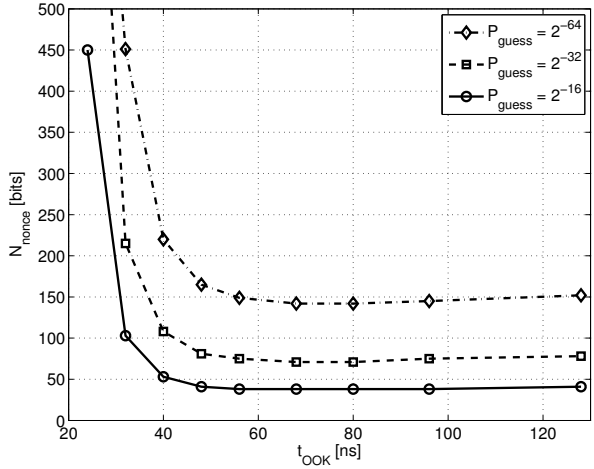

Figure 10.14: Cost of OOK countermeasure where the receiver decides on the bit value using OOK demodulation at time t_{OOK}. t_{OOK} also corresponds to the maximum theoretically achievable t_{relay}. The only cost that this countermeasure entails is an increase in the required nonce length N_{nonce}.

OOK detection times t_{OOK}. As long as $t_{\text{OOK}} \geq 32$ ns, the required length of a nonce is below the maximum IEEE 802.15.4a packet size of 1016 bits. By employing the countermeasure and increasing the number of bits per nonce from 42 to 108, we can, e.g., bring the maximum theoretically achievable distance-decrease to 40 ns (about 12 m). A realistically achievable distance-decrease will be significantly lower because this value does neither include the time it takes the adversary to perform early detection and late commit, nor any processing delays. At the same time, this countermeasure does not reduce the performance in terms of PER and we also keep the same security level against guessing attacks. The only cost that occurs is the cost of generating, sending and receiving the additional bits required for the longer nonces. Since every IEEE 802.15.4a payload carrying a nonce is preceded by a preamble of considerable length, and since a good deal of receiver complexity during reception stems from synchronization, we argue that the cost of adding a few bits to the payload is negligible.

10.5 Beyond Energy-Detection Receivers

So far, we assumed that both honest and adversarial devices are built upon a low-complexity non-coherent energy-detection architecture. Since the IEEE 802.15.4a standard also allows for coherent reception, an interesting question is how a relaxation of the assumption on the receiver architecture impacts the effectiveness of distance-decreasing attacks.

In the following, we will analyze three scenarios. First we will consider the case where the adversary employs a coherent Rake receiver, while the victim devices remain with the non-coherent energy-detection receiver (“Rake-vs-EnergyDetection”). Then, we will discuss the case where all of the devices involved use coherent reception (“Rake-vs-Rake”). Finally, we will look at an asymmetric scenario where only one of the honest devices uses a coherent architecture (“Rake-vs-Rake/EnergyDetection”). A summary of our findings, comparing the different scenarios, is given at the end of this chapter in Table 10.1.

10.5.1 Scenario “Rake-vs-EnergyDetection”: Adversary Uses Coherent Receiver for Maximum Performance

In this scenario, we assume that the adversary uses coherent Rake receivers, whereas the honest devices continue to use energy-detection receivers. This is a realistic scenario where honest devices strive for low-complexity and the adversary does not have this limitation.

At first sight it might seem that the adversary only has a limited benefit to employ Rake re-

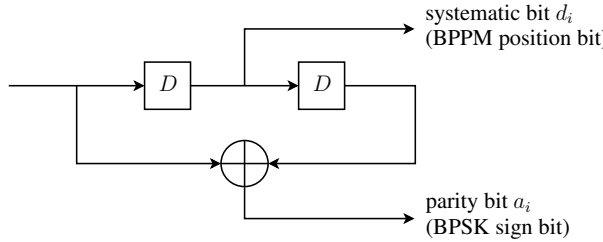

Figure 10.15: Structure of IEEE 802.15.4a convolutional code.

ceivers in this scenario. After all, we have seen in the foregoing sections that highly successful attacks can already be mounted with energy-detection receivers that have a lower complexity. Moreover, we cannot hope to achieve any performance gain for the late commit attack, since it is independent of the receiver used by the adversary. However, a closer analysis reveals a whole new space for an attack due to the convolutional code employed by IEEE 802.15.4a.

The systematic rate $1/2$ convolutional code is applied to the bits to be transmitted after RS encoding. The resulting systematic bits are modulated using BPPM and the parity bits using BPSK (see also Section 2.4.3 of Chapter 2). Both are transmitted together using a BPPM/BPSK symbol, resulting in the received signal model for the i-th symbol given by Equation (10.8) earlier in this chapter, and where d_i denotes the i-th systematic bit and a_i the i-th parity bit.

The IEEE 802.15.4a convolutional code has a constraint length of 3 and generator polynomials $g_1 = (0, 1, 0)$ and $g_2 = (1, 0, 1)$, resulting in the structure shown in Figure 10.15. The problem with this code from a security perspective is that the i-th parity bit carries information about the $i + 1$-th systematic bit. Indeed, we have that

$$a_i = d_{i-1} \oplus d_{i+1} \tag{10.12}$$

where $\oplus$ denotes modulo two addition. This implies that after demodulation of parity bit a_i, the adversary ARX can obtain full knowledge of the systematic bit d_{i+1} of the following symbol, even before it is transmitted by HTX. This allows an attack with a higher relay time-gain than in the case where the adversary uses an energy-detection receiver, which is illustrated in Figure 10.16.

The total relay time-gain of this attack is $t_{\text{relay}} = T_{\text{f}} + t_{\text{LC}} - t_{\text{ED}}$. With respect to the attacks that we have seen before, the adversary gains one frame duration T_{f} because he knows the systematic bit value one frame in advance thanks to the convolutional code.

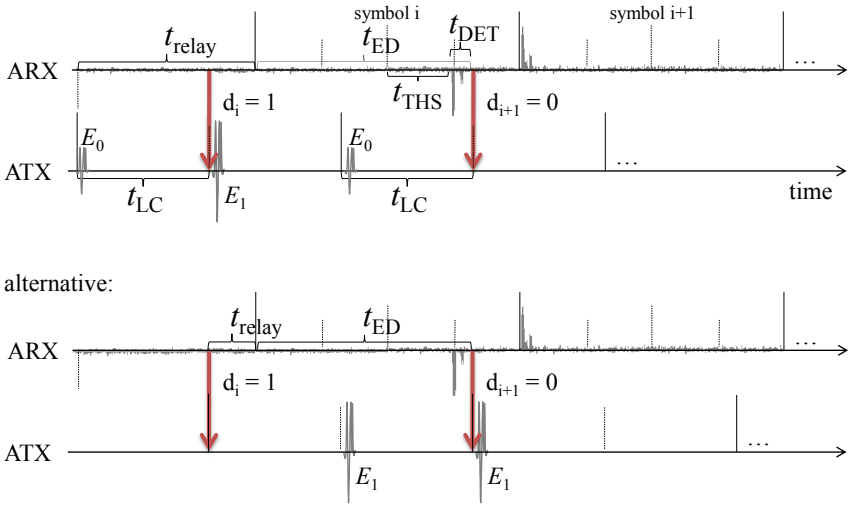

Figure 10.16: Distance decreasing relay attack on the payload symbol if the adversary uses a Rake receiver. $t_{\text{THS}} = c_{\text{THS},i} T_{\text{burst}}$ here denotes the time-hopping offset and t_{DET} the portion of the signal from the 1-block that is used in the demodulation of the symbol.

The early detection attack at ARX is slightly different than before. The adversary has to demodulate both the systematic bit and the parity bit in order to take advantage of the convolutional code. This implies that he can only decide on the bit value of the i-th bit after $t_{\text{ED},i} = T_{\text{f}}/2 + c_{\text{THS},i}T_{\text{burst}} + t_{\text{DET}}$. Here, t_{DET} denotes the portion of the signal from the 1-block of a frame that is used in the demodulation of the bit. In the following, we will refer to t_{DET} as *detection time*.

ATX can mount the same late commit attack that we have seen in Section 10.2.2, yielding a late commit delay of $t_{\text{LC},i} = T_{\text{f}}/2 + c_{\text{THS},i+1}T_{\text{burst}}$. Note that in contrast to the attacks with energy-detection receivers, here the time-hopping sequence matters since we commit to symbol $i+1$ (with time-hopping index $c_{\text{THS},i+1}$) while receiving symbol i (with time-hopping index $c_{\text{THS},i}$). The resulting achievable relay time-gain for the i-th bit is thus

$$t_{\text{relay},i} = T_{\text{f}} + t_{\text{LC},i} - t_{\text{ED},i} = T_{\text{f}} + (c_{\text{THS},i+1} - c_{\text{THS},i})T_{\text{burst}} - t_{\text{DET}} \tag{10.13}$$

The adversary has to ensure that the attack yields a consistent relay time-gain that remains realizable independently of the time-hopping indices of consecutive bits. Consequently, the maximum relay time-gain that is achievable over the entire packet is the minimum of (10.13) over all bits. [9] The minimum is obtained if $c_{\text{THS},i+1} = 0$ and $c_{\text{THS},i} = N_{\text{h}} - 1$, where $N_{\text{h}} - 1$ is the maximum time-hopping index (see Section 2.4.3 of Chapter 2). The resulting relay time-gain of this attack is then

$$t_{\text{relay}} = \frac{3}{4}T_{\text{f}} + T_{\text{burst}} - t_{\text{DET}} \tag{10.14}$$

Note that our performance evaluation in Section 10.3.1 shows that such a relay time-gain is easily achievable on the preamble where the attacker has more flexibility, even with an energy-detection receiver. In what follows, we therefore omit a further investigation of preamble attacks with a Rake receiver and focus on the payload.

Alternatively, the adversary can mount the attack completely without the late commit part and still obtain a relay time-gain of

$$t_{\text{relay}} = \frac{1}{4}T_{\text{f}} + T_{\text{burst}} - t_{\text{DET}} \tag{10.15}$$

this scenario is shown in the lower half of Figure 10.16. This alternative is particularly interesting as it is not detectable by the OOK countermeasure as we will see shortly.

9. If this were not the case, the attacker would have to guess bits for which the relay time-gain is not realizable. Such a strategy would of course be possible, yielding a slight improvement in terms of distance-decrease at the cost of a degradation of the attack's success probability.

Performance Evaluation

We evaluate the early detection attack on the payload described above with an optimal all-Rake receiver (aRake, see Chapter 2). Even though this receiver is idealistic and hardly realizable in practice, it gives the best performance an adversary can possibly hope to achieve. If the adversary were to use a suboptimal Rake receiver, the same attack still applies but the cost in terms of required SNR levels will be higher than the ones reported in this section.

We further assume perfect synchronization and that the receiver has perfect knowledge of the channel. Finally, we assume that the convolutional code is decoded with the optimal symbol-wise branch metric for BPPM/BPSK given in [110].

The results of the performance evaluation are shown in Figure 10.17, which shows the PER of the LPRF mode and for a packet size of 128 bits and nonce verification strategy I. Compared to the energy-detection receiver ("Energy-detection, benign") [10], the performance gain of the Rake receiver ("Rake, benign") is in the order of 12dB, half of which is due to the additional coding gain of the convolutional code (see also [110]).

To mount the early detection attack, the adversary has to deviate from optimal decoding of the convolutional code because each of the symbols has to be demodulated and decoded instantly. With regular decoding as in [110], the convolutional code is decoded at the end of the packet, when the full decoding trellis is available. This ensures that the maximum amount of redundancy can be taken into account. With the attack, the convolutional code can still be taken into account, however, only a partial trellis containing information about the symbols received so far is available. Such an "on-the-fly" decoding reduces the coding gain by roughly 2 dB ("Coding, on-the-fly"). Decoding in this fashion, however, is still not enough to mount the early detection attack. In addition, the adversary has to predict the position bit of the following symbol according to (10.12), which further reduces the usable redundancy provided by the convolutional code. Moreover, the symbol has to be detected as early as possible by varying the detection time t_{DET}. As long as $t_{\text{DET}} \leq 48$ ns, the additional loss is about 1.5 dB ("Coding"). With $t_{\text{ED}} = 48$ ns, we have a relay time-gain of $t_{\text{relay}} = 728$ ns. This corresponds to a distance-decrease of 218 m that costs 3.5 dB compared to normal Rake reception. For a faster decision, the adversary also has the option to decode every bit without taking the convolutional code into account ("No Coding"). The above mentioned attack that decreases the distance by 218 m costs in this case an additional 4 dB, which is still 4 dB below the performance of the energy-detection receiver.

10. The curve shown here for the energy-detection receiver corresponds to the one in Figure 10.10, i.e., perfect synchronization is assumed but not perfect estimation of the channel energy-delay profile.

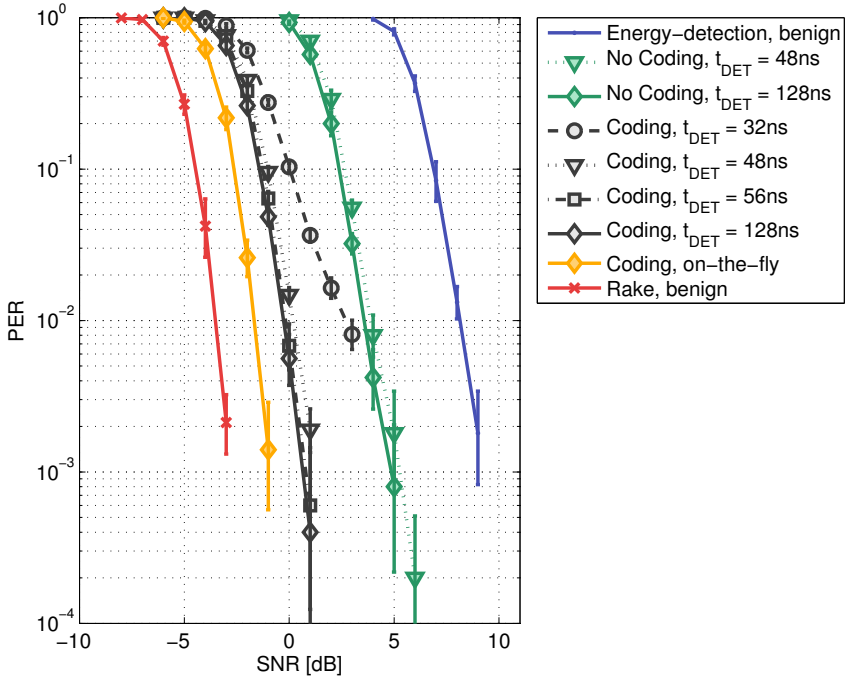

Figure 10.17: Performance of early detection attack on the payload if the adversary uses an aRake receiver. Early detection can be performed by decoding the convolutional code using partial information ("Coding") or by neglecting the convolutional code completely ("No Coding"). To achieve a relay time-gain, the current bit has to be detected within a short detection time t_{DET} allowing a prediction of the position bit of the next symbol. With respect to normal aRake operation ("Rake, benign"), this attacks costs about 3.5 dB if coding is used, and 7.5 dB if the code is neglected. In both cases, the performance is still significantly better than that of an energy-detection receiver ("Energy-detection, benign"). "Coding, on-the-fly" is a reference curve where decoding is done using partial information but future bits are not predicted, thus yielding no relay time-gain.

Countermeasures

The OOK countermeasure described in Section 10.4 can not sufficiently decrease the achievable relay time-gain to levels where this attack with a Rake receiver is not a threat. The alternative version of the attack, where no late commit is performed is not preventable since with this attack, the signal received by HRX is indistinguishable from a legitimate signal sent by HTX. With $t_{\text{DET}} = 48$ ns this still amounts to a relay time-gain of 216 ns and a distance-decrease of 65 m.

Countermeasures that aim at preventing early detection, e.g., the use of secret cryptographic preamble codes, may still be able to prevent this attack.

To counter this attack in the scope of the standard, however, we see only two possible options: 1) the honest receivers also use a coherent architecture, which makes them more complex; 2) if an energy-detection receiver is to be used, the convolutional code has to be changed or omitted. The requirement for any suitable convolutional code is that their structure needs to be such that all output bits must only depend on past information bits. However, changing the convolutional code requires a modification of the standard and is therefore an unlikely option. Omitting the code however, is realizable within the standard. IEEE 802.15.4a foresees two optional modes, one for LPRF and one for HPRF, where the convolutional code is not used. In both of these modes only a single pulse is sent per burst ($L_{\text{s}} = 1$, $T_{\text{burst}} = 2$ ns). With these modes, durations of the BPPM symbols are 256 ns for LPRF and 64 ns for HPRF. Due to the short symbol duration, the latter is only adequate for LOS environments, making the former the most suitable mode for secure ranging with energy-detection receivers.

10.5.2 Scenario "Rake-vs-Rake": High-End Coherent Ranging System

In this scenario, we assume a high-end secure ranging system, where all devices involved use coherent reception with Rake receivers.

With such a setting, the attack described in scenario "Rake-vs-EnergyDetection" is no longer possible. Since the honest devices are also capable of decoding the parity bits, the adversary now has to relay both the systematic and the parity bits correctly. However, only the systematic bit can be predicted from the previously received symbols. The maximum theoretically achievable relay time-gain is now upper-bounded by the symbol duration of a BPSK symbol, which is in the order of the channel spread and equals $60 - 70$ ns in the NLOS case. However, we have seen in the performance evaluation of scenario "Rake-vs-EnergyDetection" that a faster detection time of $t_{\text{DET}} = 48$ ns entails practically no performance loss and can

therefore also be adapted by a honest receiver. Considering additional processing delays, this leaves only little margin for the adversary, who can in the best case hope to achieve a distance-decrease in the order of 10 m. This is an order of magnitude lower than the attacks against energy-detection receivers. Similar to the countermeasure described in Chapter 10.4, the honest receivers can increase the length of the nonces which allows them to further decrease the detection time t_{DET} and thus the obtainable distance-decrease.

10.5.3 Scenario "Rake-vs-Rake/EnergyDetection": Asymmetric Scenario With High-End Reader and Low-End Tag

In this last scenario, we consider a hybrid between scenarios "Rake-vs-EnergyDetection" and "Rake-vs-Rake". We assume that all of the adversarial devices and one of the honest devices are coherent Rake receivers. The second honest device continues to use an energy-detection receiver. The motivation for this asymmetric scenario is a setting where one of the honest devices acts as a fixed reader with no constraints on complexity, while the other one is a low-end tag that must strive for low-complexity.

Ranging packets from the reader to the tag can be relayed according to the attack described in scenario "Rake-vs-EnergyDetection". Without any countermeasure in place, this results in a maximal relay time-gain given by (10.14). Packets from the tag to the reader, on the other hand, conform to scenario "Rake-vs-Rake". However, since the adversary already achieves a substantial relay time-gain in the transmission from the reader to the tag, he can relay from the tag to the reader without time-gain, or even by losing some time, and still decrease the distance. An alternative option might be to use an analog repeater that amplifies and forwards the signal pulse-by-pulse with minimal delay [190]. Finally, if tag and reader are within communication range, relaying can be omitted altogether. In all of these cases, the relay time-gain achieved is about half the relay time-gain of scenario "Rake-vs-EnergyDetection".

With the OOK countermeasure in place and a suitable mode that does not use the convolutional code, the achievable distance-decrease is reduced to the one of scenario "Rake-vs-Rake", which is also equivalent to the case where the attacker uses energy-detection receivers as well.

10.6 Conclusion

We have investigated the vulnerability of the IEEE 802.15.4a standard to physical layer distance-decreasing relay attacks. We demonstrated that without appropriate countermeasures and un-

Architecture		No counterm.	With countermeasure	
ARX	HRX	Dist.-decr.*	Type	Dist.-decr.*
Energy-detection	Energy-detection	144 m	OOK	max. 12 m
Rake	Energy-detection	218 m	OOK + Mode w/o conv. code	max. 12 m
Rake	Rake	max. 12 m	-	-
Rake	Rake/Energy-det.	109 m	OOK + Mode w/o conv. code	max. 12 m

*Not accounting for any processing delays at the adversary

Table 10.1: Achievable distance-decrease in various scenarios with different architectures at honest and adversarial devices.

less all of the honest devices perform ranging with coherent Rake receivers, an adversary can perform highly effective attacks resulting in a distance-decrease in the order of one hundred meters. E.g., in a scenarios where all honest and adversarial devices are energy-detection receivers, an adversary can decrease the distance by 144 m with an impressive success rate of 99% and at a cost of just a few dB in SNR with respect to normal system operation. A further increase in SNR allows the adversary to make the success rate arbitrarily large.

Moreover, we unveiled an anomaly in the IEEE 802.15.4a convolutional code that allows the adversary to further decrease the distance, provided that he uses a coherent Rake receiver. With this additional exploit, the achievable distance-decrease amounts to 218 m. In general, we find that the achievable distance decrease depends on the architecture of the receiver used by honest and adversarial devices. A summary of the achievable distance-decrease for each of the scenarios that we investigated is given in Table 10.1.

Nevertheless, we also find that secure ranging within IEEE 802.15.4a, keeping the maximum achievable distance-decrease at acceptable levels, is possible. Further, this is possible even if the honest receivers are based on energy-detection, under the condition that two prerequisites are met. The first prerequisite is to adopt a countermeasure that requires a change in the way symbol demodulation is performed. With the choice of a proper packet length, this countermeasure induces practically no additional cost. The second prerequisite is that the convolutional code must either be turned off or its encoder structure must be changed such that a prediction of future bits becomes impossible. Within the standard, only the former is possible and can be achieved by switching to an optional mode without convolutional code. Based on

our findings, we therefore also make a recommendation about which mode of the standard to use for secure ranging.

The evaluation performed in this chapter confirms that the physical layer, and in particular the physical layer of IEEE 802.15.4a, indeed provides a lot of space for attacks against ranging protocols if it is not handled carefully. In the design of secure protocols, a thorough analysis of the physical layer is therefore indispensable.

10.7 Acknowledgments

We would like to thank Yannick Do and Florence Le Goff for their extensive help with the simulations underlying Section 10.3.

Closing Remarks and Complementary Material

Chapter 11

Conclusion

In this thesis we demonstrated that the effect of multi-user interference (MUI) can be effectively mitigated on the physical layer (PHY). Further, we have shown this to be true also in the context of the IEEE 802.15.4a standard and even with energy-detection receivers that generally have a lower-complexity than coherent Rake receivers.

We believe that this possibility to cope with interference on the PHY is important for IR-UWB networks because of several reasons. First of all, we can never hope to absolutely contain interference. Even a coordinated system that perfectly isolates concurrent transmissions within the network, is vulnerable to interference from outside the network. The more networks are deployed, the bigger the chance that several of these networks interfere. We have seen in several chapters of this thesis that, e.g., with IEEE 802.15.4a such inter-network interference can be extremely detrimental, even if precautions such as the assignment of different preamble codes and different time-hopping sequences are taken at higher communication layers. A second point is related to the design of medium access control (MAC) protocols for IR-UWB networks. It can be shown that the optimal MAC protocols for IR-UWB are uncoordinated, allow for concurrent transmissions, and can thus be greatly simplified, provided that interference mitigation is employed on the PHY (see, e.g., [1] and the references therein). Most recent proposals for IR-UWB MAC protocols [54, 97] including the IEEE 802.15.4a MAC [35] follow this principle. We have shown in the case of IEEE 802.15.4a that the benefit of concurrent transmissions is completely lost if proper interference mitigation techniques on the PHY are not employed.

Further, we have shown that an interference robust design necessitates a system level approach, where all of the components involved in receiving a data packet are designed appropriately. Interference affects synchronization as well as it affects channel estimation or data

decoding. Increasing the robustness to MUI of only one of these mechanisms in isolation will therefore not necessarily result in a design that is more robust overall. Consequently, we presented concrete solutions, some of which have recently successfully been patented, for each of these mechanisms.

First of all, we have addressed synchronization and shown that robust packet detection and timing acquisition algorithms that rely on simple thresholding mechanisms are able to mitigate MUI, showing a near perfect capture effect. However, to allow for concurrent transmissions, the receiver also has to be able to detect the end of the preamble in a reliable fashion, even in the presence of near interferers. This necessitates a robust detection of the start frame delimiter (SFD), which is ignored by most of the related work, even in the single user case. We compared different SFD detection algorithms and showed how they can be made robust to MUI. For data decoding we have presented and compared different solutions that are robust to MUI in the context of both coherent and non-coherent receivers. Our solutions range from simple thresholding schemes to more sophisticated methods that model interference according to a non-Gaussian model and that had previously not been considered in the context of IR-UWB. What is common to all our solutions is their ability to significantly reduce the effect of MUI: in various scenarios we found an improvement in packet error rate of at least one order of magnitude, compared to traditional receiver designs. Further, none of our schemes assumed perfect knowledge about any of the involved parameters, such as, e.g., the channel delay profile. We rather proposed methods to also perform parameter estimation in a MUI resistant fashion.

Further, we have shown that even quite simple energy-detection receivers, that adapt to varying channel conditions, are vulnerable to clock drifts, especially if the oscillators driving the transmitter and receiver clocks are of average quality. We proposed two solutions that can be implemented without increasing the hardware complexity. In particular, our algorithm that is based on the Radon transform showed a performance close to the optimum. What additionally makes this algorithm promising is its potential to also perform other receiver operations such as channel and time-of-arrival estimation.

Finally, this thesis includes a detailed analysis of secure ranging over IR-UWB. We demonstrated that if no appropriate countermeasures are taken, the IEEE 802.15.4a standard does not provide security against distance-decreasing attacks on the PHY. To this end we proposed a set of attacks, targeting both synchronization and payload decoding, and showed that they allow an adversary to significantly alter the measured distance between two devices with an almost perfect success probability and without breaking any cryptographic protocols. We further analyzed different scenarios involving different receiver architectures and found that they are not

without impact on the effectiveness of distance-decreasing attacks. Based on these findings we evaluated possible countermeasures and made recommendations on how to improve the security of IR-UWB ranging.

11.1 Future Work and Possible Extensions

We proposed solutions of various levels of complexity resulting also in various levels of robustness to MUI. Most of them obviously do not come for free. Some solutions like the robust energy-detection architecture for burst transmissions presented in Chapter 5 require additional hardware circuitry, which will increase both the manufacturing cost as well as the energy consumption of the receiver. Other mechanisms, e.g., the robust synchronization algorithms in Chapter 6 require the calculation of additional thresholds.

One of the most significant advances of the work presented here would be the experimental validation of our algorithms in a real system. This should become possible as IR-UWB hardware becomes more widely available. A real-world implementation would allow for an exact quantification of the costs involved with the different trade-offs that we discussed above.

On the other hand, not mitigating interference on the PHY also has its cost: packets that are lost due to interference need to be retransmitted; preventing packet collisions requires sophisticated protocols to coordinate between nodes. There is thus another trade-off between interference mitigation on the PHY versus interference management at higher communication layers. Another possible direction for future work is therefore the evaluation of a complete network. For complexity reasons, we were restricted to a single link, when assessing the performance of IR-UWB systems in a multi-user environment. A next step would be to capitalize on the insights that we gained in order to build more sophisticated models that allow for a network-wide evaluation. Such models could also include optional IEEE 802.15.4a clear channel assessments modes and would undoubtedly lead to further insights into trade-offs between interference management at different layers.

In a less broad perspective, we have already hinted at possible improvements of our work in the individual chapters. For example, it might be worth investigating whether some of the optimal decision rules, e.g., for the energy-detection receiver assuming a gamma mixture distribution, can be simplified to yield suboptimal receivers that are easier to implement but still show a robustness to MUI comparable to the optimal solutions.

Another point that future work should address is an evaluation of clock-offset tracking under MUI. Further, we have seen that the Radon tracking algorithm proposed in Part II might

offer interesting perspectives also for time-of-arrival and channel estimation both of which are certainly worth exploring.

In terms of security in IR-UWB networks in general, and of IR-UWB ranging in particular, investigations into improved location privacy would be worthwhile. Although the IEEE 802.15.4a standard already proposes some solutions, we believe that they are insufficient due to their restricted nature. A promising direction seems to be to use cryptographic modulation, where, e.g., the time-hopping sequence depends on a shared secrets. Such a mechanism might also be an interesting option to thwart distance-decreasing attacks. Finally, understanding the effectiveness of distance-decreasing attacks that are less greedy than the ones presented here can be important for applications relying on highly accurate secure distance measurements. In such applications much harm might already be done by an adversary that decreases the distance by only a few meters, scraping together a few nanoseconds here and there.

Appendix A

Appendix

A.1 Mean Clock-Offset Estimation Error of Perfect Algorithm

Let ϵ_{tx} and ϵ_{rx} be the clock offsets with respect to a global reference clock of the transmitter and the receiver, respectively. Assume that both of these quantities are independently drawn from the same uniform distribution, i.e.,

$$\epsilon_{\text{tx}} \sim U(-\epsilon_{\text{max}}, \epsilon_{\text{max}}), \quad \epsilon_{\text{rx}} \sim U(-\epsilon_{\text{max}}, \epsilon_{\text{max}}) \tag{A.1}$$

where ϵ_{max} is the maximum absolute value of the clock offset with respect to the reference clock.

The relative offset between transmitter and receiver $\epsilon = \epsilon_{\text{rx}} - \epsilon_{\text{tx}}$ is then distributed according to a triangular distribution with PDF

$$f_{\triangle}(\epsilon|\epsilon_{\text{max}}) = \begin{cases} \frac{1}{\epsilon_{\text{max}}}(\frac{1}{\epsilon_{\text{max}}}\epsilon + 1) & \text{if } -\epsilon_{\text{max}} \leq \epsilon < 0 \\ \frac{1}{\epsilon_{\text{max}}} & \text{if } \epsilon = 0 \\ \frac{1}{\epsilon_{\text{max}}}(-\frac{1}{\epsilon_{\text{max}}}\epsilon + 1) & \text{if } 0 < \epsilon \leq \epsilon_{\text{max}} \\ 0 & \text{otherwise.} \end{cases} \tag{A.2}$$

Now assume a perfect algorithm that always estimates the correct clock offset but only has a resolution of $\Delta\epsilon$, i.e., it always finds the estimate $\hat{\epsilon}$ from the set

$$\hat{\epsilon} \in \{-\epsilon_{\text{max}}, \ldots - 2\Delta\epsilon, -\Delta\epsilon, 0, \Delta\epsilon, 2\Delta\epsilon, \ldots, \epsilon_{\text{max}}\}$$

that is closest to the true value ϵ. For simplicity we assume that the resolution is such that $\Delta\epsilon$ divides ϵ_{max}.

The expected value of the absolute error of such an algorithm is then given by

$$\mathbb{E}[|\hat{\epsilon} - \epsilon|] = 2\left[\sum_{k=0}^{\frac{\epsilon_{\text{max}}}{\Delta\epsilon}-1} \int_{k\Delta\epsilon}^{k\Delta\epsilon+\frac{\Delta\epsilon}{2}} (\epsilon - k\Delta\epsilon) f_{\triangle}(\epsilon|\epsilon_{\text{max}})\,\mathrm{d}\epsilon + \sum_{k=1}^{\frac{\epsilon_{\text{max}}}{\Delta\epsilon}} \int_{k\Delta\epsilon-\frac{\Delta\epsilon}{2}}^{k\Delta\epsilon} (k\Delta\epsilon - \epsilon) f_{\triangle}(\epsilon|\epsilon_{\text{max}})\,\mathrm{d}\epsilon\right] \tag{A.3}$$

where the factor of two follows from the symmetry of the triangular PDF, each term in the first sum accounts for the mean absolute error during the first half of the k-th interval of length $\Delta\epsilon$ of the positive part of the PDF, and the terms in the second sum accounts for the mean absolute error during the second half of the k-th interval.

Plugging (A.2) into (A.3) yields after tedious calculations involving only basic algebra

$$\mathbb{E}[|\hat{\epsilon} - \epsilon|] = \frac{\Delta\epsilon}{4} \tag{A.4}$$

Publications

Published

- M. Flury, M. Poturalski, P. Papadimitratos, J.-P. Hubaux, and J.-Y. Le Boudec, "Effectiveness of Distance-decreasing Attacks Against Impulse Radio Ranging", *3rd ACM Conference on Wireless Network Security (WiSec'10)*, Hoboken, NJ, USA, March 22-24 2010.

- M. Flury, R. Merz, and J.-Y. Le Boudec, "Robust IEEE 802.15.4a Energy Detection Receiver Using Statistical Interference Modeling", *Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, CA, USA, November 2009.

- M. Flury, R. Merz, and J.-Y. Le Boudec, "Robust Non-Coherent Timing Acquisition in IEEE 802.15.4a IR-UWB Networks" *IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC 2009)*, Tokio, Japan, 13-16 September 2009.

- M. Flury, R. Merz, and J.-Y. Le Boudec, "Clock-Offset Tracking Software Algorithms For IR-UWB Energy-Detection Receivers", *IEEE International Conference on Ultra-Wideband (ICUWB 2009)*, Vancouver, Canada, 9-11 September 2009

- M. Flury, R. Merz, J.-Y. Le Boudec, "Method for Estimating and Correcting a Drift Between Clocks of a Receiving Transceiver and a Corresponding Emitting Transceiver, and Receiver for Implementing Said Method", *Patent filed with the European Patent Office*, September 2009

- M. Flury, R. Merz, J.-Y. Le Boudec, "An Energy Detection Receiver Robust to Multi-User Interference for IEEE 802.15.4a Networks", *IEEE International Conference on Ultra-Wideband (ICUWB 2008)*, Hannover, Germany, 10-12 September 2008

- M. Flury, R. Merz, J.-Y. Le Boudec, "Method for Retrieving Data from Ultra Wideband Radio Transmission Signals", *US Patent Application 61/006,972*, February 2008

- M. Flury, R. Merz, J.-Y. Le Boudec and J. Zory, "Performance Evaluation of an IEEE 802.15.4a Physical Layer with Energy Detection and Multi-User Interference", *IEEE International Conference on Ultra-Wideband (ICUWB 2007)*, Singapore, 24-26 September 2007

- M. Flury, R. Merz and J.-Y. Le Boudec, "Managing Impulsive Interference in Impulse Radio UWB Networks", *ST Journal of Research*, vol. 4, no 1, May 2007

- M. Flury and J.-Y. Le Boudec, "Interference Mitigation by Statistical Interference Modeling in an Impulse Radio UWB Receiver", *IEEE International Conference on Ultra-Wideband (ICUWB 2006)*, Waltham, MA, USA, 24-27 September 2006

In Preparation

- M. Flury, M. Poturalski, P. Papadimitratos, J.-P. Hubaux, and J.-Y. Le Boudec, "Distance-decreasing Attacks Against Impulse Radio Ranging", to be submitted to *IEEE Transactions on Wireless Communications*.

- M. Flury, R. Merz, and J.-Y. Le Boudec, "Synchronization Algorithms for Impulse-Radio UWB Energy-Detection Receivers in IEEE 802.15.4a Networks", to be submitted to *IEEE Transactions on Wireless Communications*.

Curriculum Vitæ

Manuel Flury was born in 1979 in Switzerland. Swiss German by origin, he took the plunge and moved to French speaking Switzerland in 1999 to pursue his studies in Communication Systems Engineering. In 2005, he earned a Master of Science degree in Communication Systems from Ecole Polytechnique Fédérale de Lausanne (EPFL).

In 2001/2002 he studied one year as an exchange student at the Royal Institute of Technology (KTH) in Stockholm, Sweden, where he followed Computer Science and Electrical Engineering courses given in English and Swedish.

In 2003 he was accepted for a six month internship at the Nokia Research Center in Helsinki, Finland. During his internship he contributed to the development of a next generation base-station prototype for cellular networks.

In 2004/2005 he successfully completed his master thesis in the scope of another six month internship, this time in the United States at Qualcomm, located in San Diego, California. During his time at Qualcomm he was exposed the first time to interference in wireless networks, his thesis topic being on interference suppression in third generation cellular networks.

In 2005, he joined the Laboratory for Computer Communication and Applications (LCA) at EPFL, School of Computer and Communication Sciences, and began working on his PhD thesis under the supervision of Professor Jean-Yves Le Boudec. There, he participated to the National Center of Competence in Research on Mobile Information and Communication Systems (NCCR-MICS).

During his PhD, he was teaching assistant for courses in TCP/IP Networking, Performance Evaluation and Java Programming. Further, he supervised several students doing their semester and master thesis projects.

His current research interests are in wireless communication and computer networks.

Bibliography

[1] J.-Y. Le Boudec and R. Merz, "Concurrent and Parallel Transmissions are Optimal for Low Data-Rate IR-UWB Networks," in *IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC 2008)*, 2008.

[2] Federal Communications Commission, United States, "FCC 02–48, First Report and Order, available at http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-02-48A1.pdf," February 14th 2002.

[3] European Radiocommunications Office (ERO), Electronic Communications Committee, "ECC Decision of 24 March 2006 amended 6 July 2007 at Constanta on the harmonised conditions for devices using Ultra-Wideband (UWB) technology in bands below 10.6 GHz, Document ECC/DEC/(06)04, available at http://www.ero.dk/documentation/docs/doc98/official/pdf/ECCDEC0604.pdf," July 2007.

[4] IEEE Computer Society, LAN/MAC Standard Committee, "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), 2006.

[5] ECMA International, "Standard ECMA-368, High Rate Ultra Wideband PHY and MAC Standard, 3rd edition, available at http://www.ecma-international.org/publications/files/ECMA-ST/ECMA-368.pdf," December 2008.

[6] ——, "Standard ECMA-369, MAC-PHY Interface for ECMA-368, 3rd edition, available at http://www.ecma-international.org/publications/files/ECMA-ST/ECMA-369.pdf," December 2008.

[7] "Wireless USB," http://www.usb.org/developers/wusb/.

[8] R. Scholtz, "Multiple access with time-hopping impulse modulation," in *IEEE Military Communications Conference (MILCOM)*, vol. 2, 1993, pp. 447–450 vol.2.

[9] M. Z. Win and R. A. Scholtz, "Impulse radio: how it works," *IEEE Commun. Lett.*, vol. 2, no. 2, pp. 36–38, 1998.

[10] ——, "Ultra-wide bandwidth time-hopping spread-spectrum impulse radio for wireless multiple-access communications," *IEEE Trans. Commun.*, vol. 48, no. 4, pp. 679–691, April 2000.

[11] H. Hashemi, "The indoor radio propagation channel," *Proc. IEEE*, vol. 81, no. 7, pp. 943–968, 1993.

[12] A. F. Molisch, "Ultra-wide-band propagation channels," *Proc. IEEE*, vol. 97, no. 2, pp. 353–371, 2009.

[13] A. Saleh and R. Valenzuela, "A statistical model for indoor multipath propagation," *IEEE J. Sel. Areas Commun.*, vol. 5, no. 2, pp. 128–137, 1987.

[14] A.-F. Molisch, K. Balakrishnan, C.-C. Chong, S. Emami, A. Fort, J. Karedal, J. Kunisch, H. Schantz, U. Schuster, and K. Siwiak, "IEEE 802.15.4a channel model - final report, document 04/662r1," http://www.ieee802.org/15/pub/TG4a.html, November 2004.

[15] A. F. Molisch, D. Cassioli, C. C. Chong, S. Emami, A. Fort, B. Kannan, J. Karedal, J. Kunisch, H. G. Schantz, H. G. Schantz, K. Siwiak, and M. Z. Win, "A comprehensive standardized model for ultrawideband propagation channels," *IEEE Trans. Antennas Propag.*, vol. 54, no. 11, pp. 3151–3166, 2007.

[16] M. Chiani and A. Giorgetti, "Coexistence between UWB and narrow-band wireless communication systems," *Proc. IEEE*, vol. 97, no. 2, pp. 231–254, 2009.

[17] M. Z. Win and R. A. Scholtz, "On the energy capture of ultrawide bandwidth signals in dense multipath environments," *IEEE Commun. Lett.*, vol. 2, no. 9, pp. 245–247, 1998.

[18] A. F. Molisch, J. R. Foerster, and M. Pendergrass, "Channel models for ultrawideband personal area networks," *IEEE Trans. Wireless Commun.*, vol. 10, no. 6, pp. 14–21, 2003.

[19] M. Z. Win, G. Chrisikos, and N. R. Sollenberger, "Performance of rake reception in dense multipath channels: implications of spreading bandwidth and selection diversity order," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 8, pp. 1516–1525, 2000.

[20] D. Cassioli, M. Z. Win, F. Vatalaro, and A. F. Molisch, "Low complexity rake receivers in ultra-wideband channels," *IEEE Trans. Wireless Commun.*, vol. 6, no. 4, pp. 1265–1275, 2007.

[21] J. D. Choi and W. E. Stark, "Performance of ultra-wideband communications with suboptimal receivers in multipath channels," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 9, pp. 1754–1766, 2002.

[22] Z. Tian and G. B. Giannakis, "BER sensitivity to mistiming in ultra-wideband impulse radios-part I: nonrandom channels," *IEEE Trans. Signal Process.*, vol. 53, no. 4, pp. 1550–1560, 2005.

[23] W. M. Lovelace and J. K. Townsend, "The effects of timing jitter and tracking on the performance of impulse radio," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 9, pp. 1646–1651, 2002.

[24] H. Sheng, R. You, and A. M. Haimovich, "Performance analysis of ultrawideband rake receivers with channel delay estimation errors," in *Proc. Conf. on Information Sciences and Syst. (CISS), Princeton, NJ*, March 2004, pp. 921–926.

[25] I. Guvenc and H. Arslan, "Performance evaluation of UWB systems in the presence of timing jitter," in *Ultra Wideband Systems and Technologies, 2003 IEEE Conference on*, 2003, pp. 136–141.

[26] R. Hoctor and H. Tomlinson, "Delay-hopped transmitted-reference RF communications," in *IEEE Conference on Ultra Wideband Systems and Technologies*, 2002, pp. 265–269.

[27] M. R. Casu and G. Durisi, "Implementation aspects of a transmitted-reference UWB receiver," *Wireless Communications and Mobile Computing*, vol. 5, no. 5, pp. 537–549, 2005.

[28] H. Urkowitz, "Energy detection of unknown deterministic signals," *Proc. IEEE*, vol. 55, no. 4, pp. 523–531, 1967.

[29] Y. Souilmi and R. Knopp, "On the achievable rates of ultra-wideband PPM with non-coherent detection in multipath environments," in *IEEE International Conference on Communications (ICC)*, vol. 5, 2003, pp. 3530–3534 vol.5.

[30] C. Carbonelli and U. Mengali, "M-ppm noncoherent receivers for UWB applications," *IEEE Trans. Wireless Commun.*, vol. 5, no. 8, pp. 2285–2294, 2006.

[31] M. Weisenhorn and W. Hirt, "ML receiver for pulsed UWB signals and partial channel state information," in *IEEE International Conference on Ultra-Wideband*, September 2005, p. 6.

[32] ——, "Robust noncoherent receiver exploiting UWB channel properties," in *IEEE joint Conference on Ultra Wideband Systems and Technologies & International Workshop on Ultra Wideband Systems*, May 2004, pp. 156–160.

[33] A. A. D'Amico, U. Mengali, and E. Arias-De-Reyna, "Energy-detection UWB receivers with multiple energy measurements," *IEEE Trans. Wireless Commun.*, vol. 6, no. 7, pp. 2652–2659, 2007.

[34] "Zigbee alliance, http://www.zigbee.org," January 2008.

[35] IEEE Computer Society, LAN/MAC Standard Committee, "IEEE P802.15.4a/D7 (amendment of IEEE std 802.15.4), part 15.4: Wireless medium access control (MAC) and physical layer (PHY) specifications for low-rate wireless personal area networks," Jan. 2007.

[36] Z. Lei, F. Chin, and Y.-S. Kwok, "UWB ranging with energy detectors using ternary preamble sequences," in *Wireless Communications and Networking Conference, 2006. WCNC 2006. IEEE*, vol. 2, April 2006.

[37] R. Tesi, M. Hämäläinen, J. Iinatti, J. Oppermann, and V. Hovinen, "On the multi-user interference study for ultra wideband communication systems in AWGN and modified saleh-valenzuela channel," in *IEEE joint Conference on Ultra Wideband Systems and Technologies & International Workshop on Ultra Wideband Systems*, 2004, pp. 91–95.

[38] A. R. Forouzan, M. Nasiri-Kenari, and J. A. Salehi, "Performance analysis of ultrawideband time-hopping code division multiple access systems: uncoded and coded schemes," in *IEEE International Conference on Communications (ICC)*, vol. 10, 2001, pp. 3017–3021.

[39] G. Durisi and G. Romano, "On the validity of gaussian approximation to characterize the multiuser capacity of UWB TH PPM," in *IEEE Conference on Ultra Wideband Systems and Technologies*, 2002, pp. 157–161.

[40] B. Hu and N. C. Beaulieu, "Exact bit error rate analysis of TH-PPM UWB systems in the presence of multiple-access interference," *IEEE Commun. Lett.*, vol. 7, no. 12, pp. 572–574, 2003.

[41] B. Hu and N. Beaulieu, "Accurate evaluation of multiple-access performance in TH-PPM and TH-BPSK UWB systems," *IEEE Trans. Commun.*, vol. 52, no. 10, pp. 1758–1766, October 2004.

[42] A. Forouzan, M. Nasiri-Kenari, and J. Salehi, "Performance analysis of time-hopping spread-spectrum multiple-access systems: uncoded and coded schemes," *IEEE Trans. Wireless Commun.*, vol. 1, no. 4, pp. 671–681, October 2002.

[43] Y. Dhibi and T. Kaiser, "On the impulsiveness of multiuser interferences in TH-PPM-UWB systems," *IEEE Trans. Signal Process.*, vol. 54, no. 7, pp. 2853–2857, 2006.

[44] D. Middleton, "Statistical-physical models of electromagnetic interference," *IEEE Trans. Electromagn. Compat.*, vol. EMC-19, no. 3, pp. 106–127, 1977.

[45] ——, "Canonical and quasi-canonical probability models of class a interference," *IEEE Trans. Electromagn. Compat.*, vol. EMC-25, no. 2, pp. 76–106, 1983.

[46] A. Spaulding and D. Middleton, "Optimum reception in an impulsive interference environment–part I: Coherent detection," *IEEE Trans. Commun.*, vol. 25, no. 9, pp. 910–923, 1977.

[47] ——, "Optimum reception in an impulsive interference environment–part II: Incoherent reception," *IEEE Trans. Commun.*, vol. 25, no. 9, pp. 924–934, 1977.

[48] K. Vastola, "Threshold detection in narrow-band non-gaussian noise," *IEEE Trans. Commun.*, vol. 32, no. 2, pp. 134–139, 1984.

[49] G. D. Weeks, J. K. Townsend, and J. A. Freebersyser, "Performance of hard decision detection for impulse radio," in *IEEE Military Communications Conference (MILCOM)*, vol. 2, 1999, pp. 1201–1206 vol.2.

[50] R. Knopp and Y. Souilmi, "Achievable rates for UWB peer-to-peer networks," in *International Zurich Seminar on Communications*, 2004, pp. 82–85.

[51] J. Fiorina, "On the benefit of a one-bit sampling receiver and hard decoding in impulse radio ultra wide band communications with multi-user interferences," in *Personal, Indoor and Mobile Radio Communications, 2006 IEEE 17th International Symposium on*, 2006, pp. 1–4.

[52] W. Lovelace and J. Townsend, "Chip discrimination for large near far power ratios in UWB networks," in *IEEE Military Communications Conference (MILCOM)*, vol. 2, October 2003, pp. 13–16.

[53] W. M. Lovelace and J. K. Townsend, "Threshold discrimination and blanking for large near-far power ratios in UWB networks," *IEEE Trans. Commun.*, vol. 53, no. 9, pp. 1447–1450, 2005.

[54] R. Merz, J. Widmer, J.-Y. Le Boudec, and B. Radunovic, "A joint PHY/MAC architecture for low-radiated power TH-UWB wireless ad-hoc networks," *Wireless Commun. and Mobile Comput. J., Special Issue on Ultrawideband (UWB) Communications*, vol. 5, no. 5, pp. 567–580, August 2005.

[55] H. Shao and N. C. Beaulieu, "A novel zonal UWB receiver with superior performance," *Communications, IEEE Transactions on*, vol. 57, no. 4, pp. 1197–1206, 2009.

[56] N. Guney, H. Delic, and M. Koca, "Robust detection of ultra-wideband signals in non-gaussian noise," *IEEE Trans. Microw. Theory Tech.*, vol. 54, no. 4, pp. 1724–1730, 2006.

[57] N. C. Beaulieu and B. Hu, "An adaptive threshold soft-limiting UWB receiver with improved performance in multiuser interference," in *IEEE International Conference on Ultrawideband (ICUWB 2006)*, 2006, pp. 405–410.

[58] R. Blum, R. Kozick, and Sadler, "An adaptive spatial diversity receiver for non-Gaussian interference and noise," *IEEE Trans. Signal Process.*, vol. 47, no. 8, pp. 2100 – 2111, Aug. 1999.

[59] V. Cellini and G. Dona, "A novel joint channel and multi-user interference statistics estimator for UWB-IR based on Gaussian mixture model," in *IEEE International Conference on Ultra-Wideband (ICU)*, 2005, pp. 655–660.

[60] T. Erseghe, "A low-complexity receiver for impulse radio based upon a gaussian mixture interference model," *IEEE Trans. Wireless Commun.*, vol. 7, no. 12, pp. 4867–4876, 2008.

[61] S. de Rivaz, B. Denis, M. Pezzin, and L. Ouvry, "Performance of IEEE 802.15.4a UWB systems under multi-user interference," in *IEEE 18th International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 2007, pp. 1–7.

[62] J. Mitra and L. Lampe, "Robust detectors for TH IR-UWB systems with multiuser interference," in *Ultra-Wideband, 2007. ICUWB 2007. IEEE International Conference on*, September 2007, pp. 745–750.

[63] J. Fiorina, "A simple IR-UWB receiver adapted to multi-user interferences," in *IEEE Global Telecommunications Conference (GLOBECOM)*, 2006, pp. 1–4.

[64] B. S. Kim, J. Bae, I. Song, S. Y. Kim, and H. Kwon, "A comparative analysis of optimum and suboptimum rake receivers in impulsive UWB environment," *IEEE Trans. Veh. Technol.*, vol. 55, no. 6, pp. 1797–1804, 2006.

[65] T. Erseghe and S. Tomasin, "UWB WPAN receiver optimization in the presence of multiuser interference," *Communications, IEEE Transactions on*, vol. 57, no. 8, pp. 2369–2379, 2009.

[66] N. Beaulieu and S. Niranjayan, "New UWB receiver designs based on a gaussian-laplacian noise-plus-MAI model," in *IEEE International Conference on Communications (ICC)*, June 2007, pp. 4128–4133.

[67] S. Niranjayan and N. C. Beaulieu, "Receiver parameter estimation for the simplified gaussian-laplacian multiple access UWB receiver," in *Ultra-Wideband, 2009. ICUWB 2009. IEEE International Conference on*, 2009, pp. 707–712.

[68] ——, "A myriad filter detector for UWB multiuser communication," in *IEEE International Conference on Communications (ICC)*, 2008, pp. 3918–3922.

[69] J. Mitra and L. Lampe, "Comparison of detectors for multiple-access interference mitigation in TH-IR UWB," in *Ultra-Wideband, 2008. ICUWB 2008. IEEE International Conference on*, vol. 1, 2008, pp. 153–156.

[70] ——, "Design and analysis of robust detectors for TH IR-UWB systems with multiuser interference," *Communications, IEEE Transactions on*, vol. 57, no. 8, pp. 2210–2214, 2009.

[71] B. Seyfe and S. Valaee, "A new choice of penalty function for robust multiuser detection based on m-estimation," *Communications, IEEE Transactions on*, vol. 53, no. 2, pp. 224–227, 2005.

[72] B. Hu and N. C. Beaulieu, "On characterizing multiple access interference in TH-UWB systems with impulsive noise models," in *IEEE Radio and Wireless Symposium*, 2008, pp. 879–882.

[73] N. C. Beaulieu and D. J. Young, "Designing time-hopping ultrawide bandwidth receivers for multiuser interference environments," *Proc. IEEE*, vol. 97, no. 2, pp. 255–284, 2009.

[74] E. Gilbert, "Capacity of a burstnoise channel," *Bell Syst. Tech. J*, vol. 39, no. 9, pp. 1253–1265, 1960.

[75] M. Mushkin and I. Bar-David, "Capacity and coding for the gilbert-elliot channels," *IEEE Transactions on Information Theory*, vol. 35, no. 6, pp. 1277–1290, 1989.

[76] Z. Ahmadian and L. Lampe, "Performance of concatenated coded IR-UWB in the presence of multiple-access interference," in *Signals, Systems and Computers, 42nd Asilomar Conference on*, 2008, pp. 2001–2005.

[77] J. Mitra and L. Lampe, "On joint estimation and decoding for channels with noise memory," *IEEE Commun. Lett.*, vol. 13, no. 10, 2009.

[78] ——, "Sensing and suppression of impulsive interference," in *Electrical and Computer Engineering, 2009. CCECE '09. Canadian Conference on*, 2009, pp. 219–224.

[79] W. Turin, "MAP decoding in channels with memory," *IEEE Trans. Commun.*, vol. 48, no. 5, pp. 757–763, May 2000.

[80] J. Mitra and L. Lampe, "Convolutionally coded transmission over markov-gaussian channels: Analysis and decoding metrics," *IEEE Transactions on Communications*, 2010.

[81] A. El Fawal and J.-Y. Le Boudec, "A robust signal detection method for ultra wide band (UWB) networks with uncontrolled interference," *IEEE Trans. Microw. Theory Tech.*, vol. 54, no. 4, pp. 1769–1781, June 2006.

[82] J. Colli-Vignarelli, J. Vernez, R. Merz, C. Dehollain, S. Robert, and J. Y. Le Boudec, "Concurrent transmissions in IR-UWB networks: an experimental validation," in *Ultra-Wideband, 2009. ICUWB 2009. IEEE International Conference on*, 2009, pp. 337–342.

[83] Z. Sahinoglu and I. Guvenc, "Multiuser interference mitigation in noncoherent UWB ranging via nonlinear filtering," *EURASIP Journal on Wireless Communications and Networking*, pp. 1–10, 2006.

[84] D. Dardari, A. Giorgetti, and M. Z. Win, "Time-of-arrival estimation of UWB signals in the presence of narrowband and wideband interference," in *Ultra-Wideband, 2007. ICUWB 2007. IEEE International Conference on*, 2007, pp. 71–76.

[85] H. Zhan, J. Ayadi, J. Farserotu, and J. Y. Le Boudec, "Impulse radio ultra-wideband ranging under multi-user environments," in *Vehicular Technology Conference, 2009. VTC Spring 2009. IEEE 69th*, 2009, pp. 1–5.

[86] E. Ekrem, M. Koca, and H. Delic, "Robust ultra-wideband signal acquisition," *IEEE Trans. Wireless Commun.*, vol. 7, no. 11, pp. 4656–4669, November 2008.

[87] T. Erseghe and A. M. Cipriano, "Performance of UWB impulse radio in strong mai with frequency offsets estimation," in *Ultra-Wideband, 2008. ICUWB 2008. IEEE International Conference on*, vol. 1, 2008, pp. 213–216.

[88] S. Yoon, I. Song, and S. Y. Kim, "Code acquisition for DS/SS communications in non-gaussian impulsive channels," *IEEE Trans. Commun.*, vol. 52, no. 2, pp. 187–190, 2004.

[89] S. Verdù, *Multiuser Detection*. Cambridge University Press, 1998.

[90] Y. C. Yoon and R. Kohno, "Optimum multi-user detection in ultra-wideband (UWB) multiple-access communication systems," in *IEEE International Conference on Communications (ICC)*, no. 1, Apr. 2002, pp. 812–816.

[91] E. Fishler and H. V. Poor, "Low-complexity multiuser detectors for time-hopping impulse-radio systems," *IEEE Trans. Signal Process.*, vol. 52, no. 9, pp. 2561–2571, September 2004.

[92] I. Guvenc and H. Arslan, "A review on multiple access interference cancellation and avoidance for IR-UWB," *Signal Processing*, vol. 87, no. 4, pp. 623–653, April 2007.

[93] R. Merz, "Interference management in impulse-radio ultra-wide band networks," Ph.D. dissertation, Lausanne, 2008. [Online]. Available: http://library.epfl.ch/theses/?nr=4119

[94] A. El Fawal, J.-Y. Le Boudec, R. Merz, B. Radunovic, J. Widmer, and G. M. Maggio, "Tradeoff analysis of PHY-aware MAC in low-rate, low-power UWB networks," *IEEE Commun. Mag.*, vol. 43, no. 12, pp. 147–155, December 2005.

[95] B. Radunovic, J.-Y. Le Boudec, and R. Knopp, "Optimal PHY and MAC Protocols for Wide-Band Ad-Hoc Networks," in *Forty-Fifth Annual Allerton Conference*, 2007.

[96] B. Radunovic and J. Y. Le Boudec, "Optimal power control, scheduling and routing in UWB networks," *IEEE J. Sel. Areas Commun.*, vol. 22, no. 7, pp. 1252–1270, September 2004.

[97] M.-G. Di Benedetto, L. Nardis, M. Junk, and G. Giancola, "$(UWB)^2$: Uncoordinated, wireless, baseborn, medium access control for UWB communication networks," *Mobile Networks and Applications*, vol. 10, no. 5, October 2005.

[98] K. Witrisal, G. Leus, G. Janssen, M. Pausini, F. Troesch, T. Zasowski, and J. Romme, "Noncoherent ultra-wideband systems," *Signal Processing Magazine, IEEE*, vol. 26, no. 4, pp. 48–66, 2009.

[99] S. Mekki, J. L. Danger, B. Miscopein, and J. J. Boutros, "EM channel estimation in a low-cost UWB receiver based on energy detection," in *Wireless Communication Systems. 2008. ISWCS '08. IEEE International Symposium on*, 2008, pp. 214–218.

[100] J. Townsend, G. Weeks, and J. Freebersyser, "Quantifying the covertness of impulse radio," in *Ultra Wideband Conference, Washington D.C.*, September 1999.

[101] A. Bharadwaj and J. K. Townsend, "Evaluation of the covertness of time-hopping impulse radio using a multi-radiometer detection," in *Military Communications Confer-*

ence, 2001. MILCOM 2001. Communications for Network-Centric Operations: Creating the Information Force. IEEE, vol. 1, 2001, pp. 128–134 vol.1.

[102] Y. Souilmi and R. Knopp, "Challenges in UWB signalling for adhoc networking," in *DIMACS Workshop on Signal Processing for Wireless Transmission*, 2002.

[103] ——, "Challenges in UWB signalling for adhoc networking," DIMACS Series in Discrete Mathematics and Theoretical Computer Sciences, 2003.

[104] S. Paquelet and L. M. Aubert, "An energy adaptive demodulation for high data rates with impulse radio," in *Radio and Wireless Conference, 2004 IEEE*, 2004, pp. 323–326.

[105] S. Paquelet, L. M. Aubert, and B. Uguen, "An impulse radio asynchronous transceiver for high data rates," in *IEEE joint Conference on Ultra Wideband Systems and Technologies & International Workshop on Ultra Wideband Systems*, 2004, pp. 1–5.

[106] A. Anttonen, A. Mammela, and A. Kotelba, "Sensitivity of energy detected multilevel pam systems to threshold mismatch," in *Ultra-Wideband, 2008. ICUWB 2008. IEEE International Conference on*, vol. 1, 2008, pp. 165–168.

[107] I. Oppermann, L. Stoica, A. Rabbachin, Z. Shelby, and J. Haapola, "UWB wireless sensor networks: UWEN - a practical example," *IEEE Commun. Mag.*, vol. 42, no. 12, pp. S27–S32, 2004.

[108] A. Rabbachin and I. Oppermann, "Synchronization analysis for UWB systems with a low-complexity energy collection receiver," in *IEEE joint Conference on Ultra Wideband Systems and Technologies & International Workshop on Ultra Wideband Systems*, 2004, pp. 288–292.

[109] S. Dubouloz, A. Rabbachin, S. de Rivaz, B. Denis, and L. Ouvry, "Performance analysis of low complexity solutions for UWB low data rate impulse radio," in *Circuits and Systems, 2006. ISCAS 2006. Proceedings. 2006 IEEE International Symposium on*, May 2006.

[110] Z. Ahmadian and L. Lampe, "Performance analysis of the IEEE 802.15.4a UWB system," *Communications, IEEE Transactions on*, vol. 57, no. 5, pp. 1474–1485, 2009.

[111] A. A. D'Amico, U. Mengali, and L. Taponecco, "Ranging algorithm for the IEEE 802.15.4a standard," in *Ultra-Wideband, 2009. ICUWB 2009. IEEE International Conference on*, 2009, pp. 285–289.

[112] B. Geiger, "Ranging in the IEEE 802.15.4a standard using energy detectors," in *IEEE EUROCON 2009*, May 2009, pp. 1956–1963.

[113] S. Olonbayar, G. Fischer, and R. Kraemer, "Synchronisation performance of wireless sensor networks," in *IEEE International Conference on Ultra-Wideband (ICUWB)*, vol. 2, Sept. 2008, pp. 59–62.

[114] M. E. Sahin, I. Guvenc, and H. Arslan, "Optimization of energy detector receivers for UWB systems," in *IEEE Vehicular Technology Conference (VTC Spring)*, vol. 2, 2005.

[115] M. Nemati, U. Mitra, and R. Scholtz, "Optimum integration time for UWB transmitted reference and energy detector receivers," in *IEEE Military Communications Conference (MILCOM)*, October 2006, pp. 1–7.

[116] Z. Tian and B. Sadler, "Weighted energy detection of ultra-wideband signals," in *Signal Processing Advances in Wireless Communications, 2005 IEEE 6th Workshop on*, June 2005, pp. 10 168–1072.

[117] M. Dazhong and Q. Zhengding, "Weighted non-coherent energy detection receiver for UWB OOK systems," in *Signal Processing, 2008. ICSP 2008. 9th International Conference on*, 2008, pp. 1846–1849.

[118] T. Zasowski, F. Troesch, and A. Wittneben, "Partial channel state information and intersymbol interference in low complexity UWB PPM detection," in *IEEE International Conference on Ultrawideband (ICUWB 2006)*, September 2006, pp. 369 – 374.

[119] A. Rabbachin, I. Oppermann, and B. Denis, "ML time-of-arrival estimation based on low complexity UWB energy detection," in *Ultra-Wideband, The 2006 IEEE 2006 International Conference on*, 2006, pp. 599–604.

[120] ——, "GML ToA estimation based on low complexity UWB energy detection," in *Personal, Indoor and Mobile Radio Communications, 2006 IEEE 17th International Symposium on*, 2006, pp. 1–5.

[121] L. Stoica, A. Rabbachin, H. O. Repo, T. S. Tiuraniemi, and I. Oppermann, "An ultrawideband system architecture for tag based wireless sensor networks," *Vehicular Technology, IEEE Transactions on*, vol. 54, no. 5, pp. 1632–1645, 2005.

[122] L. Stoica, A. Rabbachin, and I. Oppermann, "A low-complexity noncoherent ir-uwb transceiver architecture with toa estimation," *Microwave Theory and Techniques, IEEE Transactions on*, vol. 54, no. 4, pp. 1637–1646, 2006.

[123] F. S. Lee and A. P. Chandrakasan, "A 2.5nj/b 0.65v 3-to-5ghz subbanded UWB receiver in 90nm CMOS," in *IEEE International Solid-State Circuits Conference (ISSCC 07)*, February 2007.

[124] C. Duan, P. Orlik, Z. Sahinoglu, and A. F. Molisch, "A non-coherent 802.15.4a UWB impulse radio," in *Ultra-Wideband, 2007. ICUWB 2007. IEEE International Conference on*, September 2007, pp. 146–151.

[125] A. Rabbachin, T. Q. S. Quek, I. Oppermann, and M. Z. Win, "Effect of uncoordinated network interference on UWB energy detection receiver," in *Signal Processing Advances in Wireless Communications, 2009. SPAWC '09. IEEE 10th Workshop on*, 2009, pp. 692–696.

[126] M. Hauske, H. Jaekel, H. U. Dehner, and F. K. Jondral, "Interference mitigation for energy detection in a multiband impulse radio UWB system," in *Vehicular Technology Conference, 2008. VTC 2008-Fall. IEEE 68th*, 2008, pp. 1–5.

[127] X. Peng, F. Chin, S. H. Wong, K. Y. Sam, and L. Zhongding, "A rake combining scheme for an energy detection based noncoherent OOK receiver in UWB impulse radio systems," in *IEEE International Conference on Ultrawideband (ICUWB 2006)*, September 2006, pp. 73–78.

[128] L. Jian Xing, F. Chin, K. Yuen Sam, W. Sai Ho, and N. Liming, "UWB piconet interference suppression using clustered ternary orthogonal signaling scheme," in *Ultra-Wideband, 2009. ICUWB 2009. IEEE International Conference on*, 2009, pp. 83–87.

[129] J. Haapola, A. Rabbachin, L. Goratti, C. Pomalaza-Raez, and I. Oppermann, "Effect of impulse radio ultrawideband based on energy collection on MAC protocol performance," *Vehicular Technology, IEEE Transactions on*, vol. 58, no. 8, pp. 4491–4506, 2009.

[130] J. Rousselot and J.-D. Decotignie, "A high-precision ultra wideband impulse radio physical layer model for network simulation," in *OMNeT++ 2009: Proceedings of the 2nd International Workshop on OMNeT++ (hosted by SIMUTools 2009)*, 2009.

[131] J. G. Proakis, *Digital Communications*, 4th ed. New York, NY: McGraw–Hill, 2001.

[132] C.-C. Chui and R. A. Scholtz, "Optimizing tracking loops for UWB monocycles," in *Global Telecommunications Conference, 2003. GLOBECOM '03. IEEE*, vol. 1, 2003, pp. 425–430 Vol.1.

[133] S. Farahmand, X. Luo, and G. B. Giannakis, "Demodulation and tracking with dirty templates for UWB impulse radio: algorithms and performance," *IEEE Trans. Veh. Technol.*, vol. 54, no. 5, 2005.

[134] B. Zhen, H.-B. Li, and R. Kohno, "Clock offset compensation in ultra-wideband ranging," *IEICE Trans. Fundam. Electron. Commun. Comput. Sci.*, no. 11, pp. 3082–3088, 2006.

[135] L. Huang, El, O. Rousseaux, and B. Gyselinckx, "Timing tracking algorithms for impulse radio (IR) based ultra wideband (UWB) systems," in *Wireless Communications, Networking and Mobile Computing, 2007. WiCom 2007. International Conference on*, 2007, pp. 570–573.

[136] A. Fort, M. Chen, C. Desset, P. Wambacq, and L. Van Biesen, "Clock offset tracking for subsampling UWB architectures in a body area network," in *Ultra-Wideband, 2007. ICUWB 2007. IEEE International Conference on*, 2007, pp. 229–234.

[137] T. Erseghe and F. Renna, "On schmidl-cox-like frequency estimation applied to UWB impulse radio systems," in *Ultra-Wideband, 2009. ICUWB 2009. IEEE International Conference on*, 2009, pp. 693–697.

[138] A. Wellig and Y. Qiu, "Trellis-based maximum-likelihood crystal drift estimator for ranging applications in uwb-ldr," in *IEEE International Conference on Ultrawideband (ICUWB 2006)*, 2006.

[139] D. R. McKinstry and R. M. Buehrer, "Issues in the performance and covertness of UWB communications systems," in *Circuits and Systems, 2002. MWSCAS-2002. The 2002 45th Midwest Symposium on*, vol. 3, 2002, pp. III–601–4 vol.3.

[140] W. Tao, W. Yong, and C. Kangsheng, "Improving the jam resistance performance of UWB impulse radio independently of time hopping codes," in *Vehicular Technology Conference, 2004. VTC 2004-Spring. 2004 IEEE 59th*, vol. 3, 2004, pp. 1475–1479 Vol.3.

[141] J. P. Lie, S. Chong Meng, and N. Boon Poh, "UWB ranging with high robustness against dominant jammer and multipath," *Microwave and Wireless Components Letters, IEEE*, vol. 15, no. 12, pp. 907–909, 2005.

[142] D. S. Ha and P. R. Schaumont, "Replacing cryptography with ultra wideband (uwb) modulation in secure RFID," in *RFID, 2007. IEEE International Conference on*, 2007, pp. 23–29.

[143] Y. Zhang and H. Dai, "A real orthogonal space-time coded UWB scheme for wireless secure communications," *EURASIP Journal on Wireless Communications and Networking*, vol. vol. 2009, no. Article ID 571903, 2009.

[144] R. Wilson, D. Tse, and R. A. Scholtz, "Channel identification: Secret sharing using reciprocity in ultrawideband channels," in *Ultra-Wideband, 2007. ICUWB 2007. IEEE International Conference on*, 2007, pp. 270–275.

[145] M. G. Madiseh, M. L. McGuire, S. S. Neville, C. Lin, and M. Horie, "Secret key generation and agreement in UWB communication channels," in *Global Telecommunications Conference, 2008. IEEE GLOBECOM 2008. IEEE*, 2008, pp. 1–5.

[146] M. G. Madiseh, H. Shuai, M. L. McGuire, S. W. Neville, and D. Xiaodai, "Verification of secret key generation from UWB channel observations," in *IEEE International Conference on Communications (ICC)*, 2009, pp. 1–5.

[147] S. Brands and D. Chaum, "Distance-bounding protocols," in *EUROCRYPT'93, Lecture Notes in Computer Science 765*, 1993.

[148] S. Čapkun, L. Buttyán, and J.-P. Hubaux, "SECTOR: secure tracking of node encounters in multi-hop wireless networks," in *ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN)*, 2003.

[149] L. Bussard, "Trust establishment protocols for communicating devices," Ph.D. dissertation, 2004.

[150] G. Hancke and M. Kuhn, "An RFID distance bounding protocol," in *Security and Privacy for Emerging Areas in Communications Networks, 2005. SecureComm 2005. First International Conference on*, 2005.

[151] J. Reid, J. M. G. Nieto, T. Tang, and B. Senadji, "Detecting relay attacks with timing-based protocols," in *ASIAN ACM Symposium on Information, Computer and Communications Security*, 2007.

[152] S. Čapkun and J. Hubaux, "Secure positioning in wireless networks," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 2, 2006.

[153] C. Meadows, R. Poovendran, D. Pavlovic, L.-W. Chang, and P. Syverson, "Distance bounding protocols: Authentication logic analysis and collusion attacks," in *Secure Localization and Time Synchronization for Wireless Sensor and Ad Hoc Networks*. Springer-Verlag, Series: Advances in Information Security, Vol. 30, 2007.

[154] D. Singelée and B. Preneel, "Distance bounding in noisy environments," in *European Workshop on Security and Privacy in Ad-Hoc and Sensor Networks (ESAS)*, 2007.

[155] Y.-J. Tu and S. Piramuthu, "RFID distance bounding protocols," in *First International EURASIP Workshop on RFID Technology*, 2007.

[156] K. B. Rasmussen and S. Čapkun, "Location privacy of distance bounding protocols," in *15th ACM conference on Computer and Communications Security (CCS)*, 2008.

[157] J. Munilla and A. Peinado, "Distance bounding protocols for RFID enhanced by using void-challenges and analysis in noisy channels," *Wireless Communications and Mobile Computing*, vol. 8, no. 9, 2008.

[158] C. H. Kim, G. Avoine, F. Koeune, F.-X. Standaert, and O. Pereira, "The swiss-knife RFID distance bounding protocol," in *International Conference on Information Security and Cryptology (ICISC)*, P. Lee and J. Cheon, Eds., 2008.

[159] C. H. Kim and G. Avoine, "RFID distance bounding protocol with mixed challenges to prevent relay attacks," in *8th International Conference on Cryptology and Network Security (CANS)*, 2009, paper http://www.uclouvain.be/sites/security/download/papers/KimA-2009-cans.pdf.

[160] N. O. Tippenhauer and S. Čapkun, "Id-based secure distance bounding and localization," in *In Proceedings of ESORICS (European Symposium on Research in Computer Security)*, 2009.

[161] Y. Zhang, W. Liu, Y. Fang, and D. Wu, "Secure localization and authentication in ultra-wideband sensor networks," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 4, 2006.

[162] J. Clulow, G. P. Hancke, M. G. Kuhn, and T. Moore, "So near and yet so far: Distance-bounding attacks in wireless networks," in *Third European Workshop on Security and Privacy in Ad-hoc and Sensor Networks (ESAS)*, 2006. [Online]. Available: http://www.crysys.hu/ESAS2006/cfp.html

[163] G. P. Hancke and M. G. Kuhn, "Attacks on time-of-flight distance bounding channels," in *First ACM conference on Wireless network security (WiSec)*, 2008.

[164] Z. Sahinoglu and S. Gezici, "Ranging in the ieee 802.15.4a standard," in *Wireless and Microwave Technology Conference, 2006. WAMICON '06. IEEE Annual*, 2006, pp. 1–5.

[165] D. Dardari, A. Conti, U. Ferner, A. Giorgetti, and M. Z. Win, "Ranging with ultrawide bandwidth signals in multipath environments," *Proc. IEEE*, vol. 97, no. 2, pp. 404–426, 2009.

[166] K. Yu and I. Oppermann, "Timing acquisition for ir-uwb systems," in *Signal Processing and Its Applications, 2005. Proceedings of the Eighth International Symposium on*, vol. 1, 2005, pp. 287–290.

[167] I. Guvenc, Z. Sahinoglu, P. Orlik, and H. Arslan, "Searchback algorithms for TOA estimation in non-coherent low-rate IR-UWB systems," *Wireless Personal Communications*, vol. 48, no. 4, pp. 585–603, 2009.

[168] N. C. Beaulieu and C. C. Tan, "An FFT method for generating bandlimited Gaussian noise variates," in *Global Telecommunications Conference, 1997. GLOBECOM '97., IEEE*, vol. 2, 1997, pp. 684–688.

[169] P. J. Rousseeuw and J. Bassett, Gilbert W., "The remedian: A robust averaging method for large data sets," *Journal of the American Statistical Association*, vol. 85, no. 409, pp. 97–104, 1990.

[170] J. Ibrahim and R. M. Buehrer, "Two-stage acquisition for UWB in dense multipath," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 4, pp. 801–807, 2006.

[171] D. V. Sarwate and M. B. Pursley, "Crosscorrelation properties of pseudorandom and related sequences," *Proc. IEEE*, vol. 68, no. 5, pp. 593–619, 1980.

[172] A. Dempster, N. Laird, and D. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the Royal Statistical Society*, vol. 39, no. 1, pp. 1–38, 1977.

[173] J. Bilmes, "A gentle tutorial on the EM algorithm and its application to parameter estimation for Gaussian mixture and hidden Markov models," 1997. [Online]. Available: citeseer.ist.psu.edu/bilmes98gentle.html

[174] J. Fessler and A. Hero, "Space-alternating generalized expectation-maximization algorithm," *IEEE Trans. Signal Process.*, vol. 42, no. 10, pp. 2664–2667, Oct. 1994.

[175] L. E. Baum, T. Petrie, G. Souled, and N. Weiss, "A maximization technique occurring in statistical analysis of probabilistic functions of Markov chains," *Ann. Math. Stat.*, vol. 41, no. 1, pp. 164–171, 1970.

[176] G. D. Forney, Jr., "The Viterbi algorithm," *Proc. IEEE*, vol. 61, pp. 268–278, 1973.

[177] Z. Liu, J. Almhana, V. Choulakian, and R. McGorman, "Recursive EM algorithm for finite mixture models with application to internet traffic modeling," in *Communication Networks and Services Research, 2004. Proceedings. Second Annual Conference on*, May 2004, pp. 198–207.

[178] J. Radon, "Über die Bestimmung von Funktionen durch ihre Integralwerte längs gewisser Mannigfaltigkeiten," *Ber. Sächs. Akad. der Wissenschaften, Leipzig, Math.-Phys. Klasse*, vol. 69, pp. 262–267, 1917.

[179] M. van Ginkel, C. L. Hendriks, and L. van Vliet, "A short introduction to the Radon and Hough transforms and how they relate to each other," Quantitative Imaging Group, Delft University of Technology, Tech. Rep. QI-2004-01, January 2004.

[180] R. O. Duda and P. E. Hart, "Use of the Hough transformation to detect lines and curves in pictures," *Commun. ACM*, vol. 15, no. 1, 1972.

[181] S. Gezici, Z. Tian, G. Giannakis, H. Kobayashi, A. Molisch, H. V. Poor, and Z. Sahinoglu, "Localization via ultra-wideband radios: a look at positioning aspects for future sensor networks," *IEEE Signal Process. Mag.*, vol. 22, no. 4, pp. 70–84, July 2005.

[182] R. Mulloy, "Ultrawide-band RFID technology," FCC Radio Frequency Identification Workshop, 2004.

[183] Z. Kfir and A. Wool, "Picking virtual pockets using relay attacks on contactless smartcard," in *Security and Privacy for Emerging Areas in Communications Networks, 2005. SecureComm 2005. First International Conference on*, 2005.

[184] P. Papadimitratos, M. Poturalski, P. Schaller, P. Lafourcade, D. Basin, S. Čapkun, and J.-P. Hubaux, "Secure neighborhood discovery: A fundamental element for mobile ad hoc networking," *IEEE Communications Magazine*, vol. Vol.46, No.2, 2008.

[185] K. B. Rasmussen, S. Čapkun, and M. Cagalj, "SecNav: secure broadcast localization and time synchronization in wireless networks," in *13th Annual ACM International Conference on Mobile Computing and Networking (MobiCom)*. ACM, 2007.

[186] J. Ryckaert, G. Van der Plas, V. De Heyn, C. Desset, G. Vanwijnsberghe, B. Van Poucke, and J. Craninckx, "A 0.65-to-1.4nJ/burst 3-to-10GHz UWB digital TX in 90nm CMOS for IEEE 802.15.4a," *IEEE International Solid-State Circuits Conference (ISSCC 07)*, 2007.

[187] "http://www.realtechnews.com/posts/1366."

[188] P. Driessen and L. Greenstein, "Modulation techniques for high-speed wireless indoor systems using narrowbeam antennas," *Communications, IEEE Transactions on*, vol. 43, no. 10, pp. 2605 –2612, oct 1995.

[189] F. Trösch and A. Wittneben, "MLSE post-detection for ISI mitigation and synchronization in UWB low complexity receivers," in *Vehicular Technology Conference, 2007. VTC2007-Spring. IEEE 65th*, april 2007, pp. 2915 –2919. [Online]. Available: http://www.nari.ee.ethz.ch/wireless/pubs/p/vtc_2007_spring

[190] C. Cho, Z. Honggang, and M. Nakagawa, "A UWB repeater with a short relaying-delay for range extension," in *Wireless Communications and Networking Conference, 2004. WCNC. 2004 IEEE*, vol. 2, 2004, pp. 1154–1158 Vol.2.