

Ultra Low Power Communication Protocols for UWB Impulse Radio Wireless Sensor Networks

THÈSE N° 4720 (2010)

PRÉSENTÉE LE 3 JUILLET 2010

À LA FACULTÉ INFORMATIQUE ET COMMUNICATIONS

LABORATOIRE DE MODÉLISATION SYSTÉMIQUE

PROGRAMME DOCTORAL EN INFORMATIQUE, COMMUNICATIONS ET INFORMATION

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Jérôme ROUSSELOT

acceptée sur proposition du jury:

Prof. A. Schiper, président du jury
Prof. J.-D. Decotignie, directeur de thèse
Prof. J.-C. Grégoire, rapporteur
Prof. J.-Y. Le Boudec, rapporteur
Dr O. Rousseaux, rapporteur



ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Suisse
2010

Abstract

This thesis evaluates the potential of Ultra Wideband Impulse Radio for wireless sensor network applications.

Wireless sensor networks are collections of small electronic devices composed of one or more sensors to acquire information on their environment, an energy source (typically a battery), a microcontroller to control the measurements, process the information and communicate with its peers, and a radio transceiver to enable these communications. They are used to regularly collect information within their deployment area, often for very long periods of time (up to several years). The large number of devices often considered, as well as the long deployment durations, make any manual intervention complex and costly. Therefore, these networks must self-configure, and automatically adapt to changes in their electromagnetic environment (channel variations, interferers) and network topology modifications: some nodes may run out of energy, or suffer from a hardware failure.

Ultra Wideband Impulse Radio is a novel wireless technology that, thanks to its extremely large bandwidth, is more robust to frequency dependent propagation effects. Its impulsive nature makes it robust to multipath fading, as the short duration of the pulses leads most multipath components to arrive isolated. This technology should also enable high precision ranging through time of flight measurements, and operate at ultra low power levels.

The main challenge is to design a system that reaches the same or higher degree of energy savings as existing narrowband systems considering all the protocol layers.

As these radios are not yet widely available, the first part of this thesis presents Maximum Pulse Amplitude Estimation, a novel approach to symbol-level modeling of UWB-IR systems that enabled us to implement the first network simulator of devices compatible with the UWB physical layer of the IEEE 802.15.4A standard for wireless sensor networks.

In the second part of this thesis, WideMac, a novel ultra low power MAC protocol specifically designed for UWB-IR devices is presented. It uses asynchronous duty cycling of the radio transceiver to minimize the power consumption, combined with periodic beacon emissions so that devices can learn each other's wake-up patterns and exchange packets. After an analytical study of the protocol, the network simulation tool presented in the first part of the thesis is used to evaluate the performance of WideMac in a medical body area network application. It is compared to two narrowband and an FM-UWB solutions. The protocol stack parameters are optimized for each solution, and it is observed that WideMac combined to UWB-IR is a credible technology for such applications. Similar simulations, considering this time a static multi-

hop network are performed. It is found that WideMac and UWB-IR perform as well as a mature and highly optimized narrowband solution (based on the WiseMAC ULP MAC protocol), despite the lack of clear channel assessment functionality on the UWB radio.

The last part of this thesis studies analytically a dual mode MAC protocol named WideMac-High Availability. It combines the Ultra Low Power WideMac with the higher performance Aloha protocol, so that ultra low power consumption and hence long deployment times can be combined with high performance low latency communications when required by the application. The potential of this scheme is quantified, and it is proposed to adapt it to narrowband radio transceivers by combining WiseMAC and CSMA under the name WiseMAC-HA.

Keywords

Ultra wideband, UWB, Impulse Radio, multiple access interference, IEEE 802.15.4, IEEE 802.15.4A, omnet++, MiXiM, wireless sensor networks, medium access control, ultra low power, WideMac, maximum pulse amplitude estimation, high availability, body area network, discrete event simulation.

Résumé

Cette thèse a pour objet l'évaluation du potentiel des radios impulsionnelles à très large bande pour des applications de réseaux de capteurs.

Les réseaux de capteurs sont des ensembles d'équipements électroniques de taille réduite, composés d'un ou de plusieurs capteurs leur permettant l'acquisition d'information sur leur milieu, une source d'énergie (souvent une batterie), un microcontrôleur dédié au contrôle de mesure, au traitement de l'information et à la communication avec ses pairs, et un émetteur-récepteur radio permettant ces communications. Ces systèmes sont utilisés pour la collecte régulière d'information dans leur zone de déploiement, souvent durant de longues périodes, pouvant aller jusqu'à plusieurs années. Le grand nombre d'équipements considéré, ainsi que les longues durées de fonctionnement, rendent toute intervention manuelle complexe et coûteuse. C'est pourquoi ces réseaux doivent être auto-configurables, et s'adapter de manière entièrement autonome à des modifications de leur environnement (évolution du canal de communication, apparition et disparition d'interféreurs) et aux évolutions de la topologie du réseau: certains noeuds peuvent épuiser leur réserve d'énergie, ou subir une avarie matérielle.

Les radio impulsionnelles à ultra large bande (RI-ULB) sont une nouvelle technologie sans fil qui, en raison de son importante largeur de bande, est plus robuste aux effets de propagation dépendant de la fréquence. Leur nature impulsionnelle les rend robustes à l'évanouissement multi chemin, car la durée des impulsions, très courte en comparaison des différences de temps d'arrivée, permet d'éviter leur superposition au récepteur. Cette technologie devrait aussi permettre la localisation à précision élevée grâce à la mesure des temps de propagation, et fonctionner à des niveaux de puissance extrêmement faibles.

Le défi principal consiste à concevoir un système qui atteigne un niveau de gestion de l'énergie comparable à ou meilleur que celui des systèmes à bande étroite existants, ceci en considérant toutes les couches de protocoles.

Comme ces radios ne sont pas encore disponibles, la première partie de cette thèse présente une nouvelle approche de modélisation au niveau symbole des systèmes RI-ULB afin de permettre la réalisation du premier simulateur réseau d'équipements compatibles avec la couche physique ULB du standard IEEE 802.15.4A pour réseaux de capteurs.

Dans la deuxième partie de cette thèse, WideMac, un nouveau protocole de contrôle d'accès à la couche physique, à très basse consommation, est présenté. Il met en oeuvre une mise en veille asynchrone du récepteur-émetteur radio pour réduire la consommation, combinée à une annonce de présence périodique permettant aux noeuds du réseau de découvrir leurs moments de réveil respectifs et d'échanger des paquets. Après une étude analytique du protocole, l'outil

de simulation réseau présenté dans la première partie de cette thèse est utilisé pour l'évaluation de la performance de WideMac dans le cas d'un réseau corporel médical. La solution proposée est comparée à deux solutions à bande étroite, ainsi qu'à une solution à bande ultra large à modulation de fréquence. Les paramètres de la pile de protocoles sont optimisés pour chaque solution, et on observe que WideMac combiné à RI-ULB est une technologie compétitive pour ce type d'applications.

Des simulations similaires, ayant cette fois pour objet un réseau statique à sauts multiples, sont effectuées. On observe que WideMac et RI-ULB fonctionnent aussi bien qu'une solution à bande étroite éprouvée et hautement optimisée (basée sur le protocole à très basse consommation WiseMAC), malgré l'absence de capacité d'évaluation de l'état du canal des émetteurs-récepteurs.

La dernière partie de cette thèse étudie un protocole bi-mode nommé WideMac - Haute Disponibilité. Celui-ci combine le protocole à très basse consommation WideMac avec le protocole à haute performance Aloha, afin de combiner d'une part une durée de vie très longue et d'autre part des communications à haut débit et faibles délais lorsque l'application le requiert. Le potentiel offert par ce protocole est quantifié avec des modèles analytiques, et il est proposé de l'adapter à des radios à bande étroite en combinant WiseMAC et CSMA sous le nom WiseMAC-HA.

Mots clés

Ultra large bande, ULB, radio impulsive, interférence à accès multiple, IEEE 802.15.4, IEEE 802.15.4A, omnet++, mixim, réseaux de capteurs, contrôle d'accès au support physique, très basse consommation, WideMac, Estimation de l'amplitude maximale d'impulsion, disponibilité élevée, réseau corporel sans fil, simulation à événements finis.

Acknowledgments

First of all, I would like to thank my thesis advisor Prof. Jean-Dominique Decotignie for offering me the opportunity to work on this thesis in his research group, and for his foresight, advice and continuous support over the past four years. I would also like to thank all my colleagues at CSEM for their friendship and technical guidance, and among them Amre El-Hoiydi who always found the time to answer my numerous questions at the beginning of this thesis.

I greatly appreciated the support of my lab colleagues: in topological order, Damien Piguet, Bernard Perrin, Laurent von Allmen, Philippe Dallemagne, Corinne Kassapoglou-Faist, Martin Sénéclauze and Antonio Restrepo Zea. I was lucky enough to be part of the CSEM FM-UWB team with John Gerrits and John Farserotu. Our standardization effort at the IEEE 802.15.6 working group was a great and successful experience. My work involved quite a lot of traveling, and I appreciated the fast and efficient help provided by our secretaries Michèle Siggen, Raquel Desaulles and Jana Neuberger. Thank you for making my professional life so much easier ! Due to the high degree of technical expertise and great cooperative atmosphere at CSEM, I unfortunately received help from too many colleagues to be able to cite them all.

I am grateful to Prof. Jean-Yves Le Boudec, Prof. Jean-Charles Grégoire and Dr. Olivier Rousseaux for having accepted to be part of my jury, and to Prof. André Schiper for presiding the jury.

The work presented in this thesis was supported (in part) by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), and I would like to thank all the fellow PhD students whom I had the pleasure to meet through this research project, especially those part of the MICS UWB cluster. I am redeemable to Hai Zhan for teaching me to count in Chinese while waiting for an airplane, to Manuel Flury for helping me in the validation of my network simulation models and to Gabriela Quintero Díaz de León for introducing me to electromagnetic waves and for her enduring friendship. Meeting James Colli-Vignarelli, Prakash Thoppay Egambaram and Ruben Merz was a pleasure. Special thanks go to Prof. Le Boudec for making possible and administering the whole MICS UWB cluster. More generally, the NCCR-MICS was an exceptional place to network with fellow PhD students from all over the world. Among them: Ali Salehi, Karl Baumgartner and Vojislav Gajic.

I would also like to thank all my friends in Lausanne: Lukas Lechner and Andrea Warmuth for all the great times in Switzerland, Belgium and Austria, Joachim Wille for the regular meetings in the world's unlikeliest places, Maud Helmstetter, Agnès Tissot, Carlos Martin, and all the members of the Mikami Judo Club Lausanne. They made my life in Lausanne a pleasurable experience.

I am also grateful to my friends from Belgium, with whom I kept in touch despite the distance. Among them: Jasper van de Beek, Antoine Hauzeur, Monica Sias, Céline Loncheval, David Glinoyer, Yves Morteahan, Perrine Willems, Marc Jaumain, Fabian Nisen, Christophe van de Walle, Philippe Némery and Jonathan Gross.

Last but not least, I want to thank my family, and in particular my parents, my sister and my grandmother for their continuous understanding, encouragement and support.

Jérôme Rousselot
Neuchâtel, June 2010.

Contents

Abstract	iii
Résumé	v
Acknowledgments	vii
Contents	ix
1 Introduction	1
1.1 Wireless Sensor Networks	1
1.2 Ultra Wideband Impulse Radio Wireless Communications . . .	2
1.3 Problem Statement	4
1.4 Contributions	4
1.5 Thesis Organization	5
2 An UWB-IR PHY for Sensor Networks	7
2.1 Ultra Wideband Communications	9
2.2 UWB-IR Modulations	14
2.3 The UWB-IR Channel	16
2.3.1 Free Space Propagation	16
2.3.2 The Ghassemzadeh Channel Model	17
2.3.3 The Saleh-Valenzuela Channel Model	18
2.3.4 The IEEE 802.15.4A Channel Model	20
2.4 Receiver Architectures	22
2.4.1 Coherent receivers	22
2.4.2 Non-coherent receivers	23
2.5 The IEEE 802.15.4A Standard	23
2.5.1 Burst Position Modulation	24
2.5.2 Frame Structure	25
2.5.3 Error Correction	25
2.6 Modeling Multiple Access Interference	28
2.6.1 Overview	28
2.6.2 MATLAB Simulations	31
2.6.3 The Gaussian Approximation	31
2.6.4 Characteristic Function	33
2.6.5 The Pulse Collision Model	33
2.6.6 Large Deviations and Importance Sampling	34
2.6.7 Cumulative Noise	35
2.6.8 CTU	35

2.6.9	Observations on Existing Models	35
2.7	Maximum Pulse Amplitude Estimation	37
2.7.1	Overview	37
2.7.2	Assumptions	38
2.7.3	Software Architecture	40
2.7.4	Data Structures	41
2.7.5	Physical Layer	41
2.7.6	Channel	42
2.7.7	Radio	43
2.8	A Model of the IEEE 802.15.4A UWB PHY	43
2.8.1	Assumptions	44
2.8.2	Transmitter	45
2.8.3	Synchronization	45
2.8.4	Demodulation	46
2.9	Validation of the Simulation Model	47
2.10	Evaluation of the IEEE 802.15.4A UWB-IR PHY	50
2.10.1	Channel Models	53
2.10.2	Multiple Access Interference	53
2.10.3	Simulation Performance	56
2.11	Observations	56
3	WideMac: An ULP MAC Protocol for UWB-IR	59
3.1	Wireless MAC Protocols	59
3.1.1	Objectives	60
3.1.2	Problems	62
3.1.3	Existing Approaches	63
3.2	Ultra Low Power MAC Protocols	68
3.2.1	Types of Energy Waste and the Ideal Protocol	70
3.2.2	ULP Scheduled Protocols	71
3.2.3	ULP Random Access Protocols	72
3.3	UWB-IR MAC Protocols	75
3.3.1	Adaptive Communication Parameters	76
3.3.2	Pulse Sensing and TDMA	78
3.3.3	IEEE 802.15.4A MAC Features	78
3.3.4	Remarks	78
3.4	WideMac	79
3.4.1	WideMac Design Rationale	80
3.4.2	WideMac Description	80
3.4.3	WideMac Beacon Collisions	82
3.4.4	WideMac Backoff Algorithm	91
3.5	Store and Forward Analytical Power Consumption Analysis	92
3.5.1	Models Objectives and Assumptions	92
3.5.2	Power Consumption Models	94
3.5.3	Results Analysis	100
3.6	Star-Topology Analytical Latency Analysis	104
3.6.1	Assumptions and Latency Models	104
3.6.2	Latency Models Validation	106
3.6.3	Latency Results	109
3.7	Observations	110

4	WideMac in a Star Topology Medical Body Area Network	115
4.1	Medical Body Area Networks	115
4.2	Compared Technologies	117
4.3	Methodology	119
4.3.1	Overview	119
4.3.2	Parameters	120
4.3.3	Metrics and Calibration	121
4.4	Results without coexistence	124
4.4.1	Packet transmission success rate	124
4.4.2	Latency	129
4.4.3	Power consumption	133
4.4.4	Summary	139
4.5	Results including coexistence	143
4.5.1	Packet Success Rate	143
4.5.2	Latency	147
4.5.3	Power Consumption	150
4.5.4	Summary	156
4.6	Observations	158
5	WideMac in a static multihop wireless sensor network	159
5.1	Routing in wireless sensor networks	159
5.1.1	Mobile Ad Hoc Networks Routing Basics	162
5.1.2	Flat topology reactive protocols	163
5.1.3	Flat topology proactive protocols	164
5.1.4	Hierarchical topology protocols	165
5.1.5	Location-based protocols	167
5.1.6	Network-level broadcast protocols	170
5.1.7	UWB Routing	171
5.1.8	Routing Metrics for WSN	172
5.1.9	WiseRoute	173
5.2	Methodology	174
5.3	Results	177
5.3.1	Packet Error Rate and Connectivity	177
5.3.2	Latency and average number of hops	180
5.3.3	Power Consumption	183
5.4	Observations	186
6	WideMac-High Availability	187
6.1	Applications for reliable wireless communications	188
6.2	Wireless MAC protocol design trade-offs	191
6.2.1	Other wireless communication systems	191
6.2.2	Multichannel MAC protocols	194
6.3	The WideMac-HA Protocol	196
6.3.1	High Availability interoperable ALOHA based mode	196
6.3.2	Multiple Channels	198
6.4	Mathematical Models	202
6.4.1	Power consumption in a one-hop star topology network	202
6.4.2	Latency	203
6.5	Evaluation	204
6.5.1	Parameters	204

6.5.2	Power Consumption	204
6.5.3	Latency	210
6.6	Observations	213
7	Conclusion	215
	Bibliography	219
	Publications	235
	Curriculum Vitae	237

Introduction

This thesis evaluates the potential of Ultra Wideband Impulse Radio, a novel short range wireless communication technology, for wireless sensor network applications. This chapter introduces the reader to the concept of wireless sensor networks in section 1.1, which is followed by a definition of ultra wideband communications and UWB-IR in section 1.2. Section 1.3 describes the problems addressed in this thesis. Section 1.4 enumerates the scientific results obtained, and section 1.5 presents the thesis organization.

1.1 Wireless Sensor Networks


The continuing trend of miniaturization in microelectronics, combined to advances in the design of wireless communication electronics has led to the emergence of a new class of cheap, durable and very specialized computing devices [6] equipped with one or more sensors. When deployed and activated, these systems form wireless sensor networks (WSN) that offer previously unknown of spatio-temporal resolution of monitored data. Figure 1.1 illustrates two WSN platforms.

They are used to evaluate air pollution, to detect and prevent forest fires, floodings and rock falls, and to monitor the structure of buildings and bridges. They can also be mobile when deployed on persons such as firefighters and hospitalized patients, or on buoys for sea monitoring, or even on animals. They constitute an essential part of the ubiquitous computing vision which predicts seamless interactions between humans and computers, enabled by a deep integration of computing devices in the objects and the environment surrounding us. Figure 1.2 illustrates various WSN applications.

Currently, wireless sensor networks are limited by their radio technology. The narrowband radio transceivers in use have been optimized as far as possible and no significant progress is expected regarding power consumption with current technology. Wireless communications are sensitive to interference, caused either by other sensor nodes or by other radio technologies such as WiFi, WiMax, microwave ovens, etc. Ultra Wideband may be a way to overcome these limitations and enlarge the application domains of wireless sensor networks.



(a) CSEM WiseNode XE1203.



(b) CSEM WiseNode XE1203 System on Chip solution.

Figure 1.1: Wireless sensor networking platforms.

1.2 Ultra Wideband Impulse Radio Wireless Communications

Wireless communications operate by generating an electromagnetic wave [156]. Information is encoded on this wave (a process named *modulation*) by modifying some physical characteristics of the wave. The modulation technique easiest to understand consists of simply acting on the presence and absence of the wave (this technique is named *On-Off Keying*). The signal detection at the receiver leads to the demodulation of a 1 value, and conversely the temporary absence of the signal leads to the demodulation of a 0.

An electromagnetic wave is characterized by its frequency f , which is related to its wavelength λ and the speed of light at which it travels in the propagation environment by $f\lambda = c$. The wavelength λ is the physical distance over which the wave amplitude oscillates over all its possible values and f is the number of oscillations per second of the field. In practice, a communication signal can be decomposed as the sum of several sinusoidal waves whose frequencies are close to each other. This allows to define the signal bandwidth B , which is the difference between the maximum and the minimum frequencies of these sinusoidal waves (in practice, a condition on the power level of these waves is also used, so that for instance 95% of the signal's energy lies inside the signal bandwidth). The amount of digital information that can be sent per unit of time is called the channel capacity (in bits per second). Its maximum theoretical value is given by the Shannon-Hartley theorem [139]:

$$C = B \log_2 \left(1 + \frac{Signal}{Noise} \right), \quad (1.1)$$

where *Signal* is the received signal strength (equal to the energy radiated at the source, attenuated by the channel) and *Noise* is the noise power (additive white Gaussian noise).



Figure 1.2: Wireless sensor network applications.

This leads to the problem of spectrum allocation: to avoid interference between transmitters (which would have an effect similar to increasing the *Noise* factor in eq. 1.1), most radio services are allocated an exclusive band of the spectrum by national authorities. Until recently, frequency blocks were reserved for different radio technologies. For instance, analog FM stations usually operate between 87.5 and 108.0 MHz (with a 200 kHz bandwidth per channel), GSM services operate between 890 and 915 MHz, 935 and 960 MHz (GSM-900, also with a 200 kHz channel bandwidth) and between 1770 and 1880 MHz (GSM-1800). To allow more flexibility and faster innovation, some reserved bands called Industrial, Scientific and Medical (ISM) allow unlicensed operation as long as emission power limits are met (to minimize interference). This is the case for the spectrum between 2400 and 2500 MHz, which is used by WiFi, Bluetooth and IEEE 802.15.4 ZigBee networks.

Recently, novel regulations in the United States [51], Europe [30] and Japan have allowed the emission of so-called Ultra Wideband signals, characterized by a very large bandwidth:

$$B > \min(500 \text{ MHz}, 0.2 F_C) \quad (1.2)$$

where F_C is the signal center frequency. To protect from interference already deployed systems, for which some companies have invested huge amounts of money, or which enable safety critical services such as aviation radar and

communication systems, very strict emission limits have been imposed on this technology. Besides access to previously reserved spectrum resources and the potential for high throughput communications, UWB offers the advantage of better propagation (the received signal being thus less degraded by the channel) as the large bandwidth makes the signal less sensitive to frequency dependent attenuation effects.

Several techniques are possible to encode information on an ultra wideband signal. The most straightforward scales existing narrowband modulations by using several carrier frequencies next to each other (e.g. OFDM UWB). This thesis focuses on a more disruptive technology named Ultra Wideband Impulse Radio. UWB-IR sends brief energy pulses, and modulates the position in time and the phase of these pulses according to the data to send. This technique promises an ultra low power consumption by using an almost all-digital chip design, which can benefit from IC process technology improvements, and by duty-cycling the transceiver at the symbol level. In addition to an ultra low power consumption, UWB-IR is potentially extremely robust to multipath propagation and to interferers, and promises high precision ranging by measuring the two-way time of flight of the emitted pulse. These features make it worth evaluating for wireless sensor networks.

1.3 Problem Statement

Wireless Sensor Networking platforms must address reliability, power consumption, size and cost constraints. The communication subsystem, which is our main concern in this work, impacts reliability and power consumption. Several ultra low power medium access control protocols have been studied in the past few years [120, 164, 66, 48, 177, 163, 112, 45, 44]. However, the introduction of a significantly different radio technology may challenge the well established solutions of the field, as the performance of MAC protocols depends largely on the underlying hardware properties. This thesis evaluates the feasibility of applying existing ULP MAC protocols on UWB-IR transceivers and considers whether a novel ULP MAC scheme specifically designed for UWB-IR could be developed. Further, the performance of the proposed ULP UWB-IR is compared to the state of the art in narrowband technology.

However, the current unavailability of UWB-IR transceivers for WSN platforms makes this technology difficult to evaluate and to compare with existing systems. Therefore, this thesis must use realistic modeling approaches to reach relevant results. Emphasis is put on model validation whenever possible.

After studying the problem of accurate modeling of UWB-IR, and identifying an ULP MAC for UWB-IR, this thesis evaluates the selected solution and compares it to narrowband alternatives: first in the context of a small star-topology MBAN, second for a multihop wireless sensor network application.

Finally, some ideas on how to further improve the reliability and the peak performance of ULP protocols are discussed.

1.4 Contributions

This thesis presents the following main contributions:

- **Maximum Pulse Amplitude Estimation**, a novel symbol-level modeling method that enabled us to implement the first network simulator of IEEE 802.15.4A UWB PHY compatible devices. This simulation model has been open sourced and is now integrated with the MiXiM simulation framework in which it was developed [125].
- **WideMac**, a novel Ultra Low Power MAC protocol that operates without relying on Clear Channel Assessment and enables fast neighbor discovery [130].
- **The optimization of complete networking stacks using UWB-IR** and narrowband transceivers for a medical body area network application, with and without coexistence, as well as the identification of their maximum performance limits [126].
- The validation of WideMac and UWB-IR for use in a static multihop network, thereby answering positively to the question of the applicability of UWB-IR to wireless sensor networks.
- **WideMac-HA**, an innovative dual-mode MAC protocol that combines WideMac ultra low power consumption with Aloha [5] low latency and high throughput [127].
- **WiseMAC-HA**, the adaptation of the WideMac-HA concept to narrowband radios by combining WiseMAC [45, 44] and CSMA. This protocol was proposed in combination with FM-UWB [58] as CSEM's standardization proposal at the IEEE 802.15.6 task group on body area networks. This is a derived contribution, and is not included in this thesis. See [50] for further information.
- The development of highly realistic narrowband simulation models (radio transceiver, multiple access interference, IEEE 802.15.4 non beacon enabled mode CSMA protocol), which have been selected as the official models in the MiXiM modeling framework [128]. These models have also been selected for inclusion in the widely used SPECint high performance computing benchmark suite of the Standard Performance Evaluation Corporation (SPEC [3]) to represent a typical load of wireless networks scientific research.

1.5 Thesis Organization

This thesis is organized as follows. Chapter 2 presents in detail the UWB-IR technology, including possible modulation types, channel effects and receiver architectures. An introduction to the IEEE 802.15.4A standard follows, and various analytical approaches to the problem of modeling Multiple Access Interference in UWB-IR networks are discussed. The general principles and assumptions of our approach, named Maximum Pulse Amplitude Estimation, are presented, followed by a description of how this approach was used to model the standard, and by a validation of this model with published MATLAB results. An initial evaluation of IEEE 802.15.4A is made using the simulator, studying the variation of packet error rate and bit error rate over distance for

various packets sizes and channel models, as well as the impact of Multiple Access Interference.

Chapter 3 consists of an analytical discussion of WideMac. It begins with an overview of wireless MAC protocols and the well-known problems of the field. A presentation of the state of the art in Ultra Low Power MAC protocols follows, and UWB-IR MAC protocols are described. The description of WideMac is accompanied by several analytical models to identify limits on its parameters and guide its design. Later, analytical models of power consumption and average latency enable a comparison of WideMac with other MAC protocols.

Chapter 4 uses the UWB-IR network simulator presented in chapter 2 to compare WideMac and UWB-IR to other wireless communication solutions (WiseMAC and a narrowband 868 MHz Texas Instruments CC1100 transceiver, IEEE 802.15.4 non beacon enabled mode and TI CC2420 2.4 GHz narrowband transceiver, and WiseMAC and FM-UWB) in the context of a star topology medical body area network application. By measuring the packet success rate, average latency, sink and sensor average power consumption for various configurations of the networking stack, optimal parameter values and performance limits are identified for all solutions. The same experimentation is performed by including this time several coexisting networks, and it is shown that each solution can cope with some level of multiple access interference if it is adequately configured.

Chapter 5 performs a similar study, but considers a static multihop wireless sensor network instead of a small MBAN. Two topologies (line and grid) are considered, with various network densities (the minimum distance between two nodes varies between 10 and 50 meters) and network sizes (from 9 to 30). It is shown that WideMac combined to UWB-IR can perform as well as a mature WSN solution such as WiseMAC combined to TI CC 1100, even in a context where the lack of clear channel assessment may penalize UWB-IR.

Chapter 6 presents an analytical study of the WideMac-High Availability protocol. The chapter begins by observing that current ultra low power MAC protocols are close to the ideal power consumption, to the point that wireless sensor networks can reach the end of their operating lifetime with significant remaining energy resources. On the other hand, there exist many applications for which the reliability of wireless communications and their peak performance can be an issue. Therefore, there is a need for ultra low power MAC protocol with higher peak performance and increased reliability, and supplementary energy resources are often available to potentially answer these needs. The description of the WideMac-High Availability protocol follows, a protocol that combines, in a compatible way, the ultra low power consumption of WideMac with the higher performance of Aloha, and proposes to increase the reliability by operating over multiple frequencies. Analytical models of power consumption and latency are presented, to quantify the potential performance improvements of this scheme. It is found that the performance improvement is significant and justifies further research in this field.

Finally, chapter 7 concludes this thesis. It restates the main contributions of this thesis, answers the questions studied in this work and highlights the novel research topics that have been identified and which are deemed worth investigating.

An Ultra Wideband Impulse Radio PHY Layer Model for Network Simulation

This chapter presents a novel modeling technique, named Maximum Pulse Amplitude Estimation, for the network simulation of UWB-IR devices.

Modeling has always been an essential aspect of scientific research [146], and the field of computer networks is no exception to this fact. In particular, the NS-2 network simulator [49], offering a programmable modeling framework and a library of reusable models, has long been the preferred tool for performance evaluation in networking research. Its simple model of the link layer allows fast execution and is sufficient for wired networks. Its use of the C++ programming language allows to simulate the evaluated protocols very closely to their actual implementation. It is thus not surprising that research in wireless networks was first based on a version of NS-2 modified to account for the characteristics of the wireless medium. Unfortunately the fundamentally different nature of this medium was difficult to capture in NS-2, leading to simulation results diverging widely from reality [22, 29, 73]. Since then, much more detailed and accurate models of the propagation effects, multiple access interference and radio state machine have been developed in various network simulators [13, 67]. However, the research community remains skeptic when considering results obtained with these tools.

Trying to simulate an UWB-IR network in today's wireless network simulators introduces challenges not unlike those posed by the adaptation of wired network models to narrowband wireless networks: the time-domain and impulsive nature of the technology makes it difficult to model using the signal to noise information commonly used for narrowband systems. Numerous analytical models of multiple access interference have been proposed [96, 97, 19, 62, 40]. However, mathematical difficulties have restricted these models to one specific type of receiver (the correlation receiver), and no model has considered the modulation specified in the IEEE 802.15.4A standard [150], despite it being the only standard featuring an UWB-IR PHY.

This chapter presents, to our knowledge, the first model of the IEEE 802.15.4A UWB PHY for network simulation. After evaluation of the related models of UWB-IR, we concluded that all of them required substantial effort to adapt to

the standard. Besides, this effort would have led to a model of reduced scope since most existing analytical models capture in one single formula the effects of the channel model, the modulation type, the multiple access interference and the receiver architecture. Therefore, we adopted a more computationally intensive but much more generic approach by developing a symbol-level UWB-IR network simulator. This was facilitated by the existence of a modeling framework precisely developed for this kind of research, named MiXiM [81]. This technique is somewhat similar to MATLAB simulation models used by researchers in radio design, which also consider each symbol, and one could wonder why these models were not used directly instead. There are several reasons for this decision, concerning execution speed, higher layer models and availability.

Regarding execution speed, MATLAB models make use of discrete signal representations of high resolution. This approach leads to highly accurate results, but it requires large amounts of memory. Increasing the traffic rate and the number of simultaneously active nodes quickly increases the processing times and memory requirements. Contrarily, a network simulator is designed from scratch for such scalability in terms of memory requirements and processing time.

The implementation of higher layer models is an important part of the simulation modeler's task. While implementing a medium access control or routing protocol in MATLAB is not impossible, the process is not trivial and the resulting implementation will likely differ a lot from an implementation in real hardware. Several components must be modeled as well, such as a radio model. A network simulator comes with a ready to use radio model, well-designed abstractions, and allows to implement the protocols in a programming language such as C, C++ or Java.

Finally, while the vast majority of MATLAB simulation models are kept in the labs, requiring large efforts to reimplement such models from scratch, there are several high quality open source network simulators available online.

This chapter builds a model of UWB-IR for the computer simulation of sensor networks. It considers the characteristics of the propagation channel, the transmitter and receiver architectures, the modulation type, the robustness to multiple access interference and the efficiency of error correction coding. Section 2.1 describes the basic operating principles of UWB-IR, section 2.2 explains the possible modulations, section 2.3 discusses the properties of the UWB-IR channel and section 2.4 presents the existing receiver architectures. Section 2.5 gives an overview of IEEE 802.15.4A, the standardized UWB PHY layer for sensor networks. After the introduction of all relevant notions in the preceding sections, the problem of multiple access interference is described in section 2.6 and the existing methods for modeling this phenomenon in the context of UWB-IR are presented. Our approach is then explained in section 2.7 and its use to model an IEEE 802.15.4A compatible energy-detection transceiver is described in section 2.8. Section 2.9 discusses the validity of the simulation model by comparing performance results with MATLAB models, with and without multiple access interference. Section 2.10 analyzes the result of extensive simulations considering several channel models and multiple access interference. Section 2.11 concludes the chapter.

2.1 Ultra Wideband Communications

In the 19th century, James Clerk Maxwell laid down the basics of a theory explaining electromagnetic phenomena by introducing the novel concepts of electric and magnetic fields [156]. These groundbreaking ideas were not immediately accepted by the day's practitioners, who preferred their more familiar concept of action at a distance. Eventually, the theory's immense success in predicting a large range of physical phenomena, including the speed of light, has led to a multitude of technological innovations with a huge influence on the way we live.

Solving Maxwell's equations in vacuum leads to wave equations for the electric field \overline{E} (Volt/meter) and for the magnetic field \overline{H} (Ampere/meter), whose fronts must travel at the speed of light c in free space. By considering Maxwell's equations in vacuum in presence of charges and currents, one can learn how to generate such a field and how to manipulate its distribution in space. Brought together, these two results enable the exchange of information at the speed of light in free space by generating and modifying an electromagnetic field.

Information is commonly modulated on this electric field by using it to slightly modify the signal frequency (digital frequency shift keying or analog frequency modulation), its amplitude (Amplitude Modulation), its phase (Phase Shift Keying), or simply by the presence and disappearance of this field (On-Off Keying). Due to propagation characteristics, channel capacity and technological limitations, wireless communication users must coexist on a fraction of the existing spectrum. Most of it is reserved for exclusive use: airport radar systems, cellular communications, FM radio stations, TV broadcasts, military systems...

Some bands, most notably the Industrial, Scientific and Medical (ISM) bands, are expressly reserved for unlicensed access. This allows greater flexibility and faster innovation, at the expense of performance: the number of users in any given place is not predictable and thus interference can be a major problem. This is the case today with WiFi networks in high density environments. Strict limits on the radiated power aim at minimizing this problem and avoiding health hazards.

Recently, advances in transceiver technology have made possible efficient communications over wide signal bandwidths. These systems, called Ultra Wide Band, are defined by the U.S. Federal Communications Commission (FCC) [51, 182], by an absolute bandwidth (the difference between the upper and lower frequencies f_H and f_L of the -10dB emission mask) larger than 500 MHz, or a fractional bandwidth ($B_{frac} = \frac{2(f_H - f_L)}{f_H + f_L}$) larger than 0.2 times their center frequency if this value is smaller than 500 MHz:

$$B > \min(500 \text{ MHz}, 0.2 F_C) \quad (2.1)$$

Current UWB systems have a center frequency between 3 and 10 GHz, a spectrum that is already in use by legacy systems. To protect those, strict power spectral density (PSD) limits on the mean Equivalent Isotropic Radiated Power (EIRP) have been defined by the FCC and by the Electronic Communication Committee (ECC) of the European Conference of Postal and Telecommunications Administration (CEPT) [30] in Europe. The FCC outdoor mean EIRP and the CEPT masks are illustrated on Figure 2.1. It can be seen that the

most favorable band for UWB communications is between 6.5 and 8.5 GHz, for which the limit is at $-41.3 \text{ dBm} / \text{MHz}$. This frequency range is in the so-called UWB high band (6-10.6 GHz) and is favored by regulators because it is less used by legacy systems (GPS, cellular services and radar systems) than the UWB low band (3.1-5 GHz).

For a 500 MHz UWB signal, the FCC and ECC PSD limits lead to a maximum emitted power of $-41.3 + 10 \log_{10} 500 = -14.31 \text{ dBm} = 10^{-14.31/10} = 37 \mu\text{W}$. As UWB-IR allows concentrating the energy in time, another limit has been defined on the maximum peak EIRP, at $0 \text{ dBm} / 50 \text{ MHz}$. For systems whose pulse repetition frequency (PRF) is higher than 1 MHz (see [18] and [51] p. 76), the maximum average EIRP is more constraining and the maximum peak EIRP limit is respected *de facto*. In this thesis, systems with a PRF lower than 1 MHz are not considered, hence the average EIRP limit applies.

This can be compared to the maximum allowed transmission power of the so-called Industrial, Scientific and Medical (ISM) unlicensed bands in use by most sensor networks at 900 MHz and 2.4 GHz. The regulation for the ISM bands requires the use of a spread spectrum modulation technique (either direct sequence or frequency hopping). The FCC allows up to 1 W maximum power on these bands. In Europe, the 900 MHz band is used by the GSM networks and it is the 868 MHz band that is used instead. The European regulation is stricter and allows up to 100 mW maximum power. Even in the least favorable case for the ISM bands, UWB has a $10 \log(100/0.037) = 34 \text{ dB}$ penalty with a 500 MHz bandwidth. We will now see how the characteristics of UWB can help to overcome this difficult situation.

Contrarily to the ISM bands, the UWB definition does not mandate any type of modulation technique. There are indeed several possibilities to generate such a signal:

- Multi-band Orthogonal Frequency Division Multiplexing (MB-OFDM-UWB [61]) uses several sub-carriers to transmit information, and is used in Wireless USB products ;
- Frequency Modulation UWB (FM-UWB [58]) modulates the data stream on a low central frequency using Frequency Shift Keying (FSK) modulation, and this signal is then spread over 500 MHz at a higher center frequency using an analog step frequency modulation ;
- Impulse Radio (UWB-IR [170]) is a communication scheme that uses short energy pulses instead of a carrier signal ;
- Chaotic UWB (see [25] and references therein) use deterministic, non-periodic and random-like signals derived from nonlinear dynamical systems. These systems require sophisticated receiver architectures ;
- Chirp Ultra Wideband [89] generates a signal by frequency sweeping at the given bandwidth.

MB-OFDM-UWB is a complex system that leads to relatively expensive chips using large amounts of power. It is best suited for short range high data rate communications, such as between a digital camera and a computer, or between a DVD player and a television screen. Thanks to broad industry support and to its use of well-known technologies, this was the first UWB technology to reach




Figure 2.1: European and U.S. UWB emission masks.

the market. Because the MB-OFDM 500 MHz UWB signal is a juxtaposition of multiple 4 MHz narrowband channels, its propagation can be modeled with the mathematical tools used for narrowband channels.

FM-UWB is a simple modulation scheme that favors transmission robustness and ultra low power consumption over high bit rates. It combines a frequency shift keying modulation to encode the digital information on a low carrier frequency with an analog frequency modulation step to reach the desired center frequency and to fill the emission mask. It is possible to send several FM-UWB signals on the same center frequency (called sub-carriers) by slightly modifying the carrier frequency used by the FSK modulation. The cross-interference is kept minimal even with four simultaneous transmissions (one per sub-carrier).

UWB-IR is a more radical departure from classical radio designs. It can provide both high data rates at short ranges and low data rates at longer ranges. The signal propagation differs vastly with narrowband signal propagation. First, it offers high robustness to frequency selective fading, as in that case most of the energy will propagate unaffected while all of a narrowband signal would be modified. Second, due to its time-based nature and the short duration of the pulses, the multipath components of the received UWB signal can be isolated, thereby greatly reducing inter- and intra- symbol interference. A particular characteristic of UWB is that the radiated energy is concentrated in time. Thus, if the modulation generates pulses only during 1% of the symbol




Figure 2.2: Impulse response of a circular monopole antenna system to a Gaussian pulse. The normalized voltage at the antenna is plotted as a function of time in nanoseconds (source: [117]).

duration, the radiated power of this pulse can be 100 times higher than the average EIRP limit. The regulation agencies were aware of this, and introduced another limit, this time on the peak EIRP. It is set at 0 dBm/50 MHz. In this case, the penalty compared to the radiated power of narrowband ISM systems is $10 \log_{10}(100/1) = 20 \text{ dB}$.

In addition, the UWB-IR signal can be used for high precision ranging, providing additional information to the application. Finally, the radio transceiver design can offer ultra low power consumption, especially for the transmitter.

The Gaussian pulse is one of the most commonly considered UWB-IR pulses. It is illustrated on Figure 2.2 and can be expressed as follows:

$$x(t) = \frac{A}{\sqrt{2\pi}\sigma} \exp\left(-\frac{t^2}{2\sigma^2}\right),$$

where A is a constant to meet the emission mask and σ is the parameter to set the center frequency.

To match the FCC spectrum emission mask, derivatives of the Gaussian pulse are preferred [141]. The spectrum magnitude of the n th order derivative of the Gaussian pulse is given by:

$$|P(f)| = A (2\pi f)^n \cdot \exp\left(-\frac{(2\pi f\sigma)^2}{2}\right).$$

σ and n must be chosen jointly so as to maximize the use of the spectrum. This problem is especially difficult in multi-band systems. Further, implementing this pulse shape using CMOS technology and at low power consumption is challenging [131].

Carrier-based UWB-IR [131] generates a triangular UWB pulse (by charging and discharging a linear capacitor) which is up-converted with a carrier wave. Figure 2.3 illustrates the architecture of such a transmitter. These two signals and their product are given below, with T being the pulse half-duration,




Figure 2.3: Carrier-based impulse radio transmitter with its three main parameters: pulse duration (T_p), carrier frequency (F_c) and amplifier gain (G).




Figure 2.4: 2 ns (top row) and 0.66 ns (bottom row) triangular baseband pulses up-converted at 4.5 GHz (left column) and 8 GHz (right column) with cosinusoidal carriers.

A the pulse peak amplitude (normalized to 1 on the figures) and f_c the center frequency. Figure 2.4 illustrates the generated signal for carrier frequencies at 4.5 GHz and 8 GHz and for two pulse durations (2 ns and 0.66 ns).

$$\begin{aligned}
 \text{triangle}(t) &= A \left(\frac{t}{T} 1_{\{0 < t \leq T\}} + \frac{2T - t}{T} 1_{\{T < t \leq 2T\}} \right) \\
 \text{carrier}(t) &= \cos(2\pi f_c t) \\
 \text{signal}(t) &= \text{triangle}(t) \cdot \text{carrier}(t)
 \end{aligned}$$

The signal bandwidth is controlled by the triangular signal duration and the signal center frequency is given by the carrier sine wave frequency. This approach is simpler to implement, provides more flexibility and allows to reach

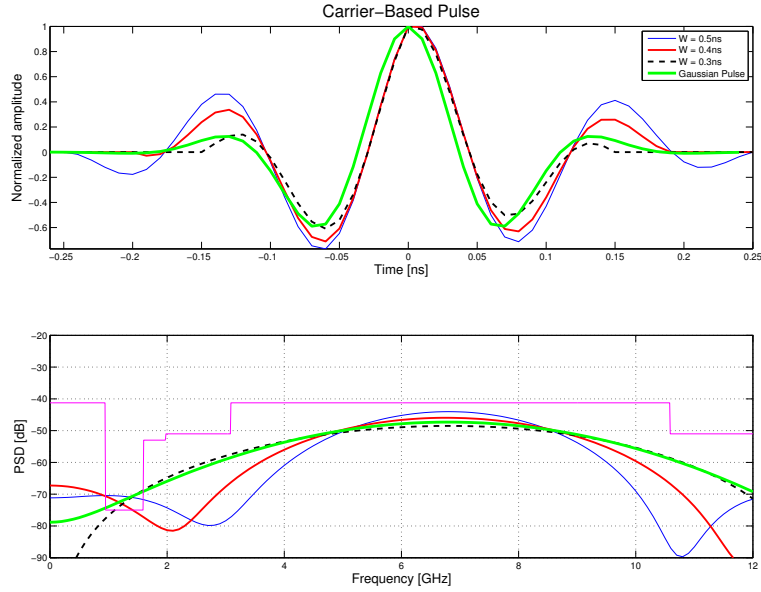


Figure 2.5: Carrier-based pulses for three pulse durations (0.3, 0.4 and 0.5 ns) and 3rd derivative Gaussian pulse with a center frequency at 4.25 GHz. Carrier-based pulses are a good approximation of the Gaussian pulse (source: [117]).

low power consumption levels. The spectrum magnitude of the generated signal is given by [131]:

$$|P(j\omega)| = \left[\frac{2\sqrt{2}Q}{t_r t_f \omega^2} \sqrt{\frac{(t_f \sin(\omega t_r) - t_r \sin(\omega t_f))^2 + (t_f (1 - \cos(\omega t_r)) + t_r (1 - \cos(\omega t_f)))^2}{2(t_r + t_f)^2}} \right],$$

which is composed of a first envelope term and of a second oscillating term. Figure 2.5 shows three carrier-based pulses with different triangular pulse durations, and a third order Gaussian pulse. It can be seen that the carrier-based UWB pulses are similar to the shape of a Gaussian pulse.

Various physical characteristics of an electromagnetic pulse can be modified to modulate data: the pulse polarity, the pulse amplitude and the pulse position in time. Each modulation has its advantages and disadvantages. Section 2.2 gives an overview of UWB-IR modulations.

2.2 UWB-IR Modulations

The process of encoding a digital stream of information into physical signals is called *digital modulation*. In practice, consecutive information bits are grouped together and each group is mapped to a specific physical *symbol*. With some simple modulations, only one bit is mapped to one symbol and we can talk of binary symbols. At the opposite, complex modulations can encode up to 64 bits in one symbol (64-QAM for ADSL communications).

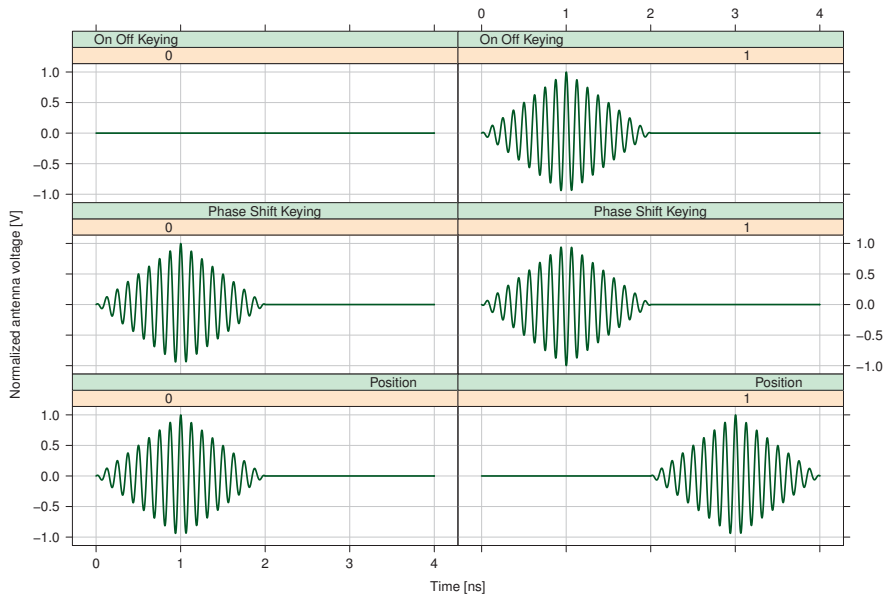


Figure 2.6: On-Off Keying, Phase Shift Keying and Pulse Position Modulation (first, second and third rows) for bit values 0 and 1 (first and second rows) for a 2 ns pulse centered at 8 GHz.

This section illustrates some important modulation techniques for UWB-IR: On-Off Keying (OOK [169]), Pulse Position Modulation (PPM [136]) and Phase Shift Keying (PSK [169]).

On-Off Keying modulates the presence of the pulse itself with the information bit. A pulse will be sent if the bit is set to one. If the bit is set to zero nothing will be sent. At the receiver, the presence or absence of a pulse is estimated and this estimation is used to recover the information stream.

The information bit can also be used to control the time at which the pulse is sent. This technique is called Pulse Position Modulation. For instance, with a bit set to zero, the pulse will be sent at the beginning of the symbol duration, and at the end of the symbol duration otherwise. The receiver estimates when the pulse was sent during the symbol to demodulate the information.

It is also possible to modify the pulse polarity (also called its phase), as with Phase Shift Keying. The phase information can only be detected by a special category of receivers, named coherent, as will be explained later in section 2.4.

Figure 2.6 illustrates the effect of On-Off Keying, Phase Shift Keying and Position Modulation on a pulse for the two possible bit values. The horizontal and vertical axes represent respectively the time in nanoseconds and the voltage applied to the antenna. Assuming free space propagation and neglecting antenna effects, this is also, at a scaling factor, the voltage that would be induced on the receiving antenna.

Other modulations have been defined, such as Transmitted Reference (TR [71]) which is similar to PPM or PSK except that a reference pulse is always sent at the symbol start.

2.3 The UWB-IR Channel

The propagation channel modifies the transmitted signal in a non-deterministic manner. These changes make the reception process prone to errors. The communication system will see its performance vary depending on the channel type: free-space, line of sight, non line of sight, indoor residential, indoor office, outdoor...

This section introduces channel models in an increasing order of complexity and realism. Subsection 2.3.1 introduces the concepts required to understand electromagnetic propagation, assuming free space and neglecting antenna effects. Subsection 2.3.2 presents the Ghassemzadeh channel model, which offers stochastic models of the path gain and of the shadowing effect. The section continues with the Saleh-Valenzuela model in subsection 2.3.3 that considers the individual multipath components of the signal. Finally, subsection 2.3.4 presents the IEEE 802.15.4A channel models, all based on a modified Saleh-Valenzuela model.

2.3.1 Free Space Propagation

Let us first consider the emission of the electromagnetic wave. The transmitter circuit applies a voltage and generates an electrical current in the antenna, which from an electrical circuit theory point of view can be modeled as a $R_{Ant} = 50\Omega$ resistor.

We assume that the antenna is isotropic, its gain G is constant and equal to 1 in all directions. We also assume that all the energy sent to the antenna is completely converted in an electromagnetic wave: the transmit efficiency η_{Tx} is equal to 1. Then, we have the following relations:

$$\begin{aligned} P_{TxAmp} &= \frac{v_{Tx}^2}{R_{Ant}} = P_{Radiated} (\eta_{Tx} = 1) \\ I &= \frac{P_{Radiated}}{4\pi d^2} (G = 1) \\ I &= E \cdot H = \frac{E^2}{120\pi}, \end{aligned} \quad (2.2)$$

where P_{TxAmp} is the power delivered by the amplifier to the antenna, v_{Tx} is the voltage applied at the antenna, $P_{Radiated}$ is the power radiated by the antenna and d is the distance from the antenna. It follows that:

$$\frac{1}{4\pi d^2} \frac{v_{Tx}^2}{R_{Ant}} = \frac{E^2}{120\pi}.$$

We observe that the strength of the electric field E is directly proportional to the voltage applied at the antenna v_{Tx} , and that the radiated intensity I decreases with the square of the distance.

Let us now consider what happens at the receiver side. Similarly, the antenna is ideal isotropic with a constant gain $G = 1$ over the considered bandwidth, and with a nearly perfect reception efficiency $\eta_{rx} = 1$.

The received power at the antenna P_{RxAnt} is:

$$P_{RxAnt} = \frac{\lambda^2}{4\pi} I = \frac{\lambda^2}{4\pi} \frac{E^2}{120\pi} = \frac{\lambda^2}{4\pi} \frac{1}{4\pi d^2} P_{Radiated} = \frac{\lambda^2}{4\pi} \frac{1}{4\pi d^2} P_{TxAmp}, \quad (2.3)$$

where λ is the signal wavelength. The voltage induced at the antenna is obtained simply:

$$v^2_{Rx} = R_{Ant}P_{RxAnt}.$$

In the case considered above, the propagation model is captured in formula 2.2. This relation, called *pathloss*, allows to compute, from the initial transmitted power, the radiated power intensity at the receiver. This intensity, when multiplied by the antenna aperture, gives the received power. Some pathloss models exclude antenna effects and take the form of equation 2.2. Others include it and take the form of equation 2.3. In other propagation environments, the relation 2.2 is modified in several ways:

- First, the wave propagation differs in the so-called short and far fields. Typically, the short field path loss is measured at a reference distance d_0 (often set to 1 meter) and an equation similar to 2.2 gives the pathloss in the far field as a function of the distance d divided by the reference distance d_0 . The losses due to the short and far fields must be summed up to evaluate the total pathloss.
- Second, the distance dependency can be much more important, by using a so-called pathloss exponent n instead of the square of the distance, that can reach values as high as 7 in some environments. With some models the path loss exponent is fixed, with others it is described by a probability distribution.
- Third, shadowing effects due to obstacles can be modeled using a random variable of zero mean.
- Finally, multipath propagation can also be modeled. This is typically done with stochastic models that give the relative arrival times of the rays as well as how the energy of the initial pulse is distributed among them. Such a model is commonly called a power delay profile.

The following section describes a more sophisticated model that implements all of these points except a power delay profile. This aspect is treated at length in the last two sections.

2.3.2 The Ghassemzadeh Channel Model

Ghassemzadeh et al. [60] have developed a statistical path loss model for UWB propagation in residential environments, for both line of sight and non line of sight environments. It is based on 300 000 frequency-response measurements of 1.25 GHz-wide pulses with a central frequency of 5 GHz, taken in 23 homes.

The mean path loss is defined by the transmit power multiplied by the transmit and receive antenna gains divided by the received power:

$$Pl = \frac{P_t G_t G_r}{P_r}$$

The authors model the path loss in dB at a distance d as follows:

$$PL(d) = \left[PL_0 + 10\gamma \log_{10} \left(\frac{d}{d_0} \right) \right] + S(d), \quad d \geq d_0$$

Table 2.1: Parameters of the Ghassemzadeh statistical channel models for LOS and NLOS residential environments.

Parameter	Description	LOS	NLOS
PL_0	path loss at d_0	47 dB	51 dB
μ_γ	mean path loss exponent	1.7	3.5
σ_γ	standard deviation of the path loss exponent	0.3	0.97
μ_σ	mean std. dev. of the shadowing	1.6	2.7
σ_σ	std. dev. of the std. dev. of the shadowing	0.5	0.98

where the reference distance d_0 is equal to 1 m, PL_0 is the path loss at the reference distance, γ is the path loss exponent and S is the shadow fading. The shadow fading is shown to be log-normal (with a standard deviation σ) in their measurements, confirming other researchers' results [134, 121].

Their path loss model considers γ and σ as random variables. PL_0 , the path loss at the reference distance, is a parameter whose values are given for both LOS and NLOS environments in table 2.1. The path loss exponent was found to have a normal distribution $N[\mu_\gamma, \sigma_\gamma]$:

$$\gamma = \mu_\gamma + n_1\sigma_\gamma$$

where n_1 follows a standard normal distribution ($N[0,1]$).

The shadow fading model follows a similar approach. Shadow fading follows a log-normal distribution, and thus the shadow fading in dB follows a normal distribution of standard deviation σ . According to the Ghassemzadeh model, this parameter has a normal distribution of parameters $N(\mu_\sigma, \sigma_\sigma)$. This can be rewritten as a function of the standard normal random variables n_2 and n_3 :

$$\begin{aligned} S &= n_2\sigma \\ \sigma &= \mu_\sigma + n_3\sigma_\sigma \end{aligned}$$

This leads to the following expression for the path loss in dB:

$$\overline{PL}(d) = [PL_0 + 10\mu_\gamma \log_{10} d] + [10n_1\sigma_\gamma \log_{10} d + n_2\mu_\sigma + n_2n_3\sigma_\sigma] \quad (2.4)$$

To prevent γ and σ from taking unrealistic values, the authors recommend the use of truncated Gaussian random variables for n_1 , n_2 and n_3 . They suggest the following bounds:

$$n_1 \in [-1.25, 1.25] ; n_2, n_3 \in [-2, 2]$$

2.3.3 The Saleh-Valenzuela Channel Model

Saleh and Valenzuela [134] measured the multipath propagation of 1.5 GHz wide, 10 ns ultra wideband pulses in an office building. The delay spread was




Figure 2.7: Representation of the Saleh-Valenzuela model, with four clusters shown. The vertical axis represents the expected average power and the horizontal axis shows the ray arrival times.

found to be as large as 200 ns with a maximum root mean square (RMS) value of 50 ns. They developed a statistical model that fits their measurements well and that can be extended to other buildings.

This model considers that multipath components (MPC, often called rays, or taps) arrive grouped in clusters, as shown on figure 2.7. The channel impulse response, summing over all clusters (variable l) and all rays in each cluster (variable k) is given by:

$$h(t) = \sum_{l=0}^{\infty} \sum_{k=0}^{\infty} \beta_{kl} e^{j\theta_{kl}} \delta(t - T_l - \tau_{kl}).$$

The models used to compute the path gains β_{kl} , the path phases θ_{kl} , the cluster arrival times T_l and the ray arrival times within each cluster τ_{kl} are described below.

The cluster arrival times T_l form a Poisson arrival process of rate Γ and the ray arrival times τ_{kl} within each cluster k form a Poisson arrival process of rate γ :

$$P(T_l | T_{l-1}) = \Lambda \exp[-\Lambda(T_l - T_{l-1})]$$

$$P(\tau_{kl} | \tau_{(k-1)l}) = \lambda \exp[-\lambda(\tau_{kl} - \tau_{(k-1)l})]$$

Each ray gain β_{kl} follows the Rayleigh distribution with a standard deviation $\overline{\beta_{kl}^2}/2$:

Model	Description
CM1	residential LOS.
CM2	residential NLOS.
CM3	indoor office LOS.
CM4	indoor office NLOS.
CM5	Outdoor LOS.
CM6	Outdoor NLOS.
CM7	Open outdoor LOS.
CM8	Industrial LOS.
CM9	Industrial NLOS.

Table 2.2: IEEE 802.15.4A channel models.

$$p(\beta_{kl}) = (2\beta_{kl}/\overline{\beta_{kl}^2}) \exp(-\beta_{kl}^2/\overline{\beta_{kl}^2}),$$

where $\overline{\beta_{kl}^2}$ is the expected value of the path power gain. This parameter depends on the average power gain of the first ray of the first cluster $\overline{\beta_{00}^2}$, on the cluster and ray arrival times T_l and τ_{kl} and on the cluster and ray power-delay time constants Γ and γ as follows:

$$\overline{\beta_{kl}^2} = \overline{\beta_{00}^2} e^{-T_l/\Gamma} e^{-\tau_{kl}/\gamma}.$$

Finally, the expected power gain of the first ray of the first cluster $\overline{\beta_{00}^2}$ is given as follows:

$$\overline{\beta_{00}^2} = (\gamma\lambda)^{-1} G(1m)r^{-\alpha}$$

where $G(1m)$ is the gain at a reference distance of 1 meter:

$$G(1m) = \frac{P_r(1m)}{P_t} = G_t G_r \left(\frac{\lambda_0}{4\pi} \right)^2$$

and P_r and P_t are the received and transmitted powers, G_t and G_r are the transmission and reception antenna gains and λ_0 is the RF wavelength. The ray phases θ_{kl} are independently and uniformly distributed in $[0, 2\pi)$. From their measurements, the authors estimate Λ^{-1} to be close to 300 ns and λ^{-1} between 5 and 10 ns.

2.3.4 The IEEE 802.15.4A Channel Model

During the development of the IEEE 802.15.4A standard, a channel modeling subgroup was created. It defined several channel models to allow fair comparison of the standardization proposals. These models can be found in their final report [100].

This section describes the IEEE 802.15.4A UWB channel model for frequency ranges between 2 and 10 GHz, which covers Line of Sight and Non Line of Sight residential, office, industrial and outdoor environments (they are listed in table 2.2). It consists of two parts. The first part can be used to compute the pathloss, and thus the energy of the emitted pulse that reaches the destination. The second part is the power delay profile, and describes how the energy from the emitted pulse arrives in time, spread over numerous multipath pulses.

The pathloss is expressed as follows:

$$PL(f, d) = \frac{1}{2} PL_0 n_{Tx}(f) n_{Rx}(f) \frac{(f/f_c)^{-2(\kappa+1)}}{(d/d_0)^n} \quad (2.5)$$

where PL_0 is the pathloss at the reference distance, $n_{Tx}(f)$ and $n_{Rx}(f)$ are the antenna efficiency in reception and transmission, κ accounts for the frequency dependence of the pathloss, f is the considered frequency, f_c is the pulse center frequency, d is the distance between the source and the receiver, d_0 is a reference distance (set to 1 meter) and n is the pathloss exponent. The antenna is an ideal isotropic antenna and thus its gain is equals to 1 and not shown in the formula. The factor $1/2$ is an approximation for some particular antenna effects due to the presence of human beings. In this work, antenna efficiency is considered equal to 1 over all frequency ranges. In addition, the signal bandwidths are relatively small (500 MHz) and the frequency dependency can be neglected. Shadowing is modeled as a Normal random variable $N[0, \sigma_S]$. While, in practice, the antenna efficiency varies with frequency, we considered in this work that it was constant and equal to 1. Over a 500 MHz bandwidth, the assumption that this parameter is constant is realistic [147].

The power delay profile model is similar to the Saleh-Valenzuela model described in the previous section. Multipath components also arrive in clusters, whose inter-arrival times are exponentially distributed. However, here ray arrival times follow a mixed Poisson distribution.

The channel impulse response is expressed as:

$$h(t) = \sum_{l=0}^L \sum_{k=0}^K \beta_{kl} e^{j\theta_{kl}} \delta(t - T_l - \tau_{kl}) \quad (2.6)$$

where L is the number of clusters, K is the number of rays in a cluster, β_{kl} is the amplitude of the k th ray in the l th cluster and θ_{kl} is its phase, T_l is the delay of the l th cluster and τ_{kl} is the delay of the k th ray relative to the cluster start. The number of clusters L is Poisson distributed, with parameter \bar{L} . The distribution of the cluster inter-arrival times $T_l - T_{l-1}$ also follows a Poisson distribution, with parameter Λ .

The IEEE 802.15.4A departs from the S-V model when considering the ray inter-arrival times $\tau_{kl} - \tau_{(k-1)l}$. Instead of considering a Poisson distribution, it uses a mixture of two Poisson processes:

$$p(\tau_{kl} | \tau_{(k-1)l}) = \beta \lambda_1 \exp(-\lambda_1 (\tau_{kl} - \tau_{(k-1)l})) + (\beta - 1) \lambda_2 \exp[-\lambda_2 (\tau_{kl} - \tau_{(k-1)l})], \quad k > 0$$

The power delay profile (PDP) is given as a function of the total energy of the ray's cluster Ω_l :

$$E \left\{ |\beta_{kl}|^2 \right\} = \Omega_l \frac{1}{\gamma_l [(1 - \beta) \lambda_1 + \beta \lambda_2 + 1]} \exp(-\tau_{kl}/\gamma_l)$$

with the cluster energy Ω_l given as a function of the cluster arrival time, of the inter-cluster decay constant Γ and of a random variable as follows:

$$10 \log(\Omega_l) = 10 \log(\exp(-T_i/\Gamma)) + M_{cluster}$$

where $M_{cluster}$ is a $N(0, \sigma_{cluster})$ random variable.

An alternative PDP is used for office and industrial NLOS environments.

Table 2.3 gives an overview of all of the IEEE 802.15.4A channel models with their key parameters. The pathloss attenuation at 1 meter, the pathloss exponent and the mean number of clusters vary greatly.

2.4 Receiver Architectures

The performance of an UWB-IR communication system depends largely on the receiver. It must detect the presence of an incoming signal, acquire its timing, correctly demodulate the encoded data and apply error correction techniques. Because of the low radiated power levels and the very short pulse durations, detecting an incoming transmission is not a trivial task. The propagation channel distorts the pulse shape, and often generates several multipath components among which the transmitted signal energy is spread.

Several receiver architectures have been proposed to address these problems. They can be separated into two groups: *coherent* and *non-coherent* receivers.

Coherent receivers can demodulate the pulse polarity, also called the phase of the signal. They reach lower bit error rates at the price of complex receiver architectures such as the rake receiver, leading to high power consumption. In addition, they require channel information: timing, fading coefficient, pulse shape of all considered rays. This information is estimated during the frame synchronization preamble.

In contrast, non-coherent receivers do not know the signal phase and do not necessarily estimate the channel. Their simpler architecture does not allow to reach the BER performance of coherent receivers, but it makes them cheaper to produce and reduce their energy consumption. The energy detector is a typical example of non-coherent receiver. It integrates the energy during a time window and squares it. The time window is often chosen as large as the typical channel delay spread in order to collect most of the signal energy. Unfortunately, a large time window also increases the collected noise [99] (a few techniques [161, 92, 168, 54] have been proposed to alleviate this problem).

The following sections 2.4.1 and 2.4.2 describe respectively some coherent and non-coherent receiver architectures.

2.4.1 Coherent receivers

Rake receivers are frequently used for coherent reception of UWB-IR signals. They are composed of several correlators, called rake fingers, and each of these fingers is configured to sample at the time delay of a particular multipath component. Increasing the number of fingers increases the system precision as it allows to capture more of the signal energy, but at the same time this




Figure 2.8: Block diagram of an energy detection receiver.

increases the receiver complexity, power consumption and channel estimation difficulties.

Three approaches exist for the number of fingers [21]. The ideal all-Rake receiver (ARAKE) uses as many fingers as there are multipath components. Unfortunately, in some environments this leads to several hundreds fingers. For a given number of fingers L , the Selective Rake (SRAKE) is the optimal rake receiver. It selects the strongest multipath components of all MPCs. A simpler, sub-optimal approach is proposed with the partial rake (PRAKE) receivers, that select only the first L MPCs. The performance difference between the PRAKE and SRAKE architectures depend on the MPCs, the number of fingers, the signal bandwidth and the center frequency.

2.4.2 Non-coherent receivers

To avoid the complexity of the rake receivers, some have considered the use of an energy detection receiver. It integrates the energy during a time window and squares it. The time window is often chosen as large as the typical channel delay spread in order to collect most of the signal energy. Unfortunately, a large time window also increases the collected noise [99]. Energy detection receivers can demodulate On-Off Keying (OOK) and Pulse Position modulation (PPM). An energy detector is illustrated on figure 2.8.

Transmitted Reference receivers [71] operate by acquiring a reference pulse which is then correlated with a second, data modulated pulse. This system is almost immune to pulse distortion as the reference pulse goes through exactly the same channel as the data modulated pulse. However, its performance is degraded due to non-linear operations on noise terms [162]. Another disadvantage is that twice as many pulses must be sent to reach the same capacity. Besides, its hardware implementation poses non-trivial challenges.

2.5 The IEEE 802.15.4A Standard

The IEEE 802.15.4 standard [151] defines physical (PHY) and medium access control (MAC) layers for wireless sensor networks. The IEEE 802.15.4A working group [150] was created to develop alternative physical layers for robust communications and high precision ranging. The standard defines two physical layers: an ultra wideband impulse radio layer and a narrow-band chirp spread spectrum layer. This section gives an overview of the former.

The standard defines the mean Pulse Repetition Frequency (PRF) as the total number of pulses emitted during a symbol period divided by the length




Figure 2.9: An IEEE 802.15.4A burst at the transmitter with carrier-based triangular pulses and using the mandatory mode parameters.

of the symbol duration. Several timing parameters are expressed as a function of the mean PRF, and compliant transceivers must support at least one of the two mandatory values of 15.6 and 3.9 MHz (the third and optional value is 62.4 MHz). We considered initially only the higher value as it allows operation at a lower peak voltage [150] and because we could find information on a transceiver using this mode in the literature [133]. According to the standard, the lower and average mean PRF should lead to better performance in environments with respectively high and low delay spreads. While the low mean PRF mode was also implemented in the simulation model presented in section 2.8, this section assumes the 15.6 MHz PRF for simplicity.

This section begins with an introduction to the IEEE 802.15.4A burst modulation in subsection 2.5.1, describes the frame structure in subsection 2.5.2 and concludes with the error correction mechanisms in subsection 2.5.3.

2.5.1 Burst Position Modulation

The IEEE 802.15.4A UWB PHY layer uses a *Burst Position Modulation* (BPM) scheme. Short impulses are sent consecutively to form a burst, and the time position of this burst in the symbol codes the symbol value (0 or 1). The number of pulses per burst and the pulse duration can take several possible values; in our work, only the values of the mandatory mode, 2 ns per pulse and 16 pulses per burst (derived from the mean Pulse Repetition Frequency), are considered. Figure 2.9 illustrates a mandatory burst.

Each symbol encodes one data bit value through the burst time position and one error correction bit value in the burst polarity. This last value may not be




Figure 2.10: An IEEE 802.15.4A UWB PHY Symbol

demodulated by the receiver, but must be present in the signal. Each symbol is divided into four equal time intervals (each of 256.4 ns in the mandatory mode), as shown on Figure 2.10. The first and third time intervals are further subdivided into time hopping positions (there are eight such positions in the mandatory mode) to smooth the signal spectrum. The time hopping position is uniquely determined by the symbol position in the frame: there is only one time hopping sequence and it does not depend on an initialization parameter or on the source or destination node.

If the burst encodes a zero, it is sent during one of the time hopping positions of the first window. Otherwise the burst is sent during one of the time hopping positions of the third window. The second and fourth time intervals are guard times that protect the signal against *Inter Symbol Interference* (ISI): while the signal propagates, multipath components are created by reflections. The guard times are sufficiently large so that the multipath components that could reach the next active time hopping position will be strongly attenuated. The relative durations of the guard times and of the bursts can be seen on figures 2.10 and 2.11. Table 2.4 gives the key physical properties of an IEEE 802.15.4A UWB PHY symbol.

2.5.2 Frame Structure

The very short duration of the pulses makes them difficult to detect. Since there is no carrier signal, the channel is empty most of the time even though a transmission is ongoing. The only part of the signal that can be reliably detected (using a dedicated algorithm) is the *synchronization preamble*, with which all transmissions begin. It consists of a deterministic sequence of isolated pulses used by all devices that are part of the same network (two synchronization preambles are defined in the standard). Table 2.5 gives some characteristic properties of an IEEE 802.15.4A UWB PHY frame.

Figure 2.12 shows an IEEE 802.15.4A UWB frame. It starts with the synchronization preamble sequence, shown in grey, and it is followed by the Start Frame Delimiter (SFD) and the data payload, both transmitted using Burst Position Modulation.

2.5.3 Error Correction

The standard proposes two error correction schemes: a mandatory one and an optional one.

The mandatory scheme is based on a Reed-Solomon code $RS_6(63, 55)$ that must be encoded as 48 additional symbols at the end of the packet by each transmitter. It can be demodulated and may be used by all receivers. It works

Table 2.3: Key parameters of the IEEE 802.15.4A statistical channel models.

Param.	Descr.	Residential		Indoor office		Outdoor		Open out-door
		LOS CM1	NLOS CM2	LOS CM3	NLOS CM4	LOS CM5	NLOS CM6	LOS CM7
PL_0	path loss (d_0)	43.9	48.7	35.4	57.9	45.6	73.0	49
n	path loss exp.	1.79	4.58	1.63	3.07	1.76	2.5	1.58
σ_S	shadowing std dev.	2.22	3.51	1.9	3.9	0.83	2	3.96
\bar{L}	Mean nb. clusters	3	3.5	5.4	1	13.6	10.5	3.31
Λ (1/ns)	inter-cl. arriv. rate	0.05	0.12	0.02	-	0.005	0.02	0.03
Γ (ns)	inter-cl. decay cst.	22.6	26.3	14.6	-	31.7	105	56
$\sigma_{cluster}$	cl. shadowing var.	2.7	2.9	-	-	-	-	-
Validity	range (m)	7-20	7-20	3-28	3-28	5-17	5-17	-

Table 2.4: Characteristic parameters of a symbol using the mandatory mode of the UWB-IR IEEE 802.15.4A PHY(15.6 MHz PRF).

Parameter	Symbol	Value
Chip duration	T_c	2 ns
Chips per burst	N_{cpb}	16
Burst duration	T_B	32.05 ns
Number of burst positions	N_{burst}	32
Number of time hopping burst positions	N_{hop}	8
Guard time interval	T_G	256.4 ns
Symbol duration	T_S	1025.64 ns




Figure 2.11: The bursts of two consecutive IEEE 802.15.4A symbols as sent at the transmitter (triangular pulses, mandatory mode).

Table 2.5: Characteristic parameters of a frame using the mandatory mode of the UWB-IR IEEE 802.15.4A PHY.

Parameter	Symbol	Value
Synchronization sequence duration	T_{sync}	63.6 μ s
Start Frame Delimiter duration	T_{sfd}	7.9 μ s
Duration of the SHR preamble	$T_{pre} = T_{sync} + T_{sfd}$	71.5 μ s
PHY header duration (two bytes)	T_{hdr}	16.4 μ s
Mandatory data bit rate		0.85 Mbps
Duration of 48 Reed-Solomon additional symbols	T_{RS}	49.23 μ s
PHY Overhead	$T_{pre} + T_{hdr} + T_{RS}$	137.13 μ s



Figure 2.12: An IEEE 802.15.4A UWB PHY Frame

by considering the data as a sequence of 6 bit symbols (data is padded to reach the desired size if necessary). This code can correct up to 8 erroneous 6 bits symbols. The number of errors in each of these erroneous symbols does not matter.

The optional convolutional code can be encoded by some transmitters (in the burst polarity) and can be demodulated and used by coherent receivers. It is mathematically more complex, but could improve greatly the robustness of the transmissions as it stores one parity bit in the burst polarity of each symbol, effectively doubling the number of bits transmitted. It is not considered in this work.

2.6 Modeling Multiple Access Interference

This section explains the problem of Multiple Access Interference (MAI) and describes several models that evaluate its influence on bit error rate in UWB-IR systems. The section concludes with a summary of the strong points and of the shortcomings of the existing models.

2.6.1 Overview

Multiple Access Interference (MAI) is a communication phenomenon that happens when two or more radio signals simultaneously reach the same radio receiver, potentially preventing the message reception. When this is the case, it is said that a *collision* happened. Frame collisions are a well-known cause of energy waste for ultra low power Medium Access Control (MAC) protocols [48, 46].

When simulating narrowband radio systems, the received signal strength can be computed for each frame arriving at the receiver [13]. These values can be used to evaluate the signal to noise ratio during a frame reception, and these values of the signal to noise ratio can be mapped to bit error rates using closed-form analytical expressions (depending on the modulation type). Figure 2.13 illustrates this process.

With UWB-IR systems, the transmitted signal is discontinuous, and thus the received signal strength varies much more frequently. It can be seen on Figure 2.10 that during most of the IEEE 802.15.4A symbol length (1025.64 ns in the mandatory mode) no signal is transmitted, since there is only one burst of 16 2 ns pulses per symbol. This allows the transmitter to be active only during a fraction of the symbol time: 3.1% in the mandatory mode. The receiver, which doesn't know the bit value *a priori*, must listen during both time windows time hopping positions and thus must be active during (at least) 6.2% of the symbol in the mandatory mode. The effective duration depends on the receiver implementation. Such duty-cycling of the transceiver has been studied in [39] to lower the power consumption.

This in-symbol signal duty-cycling offers robustness not only against inter- and intra-symbol interference but also against multiple access interference: since the transmitters are not synchronized, an interfering burst can occur with equal probability at any time during a symbol. There is thus only a probability $2T_{Burst}/T_{Symbol}$ that an interfering burst arrives during the time hopping position associated to the opposite bit value of the signal. The probability that




Figure 2.13: SNR estimate obtained with an Accumulative Noise SNR Model for a narrowband radio.

this interfering burst arrives during the other time hopping position is also equal to $2T_{Burst}/T_{Symbol}$. In both cases, a *pulse collision* happens.


When two frames interfere, for each symbol of the frame four cases can happen from the receiver's point of view. In the first case, the burst of the interfering signal falls between the two positions to code a 0 and to code a 1, and does not have any effect on the receiver.

In the second case, the burst of the interfering signal falls somewhere during the burst of the source. As illustrated on Figure 2.14a, this can have a positive effect with an energy-detection receiver, because the distortion of the pulse shape does not effect its performance. However, it can also have a negative effect: if the jamming signal of the figure is slightly shifted in time, the interference may be destructive. If this happens, the received energy is decreased. This is illustrated on Figure 2.14b.


In the last case, the burst of the interfering signal falls during the time position opposite to the value sent by the source. The value demodulated by the receiver depends in this case of the two signals relative intensities.

The exact effect of a collision depends on the receiver type. This section focused on an energy-detection receiver that listens only during the two possible burst positions. However some other receivers estimate the channel during the synchronization preamble and listen during larger time windows in order to exploit the multipath components of the signal. In that case, the probability of interference increases. For correlation receivers, a collision of two pulses in the same window makes the pulse waveform more difficult to detect, while for an energy-detection receiver the same collision can improve the demodulation.

In the case of IEEE 802.15.4A, several types of interference are possible: a



(a) The interfering pulse arrives in phase with the signal, and the received energy is increased. This leads to increased performance when considering an energy detection receiver.



(b) The interfering pulse arrives out of phase with the signal, and the received energy is decreased. This leads to decreased performance when considering an energy detection receiver. The interfering pulse arrives only 8 ps later than in the other case.

Figure 2.14: A pulse collision can lead to positive or negative interference. Both cases are shown here for a simplified 8 ns symbol structure using a 2 ns position modulated pulse. The values sampled by the receiver in the two modulation positions are shown in green and brown points (bottom).

synchronization preamble can interfere with another synchronization preamble on which the receiver is trying to synchronize, or it can interfere with a burst position modulated data sequence that the receiver is demodulating. Two other types of interference are possible: a sequence of bursts can interfere with a synchronization preamble on which the receiver is trying to synchronize, or it can interfere with a burst position modulated data sequence that the receiver is demodulating.

2.6.2 MATLAB Simulations

Research in transceiver architectures and signal processing makes extensive use of MATLAB models. These models use discretized waveforms represented at a high *simulation sampling rate*, which is typically of 10 GHz [54] (100 ps resolution). These transmitted waveforms are convoluted with channel impulse response realizations, and all arriving signals can be summed at the considered receiver.

This approach is the most realistic, but it suffers from a number of drawbacks when considered for network simulations:

- extremely large memory requirements, due to the high simulated sampling rate;
- not scalable with network size and traffic intensity;
- unadapted to the implementation of communication protocols;
- lack of network modeling concepts and tools.

Furthermore, the high accuracy of this approach requires very detailed receiver implementations. This further increases development and simulation times, while leading to potentially too specific results.

2.6.3 The Gaussian Approximation

By assuming that the sum of all interfering signals is a zero mean Gaussian random process, it is possible [40, 170, 116] to derive the bit error rate as a function of the number of currently active users for a correlation receiver and for an Additive White Gaussian Noise (AWGN) channel, for various modulations: Time-Hopping Pulse Position Modulation (TH-PPM), Time-Hopping Phase Shift Keying (TH-PSK) and Direct Sequence Phase Shift Keying (DS-PSK).

This approach, named Gaussian Approximation or Standard Gaussian Approximation (SGA), is simple and fast, and can be adapted to multipath channels [42]. However, it has been shown [4, 41] to over-estimate the performance. In addition, its adaptation to energy detection receivers is not straightforward.

The remaining of this section describes the computation of the bit error rate as a function of the interfering signals by making use of the SGA [40], assuming a pulse position modulation with a modulation shift $T_s/2$ equals to one half of the symbol duration.

Pulse position modulation can be represented as follows at the transmitter:

$$s(t) = \sqrt{E} \sum_j p_0(t - jT_{sym} - \theta_j - b_j T_{sym}/2)$$

where E is the pulse energy, $p_0(t)$ is the normalized pulse waveform, T_{sym} is the symbol duration, θ_j is the time-hopping shift for the considered symbol j , b_j is the j th bit value and $T_{sym}/2$ is the time shift for the modulation.

Considering an AWGN channel of impulse response

$$h(t) = \alpha\delta(t - \tau)$$

where α is the attenuation and τ is the delay, the energy at the receiver is $E_u = \alpha^2 E$. After propagation, the signal becomes

$$r_u(t) = \sqrt{E_u} \sum_j p_0(t - jT_{sym} - \theta_j - b_j T_{sym}/2 - \tau)$$

The received signal can be separated into three components: the transmitted signal from the source transformed by the channel $r_u(t)$, the multiple access interference $r_{mai}(t)$ caused by simultaneous transmissions from jammers (also transformed by the channel) and the thermal noise $n(t)$. The thermal noise is a zero-mean Gaussian random variable of standard deviation $N_0/2$ (where N_0 is the thermal noise given by $N_0 = k_B T$, k_B being the Boltzmann constant and T the absolute temperature) and the multiple access interference can be expressed as follows:

$$r_{mai}(t) = \sum_{n=1}^{N_i} \sqrt{E^{(n)}} \times \sum_j p_0\left(t - jT_{sym} - \theta_j^{(n)} - b_j^{(n)} T_{sym}/2 - \tau^{(n)}\right)$$

where N_i is the number of interfering signals and $E^{(n)}$ and $\tau^{(n)}$ are respectively the received energy and the channel delay for the considered signal and $\theta_j^{(n)}$ and $b_j^{(n)}$ are respectively the time-hopping shift and the bit value for the j th symbol of the considered interfering signal.

The effect of a correlation receiver (correlating the received signal $s_{RX}(t)$ with a correlation mask $m(t)$) can be expressed as:

$$Z(x) = \int_{\tau}^{\tau+T_S} s_{RX}(t) m_x(t - \tau) dt = Z_u + Z_{mai} + Z_n$$

with

$$m_x(t) = p_0(t - xT_S - \theta_j) - p_0(t - xT_S - \theta_j - T_S/2)$$

The decision variable $Z(x)$ is decomposed in three terms Z_u , Z_{mai} and Z_n , respectively representing the signal contribution, the multiple access interference and the thermal noise contribution.

The noise contribution Z_n is Gaussian distributed with zero mean and variance $\sigma_n^2 = N_0(1 - R_0(\epsilon))$, with:

$$R_0(t) = p_0(\xi)p_0(\xi - t)d\xi \quad (2.7)$$

being the autocorrelation function of the pulse waveform.

A maximum likelihood detector is assumed. It works as follows: if Z is smaller than zero, a 1 is decoded. If Z is larger than 0, a 0 is decoded. The

average BER can thus be expressed (assuming randomly distributed bit values) as:

$$\begin{aligned} BER &= \frac{1}{2}P[Z(x) < 0|b_x = 0] + \frac{1}{2}P[Z(x) > 0|b_x = 1] \\ &= P[Z(x) < 0|b_x = 0] \end{aligned}$$

Now, the SGA hypothesis is introduced. We make the hypothesis that Z_{mai} is Gaussian distributed with zero mean and variance:

$$\sigma_{mai}^2 = \frac{1}{T_S} \sigma_M^2 \sum_{n=1}^{N_i} E^{(n)}$$

with:

$$\sigma_M^2 = \int_{-\infty}^{+\infty} \left(\int_{-\infty}^{+\infty} p_0(t-\tau)(p_0(t) - p_0(t-\epsilon))dt \right)^2 d\tau. \quad (2.8)$$

The BER becomes:

$$BER = \frac{1}{2} \operatorname{erfc} \left(\sqrt{\frac{1}{2} \left(\left(\frac{E_u \gamma(T_S/2)}{N_0} \right)^{-1} + \left(\frac{T_S \gamma(T_S/2)}{\sigma_M^2 \sum_{n=1}^{N_i} (E^{(n)}/E_u)} \right)^{-1} \right)^{-1}} \right) \quad (2.9)$$

where erfc and γ are respectively the complementary error function and the complement to one of the pulse waveform autocorrelation function R_0 :

$$\begin{aligned} \operatorname{erfc}(x) &= 1 - \operatorname{erf}(x) \\ &= \frac{2}{\sqrt{\pi}} \int_x^{\infty} e^{-t^2} dt \\ \gamma(t) &= 1 - R_0(t). \end{aligned} \quad (2.10)$$

2.6.4 Characteristic Function

An alternative to the Gaussian Approximation is to compute characteristic functions (CF) [101]. This approach has been applied to DS and TH PPM and Pulse Amplitude Modulation (PAM), and to AWGN and log-normal fading multipath channels. Here again, only the correlation receiver was considered.

The CF approach, however, requires numerical evaluations of some integrals. The precision of the simulation results depends on the precision with which the integrals are evaluated, and this can greatly slow down the simulations.

2.6.5 The Pulse Collision Model

In [62], the authors consider that errors are caused by pulse collisions. They begin by modeling the problem using the approach described in section 2.6.3

on the Gaussian Approximation. After decomposing the decision variable Z in three terms Z_u , Z_{mai} and Z_n representing respectively the signal, the multiple access interference and the noise, they develop an alternative approach to evaluate the probability distribution of the random variable Z_{mai} . They evaluate for each symbol the number of possible pulse collisions (depending on the number of currently active transmissions), the probability of each case and the effect on the bit error rate for each case.

This leads to an analytic expression of the bit error rate for correlation receivers in AWGN channels for pulse position modulation and time hopping coding, as follows:

$$BER \approx Q \left(\sqrt{\frac{E_u}{N_0} \gamma \left(\frac{T_s}{2} \right)} \right) + \sum_{N_C=0}^{N_i} \frac{P_{CP}(N_C)^2}{2} \Omega \left(\frac{E_u}{N_0} \gamma \left(\frac{T_s}{2} \right), \frac{Z_{max}(N_C)^2}{N_0 \gamma \left(\frac{T_s}{2} \right)} \right) \quad (2.11)$$

where

$$\Omega(A, B) = Q(\sqrt{A} - \sqrt{B}) + Q(\sqrt{A} + \sqrt{B}) - 2Q(\sqrt{A}),$$

$$Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^{\infty} e^{-t^2/2} dt,$$

$$P_{CP}(N_C) = \binom{N_i}{N_C} P_{C0}^{N_C} (1 - P_{C0})^{N_i - N_C},$$

$$P_{C0} = \frac{\min(2T_M + T_s/2, 4T_{Burst}, T_s)}{T_s},$$

$$Z_{max} = \sum_{j=1}^{N_i} \left(\left\lceil \frac{N_C - j + 1}{N_i} \right\rceil \sqrt{E_S^{(j)}} \right),$$

and γ is defined in eq 2.10. $P_{CP}(N_C)$ gives the probability of encountering exactly N_C collisions during a symbol period and considering N_i interferers. P_{C0} is the probability that a single interferer causes a collision during a symbol period. Finally, $Z_{max}(N_C)$ is an approximation of the maximum possible value of the multiple access interference term Z_{mai} , assuming N_C realized collisions and with $E_S^{(1)}, E_S^{(2)}, \dots, E_S^{(N_i)}$ being the interfering energies sorted in descending order.

However, it is difficult to introduce multipath channels in this model or to consider energy detection receivers.

2.6.6 Large Deviations and Importance Sampling

In [97], the authors combine two methods, large deviations and importance sampling, for computing the BER of a coherent rake receiver with a correlation detector, using BPSK modulation and with arbitrary multipath channels between the transmitters and the receiver.

Large deviations is a fast technique but makes the assumption that all interferers are small. Therefore the approach taken by the authors is to use the more computationally intensive method of importance sampling for so-called large interferers, and to use large deviations otherwise.

Unfortunately it is not clear if this work can be adapted to BPM modulation or to non-coherent receivers, and if it can consider simultaneously the two types of signals present in a IEEE 802.15.4A frame: the synchronization preamble consisting of isolated pulses and the data part consisting of time modulated bursts. This model is not currently available for network simulation.

2.6.7 Cumulative Noise

The only UWB-IR physical layer model for network simulation publicly available today is described in [96]. It considers BPSK modulation, time hopping coding with variable bit rate, and a deterministic channel model without multipath inspired from the Ghassemzadeh channel model [60].

It associates an average power level to each packet, computes an average noise level from interfering packets during the reception of a packet, and uses lookup tables to convert this average signal to noise ratio into a bit error rate. The data from the lookup tables are derived from MATLAB simulations. This approach is difficult to adapt to position modulation or to energy detection receivers since the lookup tables must be regenerated, requiring access to detailed MATLAB models.

2.6.8 CTU

An analytical framework named CTU [19] aims at evaluating the saturation throughput of MAC protocols for Impulse Radio using pulse position modulation, considering log-normal fading multipath channels.

Its application is restricted to nodes deployed randomly and uniformly on a square region. In addition, this framework requires precise knowledge of the MAC protocol: the probability of being in each of the protocol's possible states (the stationary distribution) must be known. This is difficult to evaluate for today's sophisticated MAC protocols.

2.6.9 Observations on Existing Models

The problem of modeling MAI for UWB-IR systems has been studied extensively using various analytical approaches. Table 2.6 summarizes the characteristics of the existing models (excluding numerical MATLAB models). The following observations can be made:

- all models assume a correlation receiver. Energy detection or super-regenerative receiver architectures have not been considered in the literature.
- Most models use a simple additive white Gaussian noise channel. This does not reflect the reality of the multipath propagation of UWB-IR, which can have an important effect on performance.
- No model has been presented for burst position modulation, which is used by the IEEE 802.15.4A UWB PHY. Many models focus on phase shift

Table 2.6: Existing models of Multiple Access Interference in UWB-IR communication systems.

Models	Receivers		Channels			Modulations			Availability	Complexity
	Correlation	Energy- Det.	AWGN	Log-Normal	Impulse Resp.	PSK	PPM	BPM		
Gaussian Approx. [40, 170, 116]	●		●			●	●			●
Charact. Function [101]	●		●	●						●●●
PCM [62]	●		●				●			●●
Large Deviations and I.S. [97]	●				●	●				●●●
Cumulative Noise [96]	●		●			●			NS-2	●●
CTU [19]	●			●			●			●●

keying modulation. The pulse position modulation models are better candidates for adaptation to BPM.

- Only one of the models ([96]) has been implemented by its authors in a network simulator.
- The models vary greatly in complexity. Sophisticated models tend to be more accurate. Unfortunately their complexity and their mathematical assumptions also make them difficult or impossible to adapt to other receivers, other channels and other modulations.

These observations reflect the fact that modeling analytically the effect of multiple access interference on bit error rate is a challenging problem. The scarcity of hardware platforms makes the models validation difficult.

The selection and adaptation of a model to burst position modulation is not trivial. Only three models actually consider pulse position modulation. Among them, one [19] only focus on the saturation throughput and is not suitable for BER evaluation in other conditions. Furthermore, it can not be implemented in a network simulator. The two candidates left are the Gaussian Approximation [40, 170, 116] and the Pulse Collision Model (PCM) [62] by Giancola and Di Benedetto. The GA approach has been shown to greatly overestimate the system performance [4, 41]. And while the PCM approach is promising, it fails to capture the multipath characteristics of the UWB propagation channel.

Therefore, research on communication protocols for IEEE 802.15.4A UWB-IR networks lacks an adequate PHY layer model for simulation.

2.7 Maximum Pulse Amplitude Estimation

2.7.1 Overview

This section describes a flexible ultra wideband simulation model based on the Omnet++ discrete event simulation engine [165] and on the MiXiM simulation framework [81].

The Omnet++ discrete event simulation engine provides a core modeling library, text-mode and graphical front-ends to interact with the simulation during its execution, and data collection and analysis tools to study the results of a simulation run. It is used for wired network simulations, queuing networks, business processes modeling, etc... It can be used for any system that can be decomposed in *modules* that send each other *messages* through *gates* at discrete points in time. A message scheduled (in simulation time) to arrive at a destination module is called a simulation *event*. Modules can also send messages to themselves; this is called a *self-message* or a timer.

The core simulation engine creates the modules as specified by the user through configuration files, initiates the simulation time to zero, delivers the scheduled messages to their destination modules and collect statistics as specified by the user in the simulation model source code. The simulation stops once there are no more events or if the simulation time limit is reached.

The MiXiM simulation framework is a library of Omnet++ simulation models for wireless and mobile networks. It enhances the basic Omnet++ model building blocks with a set of modules that manage nodes mobility, keep track of the nodes wireless connectivity and enable detailed modeling of physical layers. This allows to focus on the physical processes instead of introducing novel concepts in the simulator as it must be done with NS-2 [96] because of software coupling problems between the simulation engine and the simulation model [165].

In MiXiM, a network host is composed of one or more network interface card modules, a mobility module, and one or more higher layer modules: routing, transport and application. The NIC module is composed of PHY and MAC modules that interact closely with each other. The MAC module receives messages from the upper layer (routing or application), encapsulates each message in a MAC message object, sets some MAC layer parameters (source and destination addresses, header length, frame type), generates a MiXiM signal object representing the physical signal to be transmitted, and forwards this to the PHY layer. The PHY module generates MiXiM AirFrames (a subclass of Omnet++ cMessage) that are handled by the MiXiM channel control module for duplication and delivery to all hosts within maximum interference distance of the source. Most of the PHY module concerns the reception of frames: it permanently tracks the AirFrames currently on air that can potentially interfere with it, it applies channel models to the incoming transmitted signals to obtain the received signals and it simulates the synchronization, reception and error correction processes.

Most wireless network simulators only store the packet size, the start and end times of the frame and the transmit power level. Knowing this informa-



Figure 2.15: A few pulses of the IEEE 802.15.4A synchronization preamble as stored in the network simulator at the transmitter.

tion for all frames intersecting with the frame currently being received, it is straightforward to derive the Signal to Noise Ratio of the received frame at any point in time (see Figure 2.13). Analytical expressions can then be used to derive the bit error rate probability for all SNR values, and the receiver can compare randomly generated values with these probabilities to decide whether the frame was correctly received or not. This is not feasible for UWB-IR.

Instead of implementing an analytical closed-form model of bit error rate, the approach taken here is to model every single pulse of each symbol. While this increases memory and processing time requirements, it offers the following advantages:

- decoupling of the modulation, channel model and receiver architecture;
- distinction between the data and synchronization signals;
- possibility to implement UWB ranging algorithms;
- possibility of considering interference from all types of narrowband and UWB signals.

This section discusses the fundamental assumptions of this approach and the overall software design. The following section describes in more detail how a model of a IEEE 802.15.4A UWB PHY energy-detection transceiver was implemented.

2.7.2 Assumptions

The simulation model makes the following assumptions: a channel coherence time larger than the maximum packet duration, no interference from other radio technologies, maximum pulse amplitude estimation and triangular pulse envelope shapes.

1. The channel coherence time is larger than the packet duration: the parameters characterizing the channel are randomly generated at the beginning of the packet reception, and are not modified during the reception. New channel parameters are generated at each packet arrival. The maximum data payload size is 127 bytes, for a maximum packet duration of 1.65 ms (including the synchronization preamble). As far as we know, all UWB simulations (MATLAB implementations) make this hypothesis. See subsections 2.7.6 and 2.3.2 for details.
2. Interference from other UWB systems such as MB-OFDM-UWB or FM-UWB, or from narrowband systems such as IEEE 802.11 or UMTS, are not taken into account. This work does not address such problems.
3. Maximum Pulse Amplitude Estimation: to allow modeling every pulse without resorting to MATLAB-like high sampling rates, only the maximum possible pulse amplitude is stored.
4. Triangular pulse envelopes: by storing only the pulse envelope, memory requirements are reduced. Intermediate points of the pulse envelope are easily obtained by linear interpolation when considering triangular pulses. Triangular pulse envelope has been implemented simply and at low power consumption [132]. See sections 2.8.2, 2.7.6 and 2.7.4 for details.

The key aspect of our approach lies in the modeling of the signal waveform. Instead of working at a high simulation sampling rate, similarly to MATLAB models, we consider only the pulse envelope. This greatly reduces the memory requirements: from 20 to 3 points for a 2 ns pulse. And by using the MiXiM Mapping data structure (see section 2.7.4), linear interpolation is performed on demand. Thus, when the signal is constant over a period of time, only two points must be stored. Assuming triangular pulse envelopes, this structure provides enough information to derive its MATLAB waveform representation at the transmitter side. Figure 2.9 illustrates a few pulses represented with this approach. However, when considering a Power Delay Profile channel model, or multiple access interference, approximations must be made. As shown on Figure 2.14, a slight shift in time between two colliding pulses can lead to dramatically different results for the final signal. While the MATLAB approach simply sums the signals to obtain the resulting signal, the approach adopted here does not allow this. Instead, an approximation must be made. We considered three options:

1. Best case: preserving the strongest pulse and ignoring the others;
2. Worst case: always assume destructive interference;
3. Maximum Pulse Amplitude Estimation: store the maximum possible pulse amplitude and let the receiver generate random values from this information.

The first approach has the advantage of simplicity and speed. The second approach leads to overly pessimistic results. The third approach was chosen in this work. The actual result depends on how this information on the maximum possible pulse amplitude is used by the receiver model: it can

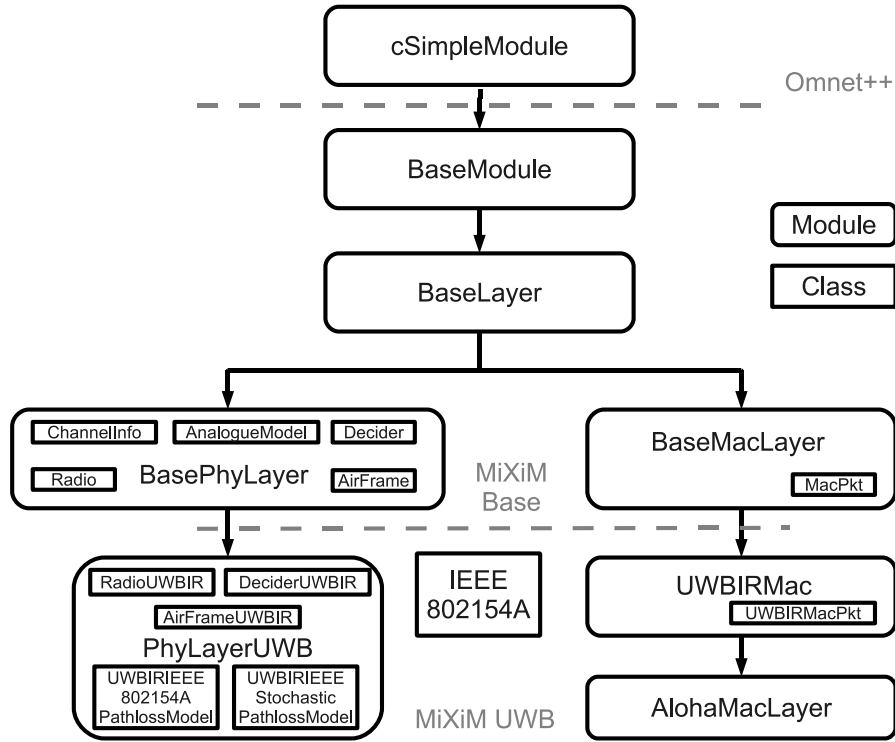


Figure 2.16: UML inheritance diagram representing the relations between Omnet++, MiXiM and MiXiM UWB classes.

consider that the stored values are the actual measured values, or it can assume that the actual pulse amplitude is randomly distributed in the interval $[-MaximumAmplitude, +MaximumAmplitude]$, where *MaximumAmplitude* is the value returned by the Signal mapping.

2.7.3 Software Architecture

The MiXiM simulation framework [81] provides good foundations for implementing detailed PHY layer models. Several MiXiM classes have been reused to implement our model. They are described below. Figure 2.16 provides an UML inheritance diagram, with Omnet++ modules represented by rounded boxes and non Omnet++ modules (simple C++ classes) represented by rectangles.

The MiXiM BasePhyLayer class has been subclassed by a PhyLayerUWBIR, and the BaseMACLayer class has been subclassed by a UWBIRMac class. PhyLayerUWBIR is configured from an Omnet++ initialization text file, and instantiates the requested channel and receiver models. UWBIRMac provides features that can be useful to all UWB-IR MAC. A IEEE802154A class contains all the standard-specific functions and allows to generate a MiXiM signal modulating bit values chosen either by the user or random.

The following subsection 2.7.4 describes the data structures used to represent radio frames and physical signals, and subsection 2.7.5 presents the objects

provided by MiXiM to model a wireless PHY and how they have been used in our work.

2.7.4 Data Structures

MiXiM defines the classes `AirFrame`, `Signal` and `Mapping` to represent respectively the radio frame, the physical signal and the actual values taken by the physical signal over time. More precisely, the `Signal` object stores a representation of the emitted signal using another `Mapping` object, the signal beginning and duration, the position and movement of its source during its emission and the bit rate information (using a `Mapping` object). The movement information is not considered in this implementation, as we assume that the node speed is negligible compared to the signal propagation speed and packet duration.

The `Mapping` class is used to represent functions, or mappings from \mathbb{R}^n to \mathbb{R} . In the simplest case, signal values only depend on time and the mapping is from \mathbb{R} to \mathbb{R} . This is the case in our implementation and the `TimeMapping` subclass of `Mapping` was used to this end. Mapping values can be read by iterating over all stored data points. It is also possible to get the value of an arbitrary point. If no data value is stored in the `Mapping` for that point, the `Mapping` will use an interpolation method to compute the result (MiXiM provides several interpolation methods in the `Mapping` class).

In our case, each pulse envelope is stored with only three points, marking the beginning, the peak and the end of the triangular pulse. The amplitude of the pulse envelope at any point in time is computed by linear interpolation.

2.7.5 Physical Layer

MiXiM `BaseMacLayer` is an Omnet++ module (class `cSimpleModule`) that can exchange Omnet++ messages (class `cMessage`) through two gates: the data gate for `MacPkt` data messages and the control gate for exchanging control messages (such as `ChannelSenseRequest` to estimate the current channel conditions). In addition to these two Omnet++ connections, the MAC can also directly call `BasePhyLayer` methods through the `MacToPhy` interface. This is useful to control the radio for instance.

When sending a packet, the following sequence of events takes place. First, the MAC layer must create the `Signal` object corresponding to the `MacPkt` to send. It is attached to the `MacPkt` through a `MacToPhyControlInfo` data structure. Upon reception of this `MacPkt`, if the radio is in the correct state (transmission), `BasePhyLayer` *encapsulates* (an Omnet++ operation that allows to cleanly isolate data from the application, transport, routing, MAC and PHY modules) the `MacPkt` in an `AirFrame` and associates the `Signal` object to this `AirFrame`. Then, `BasePhyLayer` uses its `ChannelAccess` interface to let MiXiM take care of delivering copies of this `AirFrame` to all reachable nodes (depending on a maximum interference distance).

The reception process is modeled in MiXiM as follows. Each reachable node receives an `AirFrame` from `ChannelAccess`. First, the `AirFrame` is forwarded to the local instance of `ChannelInfo`. This object is used by `BasePhyLayer` to keep track of all ongoing transmissions (from the node viewpoint). An `AirFrame` is removed from `ChannelInfo` once the simulation has progressed enough so that the `AirFrame` does not interfere with any currently active `AirFrame` (an

AirFrame is active if its end time is in the future when compared to the current simulation time). BasePhyLayer then applies to each arriving AirFrame, all loaded AnalogueModels. In our case we use only one AnalogueModel to represent all channel effects (see subsection 2.7.6).

If the radio is in the required state (SYNC, see subsection 2.7.7), BasePhyLayer gives the received signal to the Decider object. The decider will first attempt to synchronize on the incoming signal (see subsection 2.8.3 for details on modeling synchronization algorithms). It may ask the ChannelInfo object for the list of all currently active AirFrames to evaluate interference. If the synchronization is successful it will change the radio state (through the *DeciderToPhy* interface) to reception mode. Until the end of this frame, all other arriving AirFrames will be treated as noise.

At the end of the frame, BasePhyLayer gives again the received signal object to the Decider. The Decider decides whether or not it could successfully receive the frame (depending on the received signal and on the interfering signals). In case of success, BasePhyLayer decapsulates the AirFrame, attaches a PhyToMacControlInfo structure to it and sends it to the MAC for further processing.

We subclassed BasePhyLayer with a *PhyLayerUWBIR* class. It instantiates the channel model and the decider chosen by the user, and initializes the UWB-IR model. The UWB-IR decider demodulates the signal, but it needs the original bit values to evaluate the bit error rate. They are not available to it because the transmitted signal was generated (or modulated) by the MAC layer. Thus, contrarily to MiXiM intended design, the frame validation is not made at the PHY layer in the Decider but at the MAC layer. The demodulated bit values are stored and transmitted to the MAC layer (in a PhyToMacControlInfo object), which compares them with the original (randomly generated or not) bit values stored in the UWBIRMACPacket. Error correction is also implemented at this level.

2.7.6 Channel

A channel model can either provide simply a pathloss value, or additionally include a Power Delay Profile that represents the channel impulse response.

All channel effects are represented in MiXiM as so-called *AnalogueModels*. There can be more than one such model active at the same time; they are initialized and stored by the *BasePhyLayer*, and they generate *Mappings* that, multiplied with the Mapping representing the original transmitted signal, give the received signal. For a narrowband signal, one AnalogueModel can represent the pathloss, another the shadowing and a third one the small scale fading. While this approach is flexible and allows to activate and deactivate channel effects to study their effect on the system's performance, it is not well suited to the introduction of individual multipath components such as generated by a PDP. We adopted the following strategy: as the AnalogueModel has access to the whole Signal object, it can not only adds an attenuation Mapping to represent the pathloss, but also completely redefine the transmitted Mapping.

We implemented most channel models specified by the IEEE 802.15.4A working group [100]: Residential Line of Sight (LOS) CM1, Residential Non Line of Sight (NLOS) CM2, Indoor office LOS CM3, Outdoor LOS CM5 and NLOS CM6 and Open outdoor NLOS CM7. These models are described in




Figure 2.17: The UWB-IR radio state machine with an additional Sync state.

detail in subsection 2.3.4. Assumption 4 of subsection 2.7.2 was used to estimate and store the maximum pulse amplitude from the PDP. In addition, the simpler and faster stochastic channel model proposed by Ghassemzadeh [60] and also used in [96] has also been implemented (see subsection 2.3.2). This allows to run simulations faster while still keeping a good degree of precision.

2.7.7 Radio

MiXiM provides a detailed four states radio model: *Idle*, *Reception*, *Transmission* and *Switching*. This last state allows to take into account transition states such as when the radio is setting up into reception mode. This is an important property of a radio: since radios are half-duplex, they cannot receive and transmit at the same time. The time to switch between reception and transmission, called *turnaround time*, can cause serious performance degradation. Indeed, two radios attempting to send a packet to a third node can both detect an idle channel and try to send at about the same time, leading to collisions. A shorter turnaround time reduces this vulnerability window.

The four states defined by MiXiM are enough to build detailed narrowband radio models.

However, the difficulty of detecting isolated UWB pulses makes UWB receivers particularly power consuming when searching for a synchronization preamble. Therefore it is necessary to track precisely the time spent in that mode. This is why, similarly to [96], we introduced a fifth radio state, *Synchronization*, preceding *Reception*. The evaluation of the radio power consumption is done in an *RadioUWBIR* subclass of *Radio*, by counting the time spent in each state of the radio and using estimates of the power consumption in each state. The finite state machine of this new radio is illustrated on figure 2.17.

Table 2.7 summarizes all timing and power consumption parameters of the radio model. These estimates are taken from UWB-IR [133, 111, 85] and narrowband [160, 159] transceivers.

2.8 A Model of the IEEE 802.15.4A UWB PHY

This section describes the first IEEE 802.15.4A UWB PHY network simulator, based on the Maximum Pulse Amplitude Estimation (MPAE) approach developed in section 2.7. Subsection 2.8.1 discusses the additional assumptions made for this implementation, subsection 2.8.2 describes the transmitter model, and

Table 2.7: UWB-IR model parameters (sources: [133, 111, 85, 160, 159]).

Parameter	Symbol	Value
Sleep power consumption	P_Z	10 μ W
Reception power consumption (demodulation)	P_{Rx}	30 mW
Reception power consumption (synchronization)	P_{sync}	45 mW
Transmission power consumption	P_{Tx}	1 mW
Transmission setup time (from sleep mode)	$T_{SetupTx}$	100 μ s
Reception setup time (from sleep mode)	$T_{SetupRx}$	100 μ s
Switching time (reception to transmission mode)	T_{SwRxTx}	100 μ s
Switching time (transmission to reception mode)	T_{SwTxRx}	100 μ s
Shutdown time (reception/transmission to sleep mode)	T_{off}	30 μ s

subsections 2.8.3 and 2.8.4 describe respectively the synchronization model and the demodulation process of the receiver model.

2.8.1 Assumptions

To simplify the implementation, the following assumptions were made: no clock drift during a frame reception, uniform random data bits, non-coherent energy-detection receiver and a simple synchronization logic.

1. No clock drift: since each system has its own clock, in practice clock drift is unavoidable. The synchronization preamble of a packet allows the receiver to synchronize on the transmitter clock. Without clock drift, after synchronizing on the preamble, the synchronization remains accurate for the entire packet duration: the energy detection receiver always samples energy perfectly at the peak of the considered pulse. Since the packet sizes for sensor networks are small, this hypothesis is reasonable. Sophisticated receivers use special techniques (pulse tracking for coherent receivers, Radon transform for energy detection receivers [53]) to compensate the clock drift. This hypothesis only concerns the PHY layer model and does not prevent modeling clock drift at the MAC layer.
2. Random bit values: the PHY layer generates random bit values using a uniform random number generator.
3. Energy detection receiver: the low complexity of these receivers is especially attractive for sensor networks, for which low cost and low power consumption can be more important than absolute signal-to-noise performance. In addition, this demonstrates the soundness of the symbol-level

approach, that enables for the first time to consider this receiver architecture in network simulation.

4. Synchronization: since UWB-IR does not transmit continuously, the signal is difficult to detect. Interferences during synchronization can prevent successful detection. A randomized approach was taken in [96]. In this work, the radio synchronizes on the first frame whose preamble is not jammed for a certain duration and if the signal is above a threshold. See subsection 2.8.3 for details.

2.8.2 Transmitter

The generated signal implements the Burst Position Modulation of the standard, using the mandatory mode: 2 ns 500 MHz pulses, 16 pulses per burst, 1025 ns per symbol. Power consumption estimates are taken from [133], and pulse envelopes have a triangular shape. The synchronization preamble uses the default duration of 71.5 μ s. The Reed-Solomon error correction code required for standard compliant transmitters has been modeled by adding the appropriate number of symbols (48) at the end of the frame, modulated with random bit values (see subsection 2.8.4 for the error correction model at the receiver side).

The Signal object is generated by the MAC layer. A convenience function *prepareData* is provided in *UWBIRMac*: it calls the appropriate functions of the class *IEEE802154A* to generate a signal modulating the desired number of random bit values. The randomly generated bit values are stored in the *UWBIRMacPkt* to allow packet validation at the receiver. The Mapping data structure generated at the transmitter can be understood both as a representation of the maximum possible value at any given time t of the voltage applied at the antenna, of the strength of the generated electric field since they are directly proportional to each other, and of the signal intensity I , provided a distance information d and $S_{Sphere}(d)$ being the surface of a sphere of radius d :

$$I = \frac{P_{radiated}}{S_{Sphere}(d)} = \frac{1}{4\pi d^2} \frac{v^2_{Tx}}{R_{Ant}} = \frac{E^2(d)}{120\pi}. \quad (2.12)$$

2.8.3 Synchronization

The synchronization preamble of the IEEE 802.15.4A is generated and stored at the beginning of the TimeMapping object (see subsection 2.7.4) that represents the signal.

The synchronization logic is implemented in the *bool DeciderUWBIREDD::attemptSync(ConstMappingIterator* mIt)* method. It returns *true* if the synchronization was successful, and *false* otherwise. Other synchronization algorithms can be implemented by subclassing the decider and simply overloading this method. This was done for the *DeciderUWBIREDSync* and *DeciderUWBIREDSyncOnAddress* classes, which both derive from *DeciderUWBIREDD*. The three implemented synchronization models are described below.

The first synchronization model (implemented in *DeciderUWBIREDD::attemptSync()*) succeeds only if there are no interferers and if the signal power level is high enough.

The second synchronization model (*DeciderUWBIREDSync::attemptSync()*) considers all simultaneously interfering signals. If there are none, the synchronization succeeds if the signal power level is high enough. If there are interferers, the synchronization algorithm looks for an interference-free time interval during which the considered synchronization preamble is not jammed. If such a time interval is found, if it is long enough and if its power level is high enough, the synchronization succeeds. The threshold values are declared as module parameters and can be defined at run-time by the user.

The third synchronization model (*DeciderUWBIREDSyncOnAddress::attemptSync()*) succeeds if and only if the incoming frame has been sent by a specific node. This can be useful in some specific cases when the effect of the synchronization must be avoided in order to evaluate some particular aspect of the system. It can also be used to simulate other UWB modulations, for instance when using transmitter-based time hopping sequences.

Finally, the *DeciderUWBIREDSync* also offers an option to deactivate the call to *attemptSync()*. This allows to compare quickly the simulation results with and without a synchronization logic.

2.8.4 Demodulation

An energy-detection receiver has been implemented. For each data symbol, it considers the two burst positions coding respectively a zero and a one. The receiver does not attempt to collect the energy of the multipath components outside the two modulation positions. However, the rays arriving inside the integration window are considered.

In the particular case of the energy detector, two types of interference may happen: *negative interference* when the interfering pulses fall in the window coding the opposite value (e.g. coding a 1 when the transmitter sent a 0), or when the interfering pulses fall in the window coding the same value but with a time shift leading to destructive interference, and *positive interference* when the interfering pulses fall in the window in which the pulses of the transmitted signal have been sent with a constructive interference amplifying the original signal (see subsection 2.7.1 and figure 2.14).

At each pulse peak position in the burst, the receiver samples the energy on the channel and squares it. The data structures representing the maximum possible values of the signal and of the interferers instantaneous radiated powers, modified by the channel model, are used to estimate the resulting radiated power around the antenna. This value is then converted to an induced voltage value at the antenna. This is done as follows.

At each estimated pulse peak time position, each arriving signal is considered. The power level of this signal is estimated to be one half of its maximum possible value at the observed time if it is the tracked signal. This is a simplifying hypothesis to account for the sometime negative and sometime positive interference from the multipath pulse components (see subsection 2.7.6), while at the same time considering that if the receiver could synchronize on this signal then the energy cannot be too low. Hence the multiplying factor 0.5 instead of a completely random value.

If the considered signal is from an interferer, its maximum possible power level is multiplied by a random number taken between -1 and +1, again to account for a possibly destructive interference. All these power levels (one

from each concurrent signal) are summed to estimate the resulting power level at the antenna.

Since $P_r = \frac{V^2}{R_{Ant}} = \frac{V^2}{50}$ (with an antenna impedance set to 50Ω), we can deduce the induced voltage:

$$V = \sqrt{\frac{1}{50} \frac{E^2}{120\pi} \frac{\lambda^2}{4\pi}}. \quad (2.13)$$

At this point, a thermal noise value v_{noise} is generated using a normal random variable with variance $\sigma^2 = 4k_BTRB$, where k_B is Boltzmann constant, T is the absolute temperature (set to 293 K), $R = R_{Ant} = 50 \Omega$ is the resistor value and $B = 500 \text{ MHz}$ is the signal bandwidth. Finally, the signal voltage V is summed to the random noise voltage v_n and this sum is squared.

This process is repeated for the 16 pulse positions of a burst, and for the two possible burst positions. The sums of all squares for each burst position are compared, and the bit value associated to the burst position with the highest energy level is decoded.

If the received signal is low, the random values taken by the thermal noise in the window coding the opposite bit value can be higher than the received signal combined to the thermal noise random values in the window coding the correct bit value. If this happens to be the case, erroneous bit values are demodulated. The demodulation process has been implemented in

DeciderUWBIRED::decodePacket(). The demodulated bit values are forwarded to the MAC layer (*UWBIRMac*) to be compared with the original modulated bit values, stored in the encapsulated *UWBIRMacPacket*. A packet is discarded if it has at least one bit error when error correction is disabled, or if the number of errors and their positions in the frame make them impossible to correct when the Reed-Solomon error correction model is enabled. As the R-S code $RS_6(n = 63, k = 55)$ used in the standard divides a frame in blocks of 6 bit symbols, and allows to correct up to $\frac{n-k}{2} = 4$ 6-bit symbols, the simulation models counts the number of erroneous symbols in the frame (a symbol is in error if at least one of its 6 bit values is incorrect) and validates the frame if this number is lower than the maximum number of symbol errors that can be corrected.

A more realistic energy-detection receiver would instead integrate each pulse and sample at the end of the pulse, or even integrate during the whole burst duration and sample at the end. This approach was initially favored. However, it greatly slows down the simulations. Finally, simplicity and simulation speed were favored at the expense of receiver realism. The following section examines the validity of our approach.

2.9 Validation of the Simulation Model

The lack of UWB-IR transceivers compliant with the IEEE 802.15.4A standard made the validation of this work difficult. Therefore, we resorted to comparing our simulation results with MATLAB simulations. While the basic receiver structure considered in this work was not studied in the literature, Flury et al. [54] worked on the architecture of energy detection receivers for the demodulation of IEEE 802.15.4A signals. It was possible to reproduce their simulation

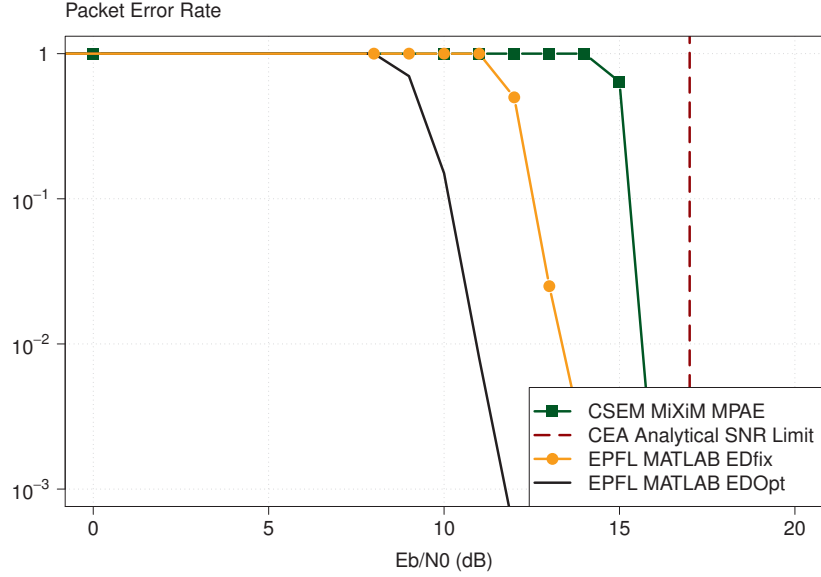


Figure 2.18: Without interferer, the performance of the MiXiM MPAE model is very close to the EPFL MATLAB energy detection model with fixed integration windows, and in line with the analytical link budget estimate.

setup in our simulator. Their receiver is referred further in this text as the EPFL receiver.

The comparison evaluates the packet error rate as a function of E_p/N_0 . E_p is the pulse power (whose maximum value is limited at 1 mW by the peak EIRP regulation) and N_0 is the thermal noise in the receiver. The thermal noise power is obtained with the following relation:

$$P_{noise} = k_B T B \quad (2.14)$$

where $k_B = 1.338E-23 J/K$, $T = 293 K$ is the ambient temperature in Kelvin, and $B = 500\text{MHz}$ is the signal bandwidth. In our simulator, these values were obtained as follows. The thermal noise was estimated by squaring the generated thermal noise voltage values v_{noise} and dividing them by the receiver resistance ($R_{Ant} = 50 \Omega$):

$$P_{noise} = \frac{v_{noise}^2}{R_{Ant}}. \quad (2.15)$$

The values obtained were equal to the analytically computed noise power, validating the noise model implementation.

The pulse power at the receiver was obtained by multiplying the pathloss values of the channel models with the transmitted pulse power equal to 1 mW. The E_b/N_0 ratio was changed by increasing the distance between the source and the destination nodes. The packet size is 127 bytes (maximum allowed by the standard) and the channel model is the IEEE 802.15.4A CM2 channel model. All simulation results use the Reed-Solomon error correction.

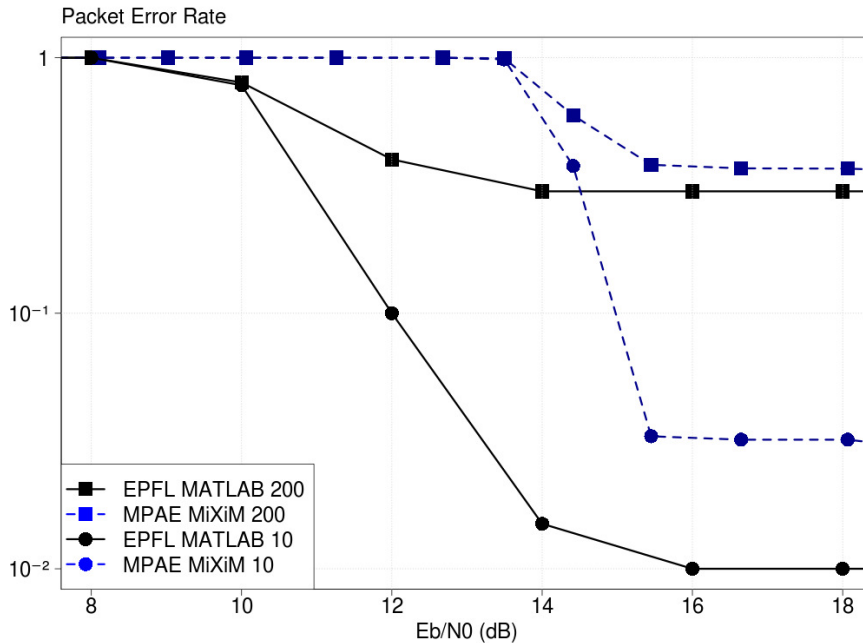


Figure 2.19: Packet error rate with two traffic sources, for two traffic intensities (200 and 10 packets per second and per source). The MPAE MiXiM simulator shows on average a 3 dB performance penalty over the EPFL MATLAB simulator to obtain similar packet error rates.

The EPFL receiver uses large integration windows to capture almost all of the arriving energy. It adapts to the channel by dividing the integration windows into bins. The integration result from each bin is weighted according to the channel estimate, so that bins containing more energy are given more importance. This allows to reduce the receiver sensitivity to thermal noise, by diminishing the importance of the integration samples with low SNR values. In addition, the result of a link budget estimate [104] has been included in the comparison.

The results are shown on figure 2.18. The EPFL receiver performance is shown for two cases. In the first case, its optimal performance is shown in black, with optimal weighting and no clock drift. The second case, shown in orange and with diamond symbols, shows the performance of this receiver with fixed weights and no clock drift correction. A red dashed line shows the performance limit obtained analytically using a conservative link budget evaluation. Our simulator performance is shown in blue circles (with error correction). As expected, the optimal receiver from EPFL performs the best. Its performance degrades when considering E_b/N_0 values around 12 dB. The second configuration of the EPFL receiver is close to this optimum, with performance degradation beginning at 14 dB. The CEA-LETI performance analysis estimates that approximately 17 dB are required for correct operation (in [56], another link budget estimate, considering another channel model, leads to a minimum E_b/N_0 value of 19 dB). The performance of our receiver effectively

begins degrading at about 16 dB, and reaches the maximum packet error rate at 14 dB. Compared to the EPFL energy-detection receiver with fixed integration window, the performance penalty is approximately 3 dB. It is close to 5 dB when compared with the EPFL receiver using optimal integration window with weighted bins.

Others [79, 171] have focused on the raw bit error rate, and considered other channel models. In [79], it was found that to reach a 1% BER, E_b/N_0 had to be higher than 20 dB for a S-Rake 10 receiver on an IEEE 802.15.4A CM8 channel. In [171], an optimal coherent receiver is considered, sampling the pulses at the Nyquist rate. This optimal architecture requires a E_b/N_0 of 8 dB to get a BER of 1%.

Figure 2.19 illustrates another comparison between our MPAE MiXiM-based simulation model and the EPFL IEEE 802.15.4A energy-detection MATLAB model. This time, two nodes periodically generate packets and send them to a third node (using the IEEE 802.15.4A Aloha backoff algorithm as specified in [54]). Both senders generate packets with the same transmitted power, and the simulations are studied with two packet rates: $\lambda = 200$ packets per second and per node, and $\lambda = 10$ packets per second and per node. When considering the low data rate case, if the signal is high enough, most packets arrive to their destination for both the MATLAB and MiXiM simulators, with respectively 1 and 4% packet error rates. The E_b/N_0 value at which this transition happens is 14 dB for the MATLAB receiver and 15 dB for the MiXiM simulator. A similar transition can be observed when considering the high packet rate. However, in that case the remaining packet error rate is still high, between 30 and 40% for both receivers. The points at which the PER begins to increase are the same than in the low data rate case, and correspond to the receiver performance analysis without interference illustrated on Figure 2.18.

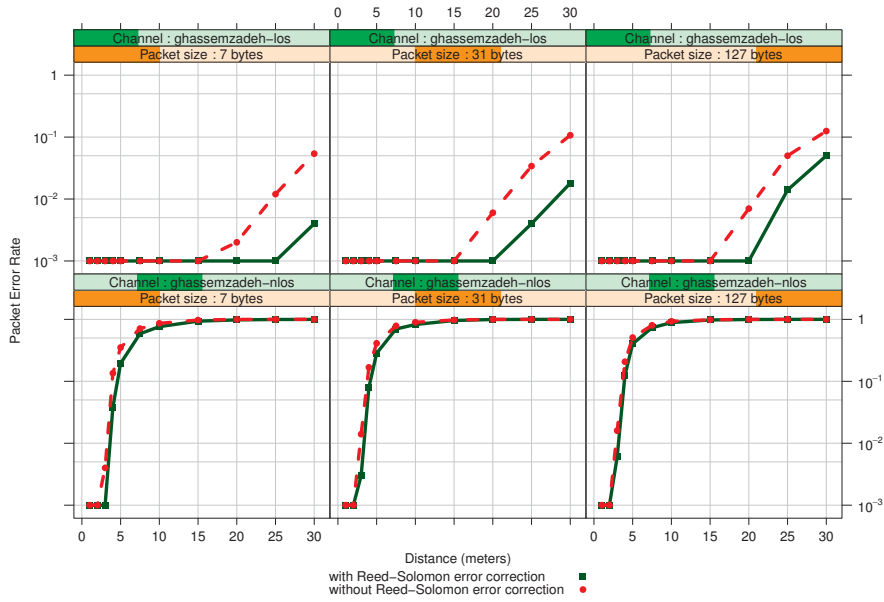
Overall, these results show that while our receiver is not the best, it behaves similarly to more detailed MATLAB implementations. The results obtained are also in line with an analytical link budget.

2.10 Evaluation of the IEEE 802.15.4A UWB-IR PHY

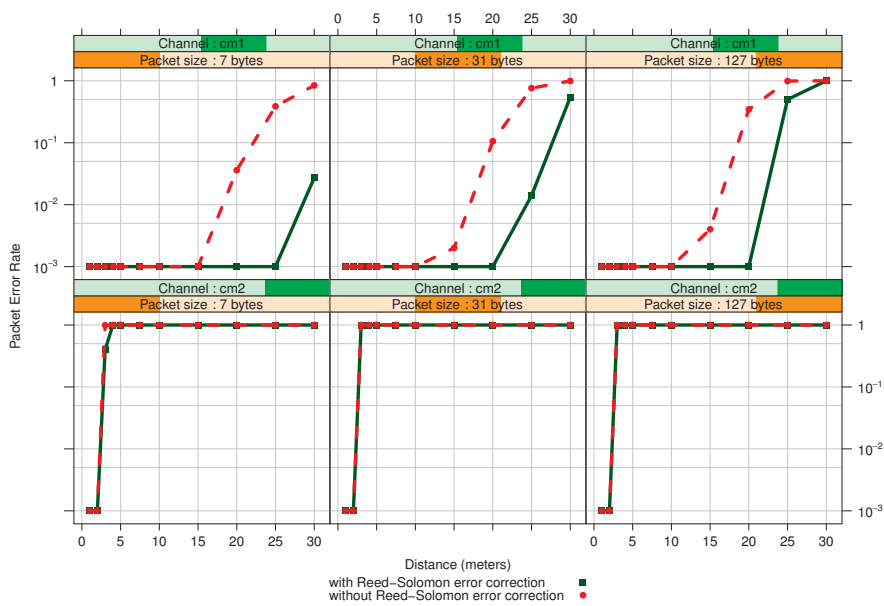
This section presents network simulation results considering various configurations. Subsection 2.10.1 focuses on the effect of the channel model choice, by comparing the packet error rates obtained with various channel models as a function of the distance between the source and the destination. The packet size varies between small (7 bytes), medium (31 bytes) and large (127 bytes, the maximum allowed by IEEE 802.15.4A).

Subsection 2.10.2 considers the effect of multiple access interference on the packet error rate. It focuses on two types of interference: burst on burst and sync on burst. Thus it is the bit error rate that is really considered, at the exclusion of the robustness of the synchronization logic. Four channel types, two packet sizes and two traffic types are considered: continuously transmitting interferers and Poisson traffic.

Subsection 2.10.3 discusses the execution speed of the model, the main factors that affect it and how to configure a simulation so that it runs faster.



(a) Ghassemzadeh LOS and NLOS channels.



(b) IEEE 802.15.4A channel models CM1 and CM2.

Figure 2.20: Packet error rates with and without Reed-Solomon error correction for an energy detection receiver (synchronization always succeeds). Application payloads of 7, 31 and 127 bytes.

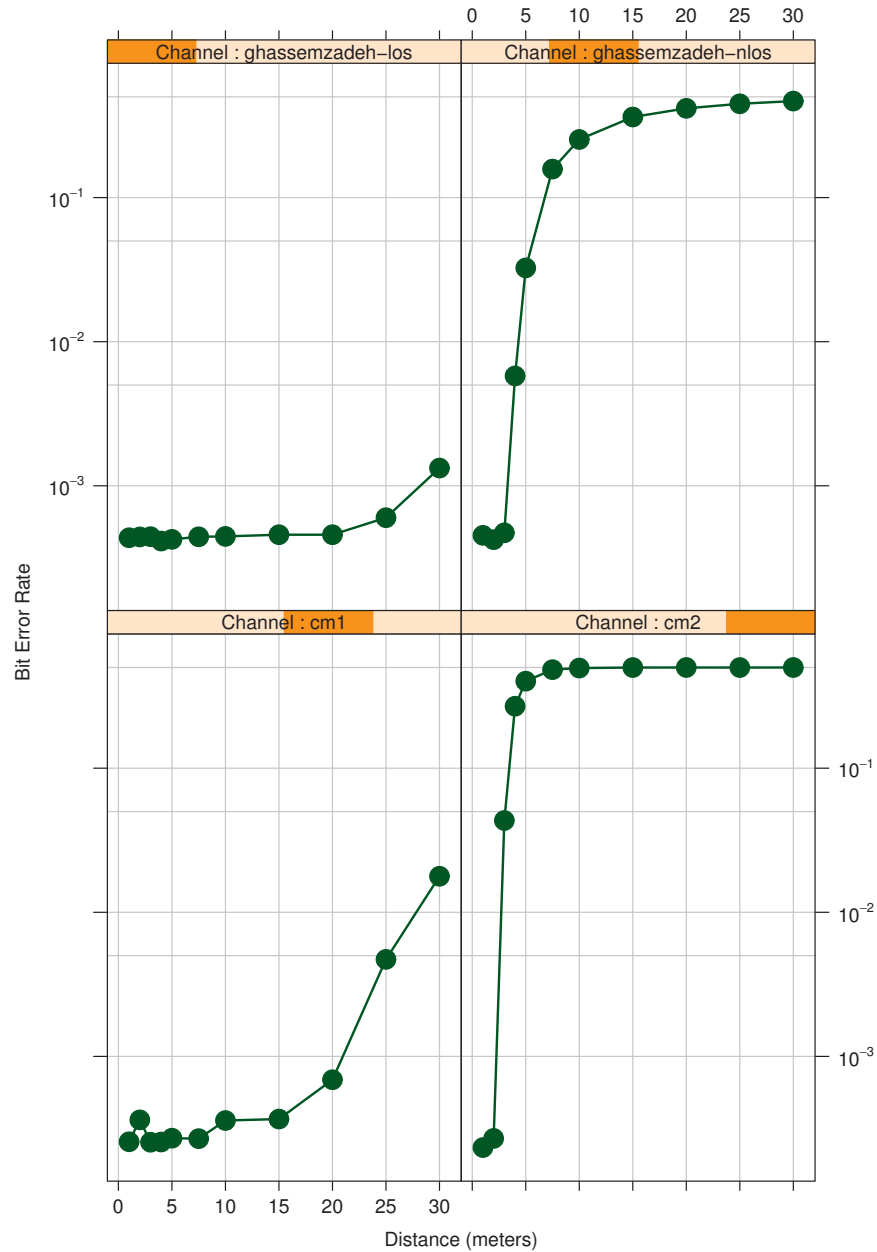


Figure 2.21: Bit Error Rate as a function of the link distance with channel models Ghassemzadeh LOS, Ghassemzadeh NLOS and IEEE 802.15.4A CM1 (residential LOS), CM2 (residential NLOS) and CM5 (Outdoor LOS). Data aggregated from numerous simulation runs with various packet sizes (7, 31 and 127 bytes).

2.10.1 Channel Models

Figure 2.21 shows the bit error rate as a function of the distance between a source and a receiver. The receiver uses energy detection. Four channel models are evaluated: the Ghassemzadeh Line of Sight and Non Line of Sight models, and the IEEE 802.15.4A residential LOS (CM1) and residential NLOS (CM2) models.


We observe a strong performance degradation for the NLOS channels when compared to their LOS counterparts, for both channel model families. In NLOS conditions, the system is unusable for distance larger than five meters, while LOS conditions allow communications over 30 meters. These differences are caused mainly by the difference in pathloss exponent, with values of 3.5 and 4.58 in NLOS conditions for respectively the Ghassemzadeh and IEEE NLOS models, and 1.7 and 1.79 for the LOS models. Thus, the considered NLOS models degrade the signal energy more strongly, to the point that the receiver thermal noise causes erroneous demodulation.

Figures 2.20a and 2.20b provide more information for respectively the Ghassemzadeh and the IEEE models by displaying the packet error rate as a function of distance for three packet sizes: 7, 31 and 127 bytes, with and without Reed-Solomon error correction. While the error correction mechanism cannot correct the packets received on the NLOS channels, due to too many bit errors and a too low signal to noise ratio, its effect can be observed in the case of LOS channels. As the packet size increases, the PER improvement due to the R-S code decreases. This is an expected result as for an identical BER, longer packets are more likely to have 5 or more 6 bit symbols in error, simply because they are composed of more symbols.


2.10.2 Multiple Access Interference

Figures 2.22a and 2.22b show how the packet error rate is influenced by data demodulation errors caused by simultaneous transmissions, for various channel models and two packet sizes (7 and 31 bytes payloads). To isolate the effect of data symbol demodulation errors, an ideal synchronization model, that always successfully synchronizes on the first arriving packet independently of its signal strength or of any ongoing interference, was used. The simulation setup is the following: all transmitters are located on a 6 meter diameter circle centered on the receiver and they all continuously send packets with 7 and 31 bytes payloads. The packet error rate (PER) is computed at the receiver by dividing the number of packets considered correct by the PHY layer by the number of packets on which the receiver has synchronized.

On Figure 2.22a, the Ghassemzadeh NLOS channel attenuates the signal so strongly that even without interference (1 traffic source), almost all packets are lost. The Ghassemzadeh LOS channel leads to much better results, with a stark difference of an order of magnitude between the R-S and non R-S PER. With error correction, the PER remains below 5% even with four simultaneous traffic sources. The CM1 and CM5 channels degrade the PER stronger. This is due to their PDP that spread the interfering signals energies and increase the bit error probability. For a larger packet size (31 bytes), the PER degrades much faster as expected. The CM1 and CM5 channels produce a PER at about 50%, and the Ghassemzadeh LOS channel PER increases from 10% to 50%



(a) 7 bytes payload.



(b) 31 bytes payload.

Figure 2.22: Multiple access interference with all transmitters continuously sending, at 3 meters from the receiver (to isolate the effect of interference on data demodulation, synchronization always succeeds).




Figure 2.23: MAI with Poisson traffic for the Ghassemzadeh LOS and CM1 channel models, homogeneous interference.

when the number of traffic sources goes from 2 to 4. This type of continuous interference is not realistic, however. These results hold for the demodulation of data symbols in presence of interferers with the same power levels. In practice, nodes are not always active and collisions are more likely to effect only partially a packet reception: the end of a packet will be jammed by the beginning of another, for instance. In presence of interference, the synchronization process is more likely to fail.

Figure 2.23 shows the error rates for the same circle topology considered previously (radius of 3 meters), with 31 bytes data payload and for two Poisson packet generation rates (10 packets per second and 200 packets per second per node). These traffic rates were taken from [54]. In the high data rate case ($\lambda = 0.005$), the performance degrades quickly. For both the CM1 and Ghassemzadeh LOS channels, two traffic sources lead to a 10% PER, while with four traffic sources the PER reaches 50%. Even with nine traffic sources, about 20% of the transmitted packets can successfully be received. The low data rate cases show similar results, with lower PER since the reduced traffic intensity leads naturally to fewer collisions.

In all cases, the effect of the Reed-Solomon code is not noticeable. We deduce from this observation that packets arrive either without errors or with too many bit errors for the R-S algorithm to be effective (similarly to Figure 2.22b).

2.10.3 Simulation Performance

The channel model has a major effect on the simulation speed. We recommend to start working with the simpler Ghassemzadeh channel models for exploratory simulations, without power delay profile model, and to switch to the more complex IEEE 802.15.4A channel models to increase the accuracy of the results when required. On average, simulations using the Ghassemzadeh channel models were an order of magnitude faster to complete. The use of a Power Delay Profile significantly increases memory usage and processing time. The number of simultaneous transmissions also effects the processing time. The maximum interference distance can be configured in order to avoid unnecessary computations. As the low Pulse Repetition Frequency mode is faster as it uses only four pulses per burst instead of the 16 pulses of the 15.6 MHz PRF mode, it is the default mode of the simulator.

Fortunately, Omnet++ simulations can be run in parallel with zero overhead. Most of the time the user would like to study the variation of a metric as a function of a simulation parameter, for instance the variation of the packet error rate as a function of the distance between two nodes or as a function of the traffic intensity. Each considered value of this parameter requires an independent simulation run, which can be executed on any computer system. We adopted this approach to process simultaneously eight runs at a time on a eight-core computer.

Even when using the Ghassemzadeh models, the simulation speed is much slower than what can be obtained with a narrowband network simulation model. The simplest way to improve the performance is to make sure that the simulator is compiled in release mode, as this eliminates all assertions and activates several performance optimizations in Omnet++. A second step is to identify the model bottlenecks with a profiler, and then optimize these issues. For instance, we found that our simulations were slowed down by calls to `RSAMMapping::getValue()`, a base method of MiXiM. A third way to optimize the simulator speed is to improve the implementation of the IEEE PDP channel model: as this model generates a large quantity of points, some of those are very close to each other in simulated time. This leads to unnecessarily large data structures. Replacing such points with a single point is a process called binning. It was not implemented because it requires to be able to remove points from Mapping objects, a feature currently not available in MiXiM.

2.11 Observations

This chapter presented, to the author's knowledge, the first generic Ultra Wideband Impulse Radio physical layer model for network simulation, and the first network simulation model of the IEEE 802.15.4A UWB PHY layer. This model innovates by working at the symbol level instead of using complex and limiting mathematical approximations. Albeit slower than an analytical model, its advantages are numerous. It allows to study all types of receivers, *coherent* as well as *non-coherent*, as illustrated by the energy detector described in this work (analytical models focus on the correlation receiver). It enables to study the problem of synchronization and in particular the effect of interference on synchronization performance. It also possible to switch between channel models as needed, from simple models to evaluate large parameter spaces to highly

detailed models such as the IEEE 802.15.4A channel models when high accuracy is desirable. The effect of the Reed-Solomon error correction code can also be evaluated.

In this work, an energy-detection receiver was considered because of its low complexity, making it an ideal choice for sensor networks. The implementation of a correlation receiver would also be interesting for comparison purposes, especially one using the full error correction capabilities of IEEE 802.15.4A UWB PHY. Correctness, accuracy and validity were the priorities during the development of this simulator. Therefore, there are still several options available to optimize the simulator speed. The model itself can be optimized by implementing channel PDP binning, and the underlying MiXiM Mapping objects could be specialized to better support Maximum Pulse Amplitude Estimation. In this context, the use of a call profiler can provide valuable information.

The proposed modeling approach, named Maximum Pulse Amplitude Estimation, provides a signal representation close to the so-called Complex Baseband Equivalent Representation (CBER) [115, 123] which is used in signal processing. Indeed, the transmitted signal maximum amplitude representation is equivalent to the absolute value (or magnitude) of the CBER. MPAE differs from CBER mainly by making simplifying assumptions on the phase information. Adding the phase information to the signal representation by extending the TimeMapping class to store complex values instead of real numbers would allow the implementation of CBER in the network simulator, leading to greater accuracy, enabling the modeling of coherent receiver architectures, and extending the application potential of this tool to model any signal which can be represented using the complex baseband equivalent representation.

As the physical layer has an important effect on the performance of a communication system, it is hoped that this tool will allow more research in communication protocols for ultra wideband impulse radio wireless sensor networks, a field that has received little attention compared to narrowband wireless sensor networks.

WideMac: an Ultra Low Power Medium Access Control Protocol for UWB-IR

This chapter presents WideMac, an ultra low power MAC protocol based on asynchronous periodic beacon emissions and developed specifically for UWB-IR transceivers.

The chapter begins with an overview in section 3.1 of the problems faced by wireless MAC protocols and how they have been addressed by the research community. Section 3.2 follows, devoted to the specific challenges posed by ultra low power MAC protocols for wireless sensor networks. Various existing approaches to the problem are presented and their respective pros and cons are discussed. An ideal protocol is introduced for comparison purposes. Characteristics of IEEE 802.15.4A UWB-IR distinguishing it from narrowband radios are recalled from chapter 2. Section 3.3 gives an overview of UWB MAC protocols. WideMac is presented in detail in section 3.4. Analytical models are used in section 3.6 to compare it to the state of the art in terms of power consumption and latency. The chapter ends with observations in section 3.7.

3.1 Wireless MAC Protocols

While they are often evaluated with the same metrics of packet error rate (PER), throughput and latency, medium access control protocols operating over a wireless communication channel differ from their wired counterparts on many points. First, connectivity is essentially an unknown function of the environment, that fluctuates with time. Second, the link quality is also dynamic and the PER can be several orders of magnitude higher than the PER of wired communications. Third, the medium is intrinsically shared and open, thereby posing significant new security challenges that are not considered in this thesis. Fourth, the communication links can be *simplex* (in which case it is possible to communicate in only one direction between two radios), or *half duplex* (in which case it is possible to communicate in both directions but not at the same time), but never *full duplex* (in which case it is possible to communicate in both directions at the same time) which is typical for wired networks. Fifth, *listen-while-talk* capability is usually unavailable in radio devices. This useful

feature for the MAC layer is difficult to implement in a radio device due to the widely different power levels between a signal being transmitted and a possibly incoming signal, and the fact that part of the transmitted signal energy would leak into the reception circuit. Therefore, well-known techniques from wired MAC protocols that depend on this capability, such as Collision Detection, cannot be reused.

Table 3.1: A comparison of wired and wireless communication links from the MAC perspective.

	Wired links	Wireless links
Connectivity	User controlled	Unknown and dynamic function of the environment
Confidentiality	Good	Bad
Typical bit error rates	$< 10^{-6}$ (IEEE 802.3 objective ternary symbol error ratio is less than 10^{-8})	between 10^{-6} and 10^{-2}
Typical link type	half or full duplex	half duplex or simplex

The strong links between the wireless channel and the spatial environment influences not only the link types (simplex or half duplex) and their quality (packet error rate), but also the protocol operation. Indeed, the knowledge of the channel state at a network source node gives only limited information on the channel state at a network destination node. Furthermore, even this local knowledge is usually very limited: it often consists of a single real number varying over time called the *received signal strength indication* (RSSI), without giving any clue to what is the cause of this value. One way to use this information at the MAC layer is to compare it to some predefined threshold and to classify the communication channel as busy or idle if the value is respectively higher or lower. This channel state evaluation is sometimes called *Clear Channel Assessment* (CCA) in the literature. CCA support in the IEEE 802.15.4A standard is very limited: the impulsive nature of the transmissions make them very difficult to detect. Table 3.1 summarizes the key differences between wired and wireless links from the medium access control perspective while table 3.2 compares the narrowband and impulse radio ultra wide band communication channels.

This section begins by describing the objectives shared by most wireless MAC protocols. It continues by defining some well-known problems of this field, and ends with an overview of the existing solutions.

3.1.1 Objectives

The system throughput (the quantity of information that can be exchanged per unit of time) has received a lot of attention from the research community during the past few years, since the success of IEEE 802.11 (a,b,g) devices has put this limitation under the spotlight when compared to wired Ethernet connections. The necessity of sharing the spectrum with other systems, thereby

Table 3.2: A comparison of typical narrowband and UWB-IR links from the MAC perspective.

	narrowband [160]	UWB-IR [133, 99]
Propagation effects	Sensitive to fading	Robust to frequency selective fading
Spectrum availability	Low	High
Implementations	Mixed analog / digital	Possibly full digital
Theoretical ranging relative error (MSE)	27% [106, 63]	0.01% [59]
Transmitter power consumption (typ.)	40 mW	1 mW
Receiver power consumption (typ.)	40 mW	30 mW
bit rate (typ.)	250 kbps	0.85 Mbps
bandwidth (typ.)	1 MHz	500 MHz
center frequency (typ.)	868 MHz, 2.4 GHz ISM bands	4 GHz, 7 GHz

reducing the maximum possible throughput, has reinforced this need for higher bit rates at the physical layer.

But speed is not everything, and the communication latency, or the delay required for the transmission of an information, is another important metric to take into account. This is becoming more significant with the advent of Voice over IP (VoIP) communication systems which have tight constraints on the end-to-end application level latency (for instance, the acceptable one-way latency without echo cancellation technology for the telecommunication industry is 50 ms, and 100 ms with echo cancellation [82]).

Some other desirable properties are more difficult to quantify. The protocol *fairness*, or characterizing how well the protocol shares equally the system communication capacity among the competing devices, is difficult to evaluate since it depends highly on the particular configuration (application traffic load and number of devices) and on the resources to share (throughput, latency). For instance, a protocol that shares bandwidth equally well between all devices may do so by greatly increasing latency for some devices, at the benefit of some others. In that case, from the throughput point of view the protocol would be completely fair, while from the latency point of view it would not.

Another hardly quantifiable property, the simplicity of implementation, is often overlooked. In addition to reducing implementation efforts, this property also helps to understand, evaluate and fine-tune the protocol parameters as well as to predict its performance in various configurations.

Finally, a low power consumption is desirable as it allows greater mobility and improves usability. It is of course even more important when considering wireless sensor networks due to their particularly small form factor, low cost and requirements to operate several years unattended. This subject has been extensively covered in the literature and several representative protocols are

explained and compared in section 3.2.

3.1.2 Problems

Research in wireless MAC protocols has identified the following problems:

- Hidden Terminal (HT) [17], in which case the transmission attempts of two nodes cause collisions at the receiver, thereby preventing reception of both messages (see figure 3.1a). This phenomenon is also called Hidden Transmitter, or Hidden Node in the literature;
- Chained Hidden Terminal (CHT) [108], a more complex case of the HT that happens when a specific resource reservation technique (RTS/CTS) is applied to address the HT problem;
- Exposed Terminal (ET), in which case a node can never access the channel because it is always kept busy by another node, even though the transmissions could be safely operated concurrently as they address two different destination nodes (see figure 3.1b);
- Capture, in which case a destination node always synchronizes on one of two sender nodes because its received signal strength is much higher than the other (typically on a narrowband spread spectrum FM system);
- Collisions, which is a generalized case of hidden terminal and concerns all reception attempts that ultimately fail because of another concurrent transmission reaching the destination node;
- Interference, which is a generalized case of collisions and exposed terminal, and considers all packet transmissions in the vicinity of a node that prevent either transmission or reception.

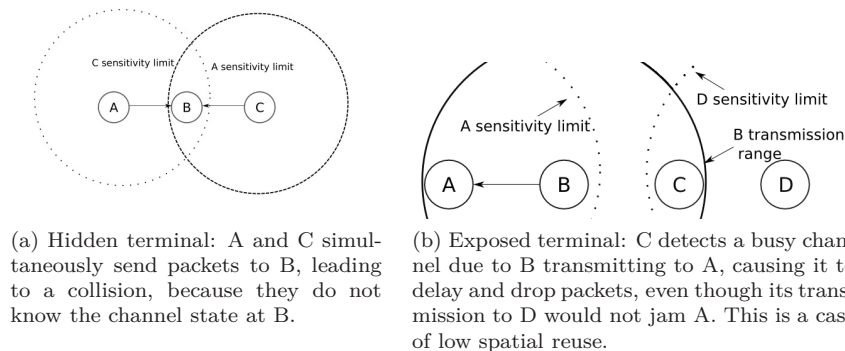


Figure 3.1: The hidden and exposed terminal problems are both caused by erroneous assumptions on the channel state at the destination node.

Figure 3.1 illustrates the hidden and exposed terminal cases and figure 3.2 illustrates the relationships between all of these problems. The influence of the underlying assumptions is important: while the HT, ET and CE scenarios can




Figure 3.2: Problems identified in the Wireless MAC literature.

exist in a propagation disk model (in which transmissions always succeed if the distance between two nodes is lower than a radius R , and fail otherwise), to correctly account for interference one must consider the probabilistic nature of reception (with a Signal to Noise Ratio to Bit Error Rate model) and the channel attenuation (with either a constant or stochastic pathloss exponent) in the reasoning.

When a protocol prevents transmissions that would not negatively effect any ongoing or future transmission, because e.g. of a hidden terminal problem or scheduling errors, it is said that the protocol has *low spatial reuse*. [8] studies extensively the spatial reuse maximization problem in multihop wireless networks.

3.1.3 Existing Approaches

There are two broad categories of protocols: scheduled protocols and random access protocols. Scheduled protocols attempt to organize the traffic so as to avoid collisions, hidden terminal, exposed terminal and capture effect. To do so, they need to exchange control information to perform neighbor discovery and traffic requirement estimates. While this approach can work in some cases, it has several shortcomings: inability to scale with network size, difficulty to operate in a completely distributed mode, and lack of robustness to interference from other systems. Indeed, since these protocols assume that by exchanging control information they can completely eliminate collisions, they often lack to offer acknowledgments and retransmission mechanisms if a packet were not to be received. Such an approach can work for wired networks, but the wireless medium is by nature shared and thus more sensitive to interference from other systems. Other networks operating with the same technology could also be nearby, leading to collisions and packet losses. Thus, this thesis focuses on random access protocols since they are more robust to packet losses.

ALOHA [5, 124, 122] was one of the first deployed wireless computer networks, between Hawaii islands. With ALOHA, a node having a packet to send waits for a random time interval (called a *backoff* time), sends its packet and listens for an acknowledgment. If no acknowledgment arrives after a certain time (this event is sometimes called a NACK), it is assumed that the packet was not correctly received by the other side. In that case, the node waits for a new random backoff time and retries the transmission. ALOHA's main drawback is that the probability of collisions quickly increases with the number of users attempting transmission. This is due to the long *vulnerability window*, or the

time interval during which another transmission can collide with an ongoing transmission. If all packets are of equal duration t , the ALOHA vulnerability window is equal to $2t$.

The introduction of Slotted ALOHA [122] doubled the theoretical maximum attainable throughput by reducing this vulnerability window by a factor two (assuming constant duration packets). This is achieved by requiring network-wide time synchronization, dividing the time in slots of equal duration (equal to the sum of the maximum packet size and of the maximum time to send an acknowledgment) and allowing nodes to compete for channel access only at the beginning of each slot. This reduces the vulnerability window to t , the maximum duration of a packet. However, in case of low traffic the latency is increased because nodes must wait for the next slot to transmit. Slotted ALOHA, indeed, trades latency for throughput. The vulnerability windows of ALOHA and Slotted ALOHA are depicted respectively on figures 3.5a and 3.5b.

CSMA (Carrier Sensing Multiple Access) [122] achieves to further reduce the vulnerability window through another approach. Instead of assuming network wide time synchronization, CSMA requires all nodes to be able to evaluate whether the channel is available (Clear Channel Assessment capability). A candidate sender evaluates the channel state. If it is idle, the node can transmit a packet. If it is busy, the node will retry later. Several variations exist for the channel access algorithm:

- 1-persistent CSMA, which requires nodes to listen constantly to the channel and let them transmit with probability 1 as soon as it is found idle, as illustrated on figure 3.3a;
- non-persistent CSMA, which waits for some random time between two clear channel assessments to reduce collisions, increasing the minimum latency to increase the performance in case of multiple access, as illustrated on 3.3b;
- p-persistent CSMA, which repeatedly perform carrier sensing to detect the channel availability as quickly as 1-persistent CSMA but avoids systematic collisions by using a backoff timer with a probability p , which must be tuned to the expected number of simultaneous senders for best performance (see figure 3.3c);
- CSMA with Collision Avoidance (CSMA/CA), which defers the next carrier sensing by a random time interval if the channel is found busy, in order to reduce the probability of collisions.

In addition, there is another version of CSMA named CSMA/CD, which uses Collision Detection to abort the transmission as soon as a collision occurs (instead of continuing the transmission and waiting for an acknowledgment), but is rarely available for wireless communications because it requires a second radio for *listen-while-talk* / full duplex operation. Some have considered the use of a second radio operating on a dedicated signaling channel to reproduce the behavior of CSMA/CD networks [108].

Like ALOHA, all these variants of CSMA can be *slotted* or *unslotted*, depending on the possibility to synchronize all nodes.

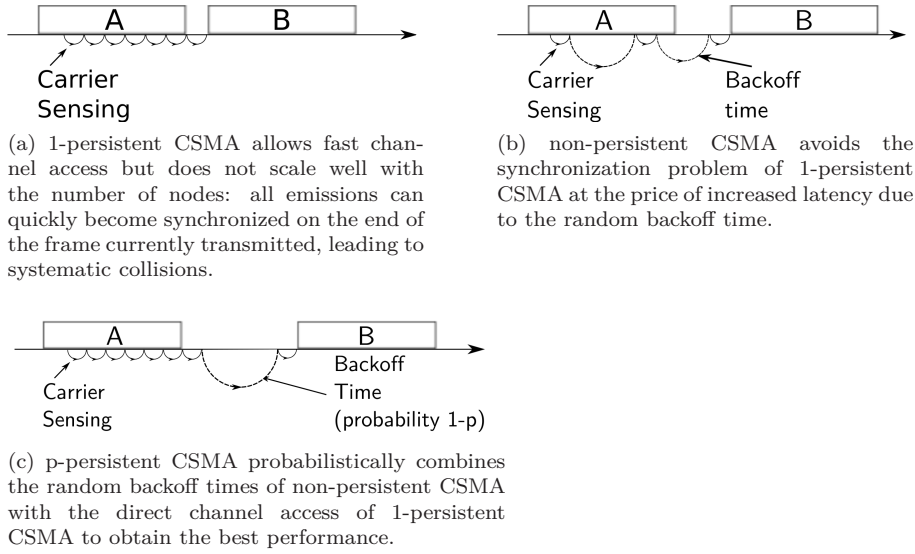


Figure 3.3: The CSMA channel access algorithm has a significant effect on collisions, latency and throughput. Three strategies are illustrated here, with a node A transmitting and a node B applying the algorithm to send another packet.




Figure 3.4: CSMA with resource reservation: CSMA/RTS/CTS/Data/ACK as defined in MACAW and IEEE 802.11 virtual carrier sensing (same as MACA but with acknowledgments).

With CSMA, the vulnerability window t_{vuln}^{CSMA} is reduced to the time required for the CCA t_{CCA} summed to the time interval between the beginning of a frame transmission and the time at which this transmission is detectable by the other nodes (composed of the propagation time t_{prop} and a hardware minimum signal detection time $t_{SigDetect}$):

$$t_{vuln}^{CSMA} = t_{CCA} + T_{prop} + T_{SigDetect}.$$

The signal detection time $T_{SigDetect}$ will usually depend on the received signal strength, as a stronger signal can be detected faster, and on the bit rate, as a higher bit rate generally leads to more complex signals. [181] describes the most common techniques for detecting channel sensing. Note that the CCA time t_{CCA} includes here the received signal strength integration time as well as the time to switch the radio from reception to transmission, which usually lasts 100 μ s on a modern radio. Figure 3.5c illustrates the vulnerability window of

CSMA.

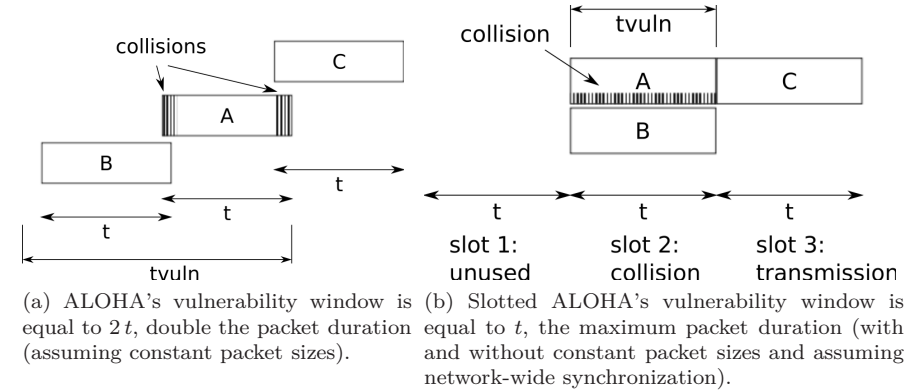


Figure 3.5: Vulnerability windows of three well known wireless MAC protocols, with three nodes A, B and C each transmitting a message (see [122]).

While CSMA leads to significant performance improvements compared to ALOHA and Slotted ALOHA, the CCA operation can mislead the candidate sender node since it evaluates the channel state at the sender node while what matters is the channel state at the destination node [77]:

- By assuming that the channel is idle at the receiver while it is actually only idle at the sender, CSMA is vulnerable to the hidden terminal problem (see Fig. 3.1a);
- By assuming that the channel is busy at the receiver while it is actually only busy at the sender, CSMA is vulnerable to the exposed terminal problem and suffers from low spatial reuse (see Fig. 3.1b).

To enable more informed decisions, MACA [77] replaces carrier sensing by a three way handshake: the candidate sender node firsts sends a Ready To Send (RTS) message to the destination node. If it can receive this message and if the channel is free at its location, the destination node authorizes the transmission by sending a Clear To Send (CTS) message to the candidate sender, which replies with the data packet. This scheme is also called RTS-CTS-Data or

CSMA with RTS-CTS, and is part of the reservation protocols family that attempt to reserve resources before using them. This way, the sender is informed of the channel condition at the receiver before transmitting, and collisions should only happen between short RTS and CTS control packets (this protocol assumes long data packets, which does not hold in wireless sensor networks). The idea behind MACA is that the destination node will inform all possible interferers of the incoming transmission by sending the CTS message, and thus prevent the Hidden Terminal problem. However, this reasoning assumes that:

- the interference range is equal to the transmission range, which means that if two nodes are close enough to jam each other, they can also exchange packets;
- bidirectional communication links, and thus that any node that can send a packet to the destination node can also receive packets from this destination;

Unfortunately, these two assumptions are false. First, the interference and transmission range depend on the deployment environment and on the radio technology in use, and for most if not all technologies, the interference range is much larger than the reception range. Second, radio links can often be asymmetric due to propagation effects. Therefore, while the RTS-CTS exchange introduced in MACA reduces occurrences of the hidden terminal problem, it does not eliminate it since a node that cannot receive a CTS message can still interfere with the reception process. A transmission using this mechanism is illustrated on figure 3.4. The RTS-CTS exchange can also be used to offer a more general *Virtual Carrier Sensing* abstraction (as in IEEE 802.11 [152]), that tells the node to defer communications for some duration without performing an actual CS operation: these control messages let all nodes know that the channel will be used for some time.

MACAW [17] aims at improving MACA's performance in high throughput and improving its fairness (ET). It does so by modifying the backoff algorithm so that congestion at one point is spread in the network through exchanges of the nodes' backoff exponent values, and it smooths the fluctuations of the backoff exponent (BE).

MACA uses a Binary Exponential Backoff (BEB) algorithm: the backoff time interval is obtained by multiplying a unit backoff duration $T_{Unit}^{Backoff}$ with an integer chosen uniformly between 0 and $2^{BE} - 1$, where BE is the backoff exponent. The backoff exponent fluctuates between an initial minimum value BE_{min} and a maximum value BE_{max} . After each unsuccessful transmission, BE is increased by 1, and each new packet begins with a backoff exponent set to BE_{min} . In a saturated network, with all nodes continuously having packets to transmit, a node that succeeded in transmitting one packet is advantaged because (1) it is not in backoff state contrarily to most of its neighbors and (2) it will reset its backoff exponent to the minimum value, and spend little time in backoff state. This is why in such situations, one or more nodes will be in backoff state most of the time and drop most of their packets while some other nodes will be able to transmit all of their packets. MACAW fixes this by adding in each packet sent, a header field that announces the current value of the backoff exponent at the sender. A node receiving this packet copies the received backoff exponent value and uses it as its own new BE value (as

a consequence, in a small MACAW network in which all nodes can receive all packets, there is a network-wide unique BE value). This improves the fairness of MACA.

Because it always uses the minimum BE for the first transmission attempt of a packet, and grows the backoff window exponentially, MACA's strategy leads to wide fluctuations of the backoff time. MACAW introduces a Multiplicative Increase and Linear Decrease (MILD) backoff algorithm. Instead of always using the minimum backoff exponent for new packet transmission attempts, the backoff exponent is decreased by a constant value for each successful transmission, until it reaches its minimum value. And instead of adopting an exponential growth, the backoff window is increased by a multiplicative factor each time that the channel is found busy. MACAW also reintroduces acknowledgments, which were erroneously assumed superfluous in MACA, to avoid the overhead of transport-level retransmissions.

The reservation mechanism introduced in MACA was later adopted in IEEE 802.11, which is based on CSMA. Before transmitting, a node using IEEE 802.11 must first wait for the channel to be free during a Distributed Inter Frame Space (DIFS). When this condition is fulfilled, the node sends its packet and waits for an acknowledgment. The receiver node sends it after a Single Interframe Space (SIFS) time, long enough for the sender's radio to switch between transmission and reception. A binary exponential backoff algorithm is used if no acknowledgment is received. The RTS/CTS mechanism is optional and is used only for frames larger than a user-defined threshold [175]. This enables the nodes to communicate quickly small packets and to protect long packets from the hidden terminal problem. In addition, a cooperative mechanism is implemented: RTS/CTS packets include a time field that informs overhearing nodes of the incoming exchange duration. This information allows these nodes to defer transmission attempts after the end of the current exchange. The timer used for this mechanism is called the Network Allocation Vector.

While this section is not exhaustive, it presented some largely used techniques. For instance, the IEEE standard for narrowband wireless sensor networks, IEEE 802.15.4 [151], defines a CSMA/CA mode with binary exponential backoff for full mesh networking operation.

3.2 Ultra Low Power MAC Protocols

Reducing the power consumption of wireless communication systems is mainly achieved by duty-cycling the radio and keeping it in sleep mode as much as possible. The various types of energy waste at the MAC layer have been clearly identified in the literature and are described in subsection 3.2.1. The existing protocols can be divided in two categories: random access and scheduled access [83]. Figure 3.6 represents several well-known ultra low power wireless MAC protocols and separates them depending on their channel access type (scheduled or random) and the network structure they require (centralized or distributed). The IEEE protocols considered here come in two versions, a centralized one and a distributed one. Although the distributed (or Ad Hoc, or non beacon-enabled) mode of these protocols does not allow low power operation, they are included on the figure for completeness and regrouped in a dashed box. As

IEEE 802.15.4 in its beacon-enabled mode allows both scheduled and random access, it overlaps the two categories on the figure.

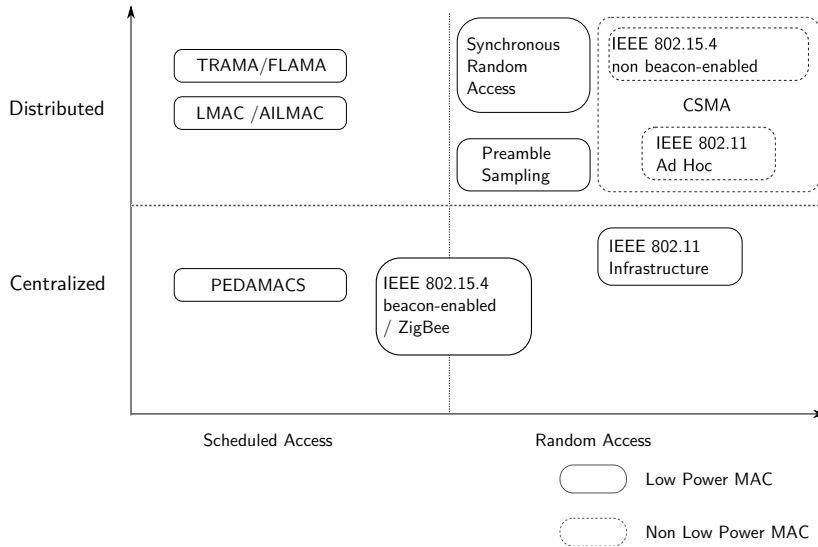


Figure 3.6: Low Power wireless MAC protocols.

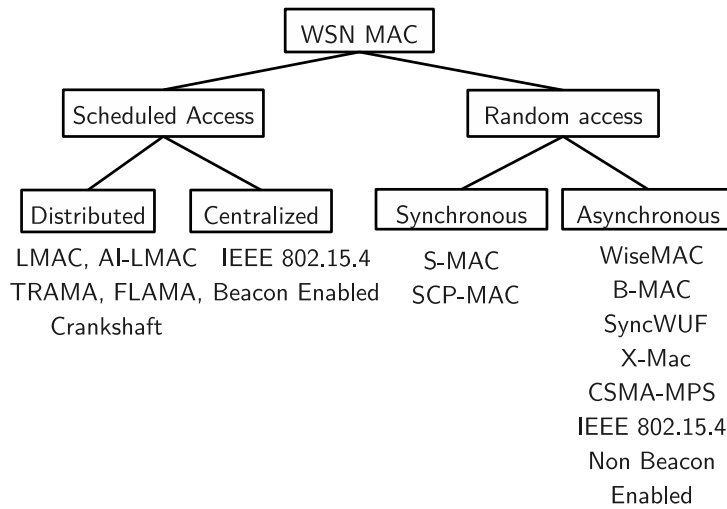


Figure 3.7: Classification of low power MAC protocols.

Scheduled access protocols organize communications through a Time Division Multiple Access (TDMA) approach. These protocols aim to schedule communications in a manner that prevents collisions, overhearing and idle listening. They can be either distributed, or centralized. Centralized scheduled access protocols such as the IEEE 802.15.4 beacon-enabled mode delegate the coordination task to a central authority.

Random access protocols grant channel access through a contention resolution mechanism. Nodes wake up periodically to briefly enter reception mode.

If the channel is idle, the radio enters low power mode again. If the channel is busy, the nodes try to receive a message.

Several well-known ULP MAC protocols are classified on figure 3.7. Subsection 3.2.1 describes the major causes of energy waste at the MAC layer and introduces an ideal protocol. Scheduled access protocols and random access protocols are presented respectively in subsections 3.2.2 and 3.2.3.

3.2.1 Types of Energy Waste and the Ideal Protocol

All protocols must deal with five types of energy waste [176, 46]:

- collisions,
- overhearing,
- idle listening,
- signaling overhead,
- over-emitting.

Collisions happen when transmissions occur simultaneously in such a manner that message reception fails for at least one of the intended recipients. *Overhearing* occurs when a transceiver uselessly listens to a message. *Idle listening* happens when the radio is in reception mode while no transmission takes place. *Signaling overhead* is the energy cost incurred by signaling data such as acknowledgments and synchronization packets. Lastly, *over-emitting* is the energy cost associated to transmitting packets when the destination node is not ready to receive them.




Figure 3.8: Types of energy waste in ULP MAC protocols and associated radio states.

Figure 3.8 represents these five cases and the associated radio states.

To compare protocols to an objective reference, it is useful to introduce an ideal protocol. The ideal protocol does not suffer from any type of energy waste: there are no collisions, nodes never listen to the channel uselessly, transmitted packets are always received correctly, and there is zero signaling overhead. Thus, nodes using the ideal protocol always stay in sleep mode, except when they have a message to send, which they can do immediately, or when they must receive a message, in which case they enter reception mode just before the message arrival. Packets transmissions are always successful and therefore never acknowledged. This protocol is highly unrealistic for three reasons:

1. all nodes know all traffic generation events (so that they can enter reception at the correct time);
2. the nodes can simultaneously transmit without ever suffering from collisions;
3. the propagation channel does not impair reception.

Despite this, the ideal protocol is useful in the sense that it provides intuitive lower bounds on a communication protocol's performance.

3.2.2 ULP Scheduled Protocols

ULP scheduled protocols all attempt to schedule communication between nodes so as to avoid collisions. PEDAMACS [47], LMAC [164], AI-LMAC [23], TRAMA [120], FLAMA [119], ISOMAC [178] and Crankshaft [66] all belong to this category. In general, the required signaling traffic to achieve the goals of distributed scheduled access protocols generates some overhead. In addition, time synchronization is often assumed in the evaluation and its difficulty underestimated, despite the fact that its energy and bandwidth costs can be significant and that scalability can be an issue. This subsection describes some of the most famous protocols, and highlights the design challenges of the scheduled strategy.

PEDAMACS [47], for Power Efficient and Delay Aware Medium ACcesS protocol, is a scheduled MAC protocol that centralizes the scheduling task at the sink, and assumes that the sink is equipped with a more powerful radio that allows it to directly reach all nodes in the network. The sink first initiates a spanning tree and collects information from all nodes about their neighbors. Based on this topology information, it computes a collision-free global schedule. This schedule is delivered to each node, and data collection begins. PEDAMACS suffers from several problems. First, the assumption that the sink can reach all nodes is difficult to guarantee in practice, due to the dynamics of radio propagation and regulatory constraints. Second, the validity of the topology information on which the global schedule is built is not guaranteed to be complete, and can vary over time. Third, the protocol is relatively complex, which makes it difficult to implement and operate over resource constrained sensor network devices.

TRAMA [120], for TRaffic Adaptive Medium Access protocol, is a fully distributed scheduled protocol. Each node periodically announces its presence, its traffic generation pattern and the list of its neighbors. After some time, this allows all nodes to build a list of their two-hop neighbors. A collision-free schedule is obtained locally at each node by using a hash function to attribute a unique sender to each slot. Sender nodes can generate a traffic indication map, that associates the destination nodes with their next transmission slot. This information allows other nodes to enter sleep mode and save energy. FLAMA [119] reduces the signaling overhead of TRAMA. Instead of depending on periodic broadcasts from all nodes, it uses a polling-based mechanism to retrieve the necessary information only when needed. It introduces random access slots to allow new nodes to join the network. Despite these improvements, its energy consumption is barely better than S-MAC [119].

LMAC [164], for Lightweight MAC, also distributes slots based on two-hop neighborhood information. It attempts to minimize transitions between reception and transmission modes, acknowledging the cost of the radio switching times. All slots are of fixed duration. A node can only transmit during the slot that it owns. Each node begins transmission during its slot with a header (followed or not by a data payload). This header includes information on the next slots taken by this node and its one-hop neighbors. New nodes joining the network listen to the transmitted headers to build a list of still available slots, and simply choose randomly one of the available ones. Messages are not acknowledged, which can be problematic in wireless networks. The cost of periodically sending the header information penalizes LMAC. Further, the number of slots must be larger than the number of nodes in any two-hop neighborhood, leading easily to over-provisioning and to capacity constraints. Adaptive Information-Centric LMAC (AI-LMAC) [23] partly addresses this problem by allowing nodes to claim more than one slot per frame. The use of periodic announcements of the already used slots to avoid collisions between two hop neighbors in a distributed scheduled protocol was later reused for ISOMAC [178], which claims in addition to operate without network-wide time synchronization.

Finally, Crankshaft [66] implements a hybrid approach. It operates as a scheduled protocol because it divides time into frames of duration T_{Frame}^{CS} , which are further divided into slots of duration T_{Slot}^{CS} . The first slots are unicast slots and each node is associated to one slot during which it must listen. At the beginning of a unicast slot, nodes wanting to address the slot owner contend for channel access during a slotted contention window of duration T_{Cw}^{CS} . Slots are chosen according to the Sift distribution [74] which minimizes the collision probability. Candidate transmitters send a preamble until the slot owner polls the channel, at which point they transfer the data and normally receive an acknowledgment message. The last slots are broadcast slots and all nodes must listen to all of these. Access to broadcast slots is contention based.

3.2.3 ULP Random Access Protocols

Some protocols such as S-MAC [176], T-MAC [163] and SCP-MAC [177] synchronize the time at which all nodes perform this carrier sensing. They are called here *synchronous*. Others let nodes perform carrier sensing independently, and they are called *asynchronous*.

S-MAC introduced the concept of synchronized sleep. All nodes periodically wake up and exchange RTS/CTS signaling messages. Afterward, inactive nodes enter sleep mode again while communicating nodes exchange data messages. The nodes stay synchronized by listening for a periodic beacon emission from one of the nodes, preceding the RTS/CTS exchange. Since all nodes wake up at the same time, this scheme creates a lot of signaling collisions. In addition, the throughput is limited to one packet per wake up time interval for all nodes. The third problem is related to the synchronization method. If a node does not receive a synchronization beacon, it will start sending its own after some time. Thus, the network becomes partitioned in multihop configurations and some nodes must act as relays and listen to two beacons per period. When considering the clock drift due to the imprecision of the quartz crystal used as time reference, inevitably these beacons overlap regularly and the relay node /

nodes cannot communicate (in addition to their increased power consumption). Throughput was later improved both in T-MAC and in a second version of S-MAC [48] by adding an RTS/CTS period after each data exchange, enabling more than one data transmission in the network per S-MAC wake-up cycle. This later version is referred to as S-MAC/T-MAC in this work. SCP-MAC, by the authors of S-MAC, is one of the most recent synchronous random access protocols. All nodes periodically and simultaneously wake up and poll the medium. Candidate transmitters contend for medium access just before this periodic polling. Nodes switch between two fixed polling rates depending on current traffic load (this mechanism is called *adaptive channel polling*). The value of the SCP-MAC polling rates and of the S-MAC/T-MAC polling rate have an important effect on the system performance. High wake-up rates lead to large power consumption, and low rates limit the throughput. SCP-MAC attempts to reduce S-MAC's signaling overhead and throughput limitations, but does not really address the multihop issues and requires design-time precise parameters tuning for the expected network traffic.




Figure 3.9: Random access ULP MAC protocols: no traffic.

Asynchronous random access protocols do not synchronize nodes on a common schedule. Instead, all nodes periodically wake up (every period T_W) to perform carrier sensing. In Low Power Listening [70], B-MAC [112] and NP-CSMA-PS [43], when a node has a message to transmit, it sends a long wake-up preamble (of duration T_W) so that all nodes in the vicinity are aware of the transmission. To avoid the Hidden Terminal problem, without introducing a reservation mechanism such as RTS/CTS, these protocols can use two different received signal strength thresholds to detect (1) an incoming transmission and (2) an ongoing transmission to another node. The value used to detect an incoming transmission determines the *reception range* while the value used to detect an ongoing transmission determines the *interference range*. To reduce

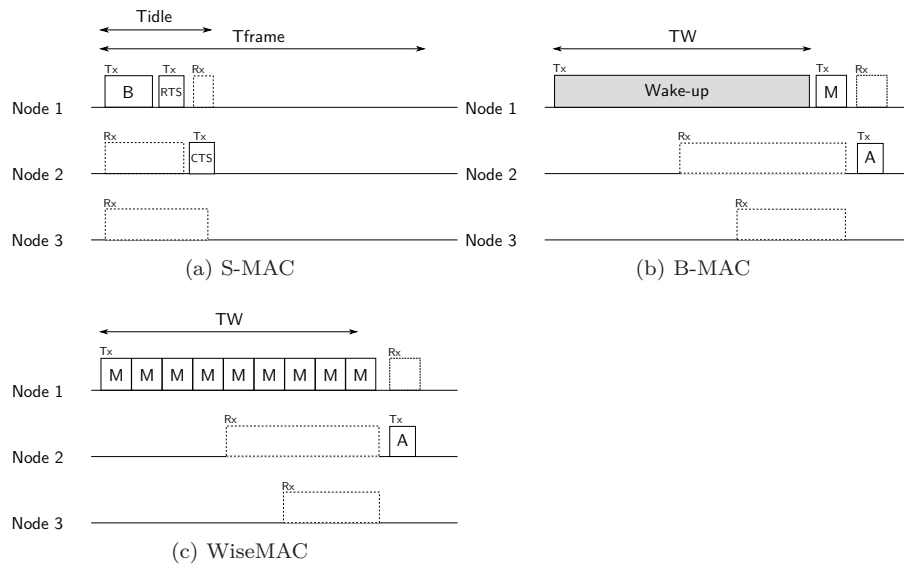


Figure 3.10: Random access ULP MAC protocols: first packet exchange from node 1 to node 2.

collisions, the considered interference range must be larger than the considered reception range, and thus the threshold must be lower. While this scheme, sometimes called more generally *preamble sampling*, avoids depending on a synchronization mechanism, it causes overhearing and signaling overhead.

WiseMAC [44, 45] addresses these problems by introducing adaptive length wake-up preambles: after a data packet transmission, the recipient sends back an acknowledgment which provides timing information on its next channel sampling. The next time a transmission occurs on the same link, the wake-up preamble can be shortened because the sender node can predict more or less accurately the time at which the receiver will wake up again. The more frequent the transmissions on a link, the better the prediction and the higher the reduction in wake up preamble length, and thus in overhearing and signaling overhead. WiseMAC further optimizes the scheme by repeating copies of the message in the wake-up preamble. This allows overhearing nodes to quickly determine that the transmission is not addressed to them and go back to sleep. This method works best for packet sizes relatively small compared to the wake-up preamble (which is typical of wireless sensor networks): if the packet is long, the overhearing node must listen for a longer time.

SyncWUF [144], CSMA-MPS [94] and X-MAC [20] all propose various optimizations on top of WiseMAC to reduce the time spent transmitting and listening to wake-up preambles. By repeating the message destination address in the wake-up preamble when the message itself is too long to be repeated, SyncWUF reduces the overhearing caused by WiseMAC long preambles. CSMA-MPS and X-MAC both propose to replace the wake-up preamble by alternating between the transmission of short wake-up packets and channel polling at the message source. As soon as the destination receives one of the small wake-up preamble packets, it sends a ready-to-receive packet and the message can be transmit-

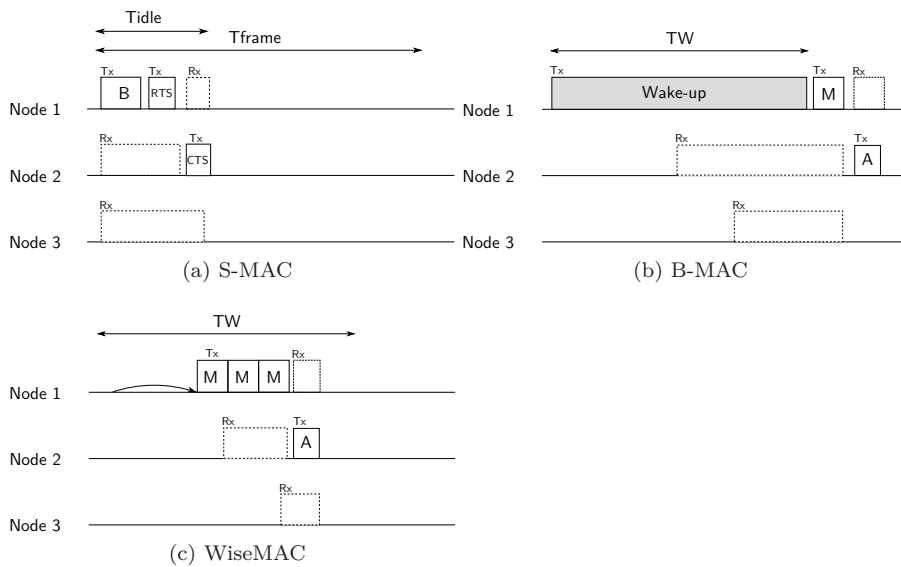


Figure 3.11: Random access ULP MAC protocols: second packet exchange from node 1 to node 2.

ted immediately afterward. This scheme, similar to an asynchronous RTS/CTS reservation scheme, minimizes overhearing at the receiver side since the preamble length is reduced. Its influence on idle listening is more discussable as it increases the cost of the periodic channel polling: since the destination node can perform the channel polling during one of the idle slots reserved for CTS messages, the channel polling must be longer than this period. And the objective of reducing the wake up preamble further than WiseMAC can only be worth in low traffic since otherwise the wake up preambles are automatically reduced, in which case the cost of the periodic channel polling already dominates WiseMAC's power consumption. Detailed analytical power consumption models were developed to compare all these protocols. This work is presented in section 3.5.

3.3 UWB-IR MAC Protocols

The introduction of an ultra wideband impulse radio PHY layer involves several significant changes compared to narrowband radios:

- no Clear Channel Assessment capability;
- better robustness against interference thanks to the large bandwidth;
- higher bit rates (large bandwidth);
- ranging capabilities;
- modification of the relative power consumption levels: the power consumption in synchronization mode is significantly higher than in reception mode, which is itself much higher than in transmission mode.

The last point is a direct consequence of the relative complexity of these operations. In a narrowband radio, the power consumption in synchronization mode is equal to the reception power consumption, and of the same order of magnitude as the transmission power consumption. The first point, no CCA capability, renders most wireless MAC protocols unsuitable for UWB-IR as is because most depend at least partially on this feature. The higher robustness against interference and the higher bit rates could allow new, previously impractical designs to reach high performance levels. The new relative power consumption levels have the potential to modify the relative performance levels of ULP MAC schemes, and also offer novel possibilities.

Numerous publications are available on UWB MAC protocols [65, 140, 87, 37, 185, 33, 32, 27, 86, 103, 38, 143, 142, 36]. While most of them consider UWB-IR, we could not find any MAC designed specifically for the IEEE 802.15.4A PHY layer (excepted an adaptation [36] of an already presented UWB MAC to the standard). This section regroups existing UWB MAC protocols in two groups: adaptive communication parameters and pulse sensing. Protocols in the first group all attempt to minimize collision and interference by modifying the parameters of the PHY layer: time hopping sequences, CDMA, bit rate and power level. They are presented in subsection 3.3.1. Protocols in the second group all assume the availability of a pulse sensing device that allows detection of an ongoing transmission. They are presented in subsection 3.3.2. Subsection 3.3.3 describes the MAC features included in the IEEE 802.15.4A standard, and subsection 3.3.4 closes this section with observations on existing UWB MAC protocols.

3.3.1 Adaptive Communication Parameters

The WHYLESS project developed a Time Hopping based distributed MAC protocol [33]. It allows fully distributed operation, and adapts the transmission rate and emission power to the channel conditions. It uses pseudo random Time Hopping Sequences (THS), and pulse position modulation to encode the information. The THS define quasi orthogonal channels. A common control channel is used to setup communications, dedicated control channels are used for the exchange of control information related to ongoing data exchanges, and dedicated traffic channels are used for the transmission of data. Figure 3.12 illustrates this time hopping mechanism and the fine time resolution that it requires among nodes to avoid pulse collisions.

Chu [27] also proposes to adapt the transmission rate and emission power, but adopts RTS-CTS exchanges on the control channel and requires all nodes to see each other. All nodes periodically announce themselves on the control channel with HELLO messages, that include the node ID, its path gain, and its interference margins. This information is used during the link establishment procedure (RTS-CTS exchange, called Link Request packets) to allocate adequate communication resources (power level, bit rate) for the new link, using heuristic methods. It distinguishes two classes of traffic to prioritize more important packets.

U-MAC [76], for Ultra Wideband MAC, makes all nodes periodically broadcast HELLO packets that are used by neighbor nodes to evaluate the link quality. As in [27], all nodes must see each other, there are two classes of traffic and the physical layer is UWB-IR with Time Hopping Pulse Position Modulation

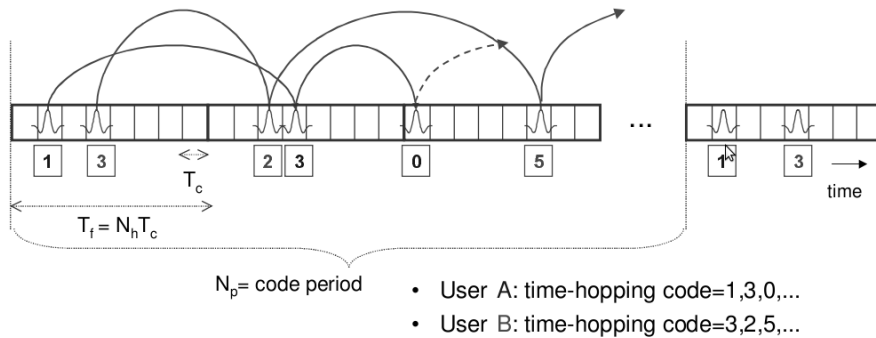


Figure 3.12: The Time-Hopping based distributed MAC protocol developed in WHYLESS (source: [33]) illustrates how time hopping can prevent pulse collisions, when nodes are synchronous at the pulse level and when multipath propagation effects are neglected.

(TH-PPM). The frequency of the HELLO messages depends on the node stability, a concept depending on node mobility and node communication activity. Listening nodes deduce the interference level from the reception of the HELLO packets. All the collected information is then used to adapt the transmission rate and power levels. The periodic listening that it requires leads to significant energy consumption.

DCC-MAC [86] allows multiple hop operation and is fully distributed. It also adapts the transmission rate to the link quality but always operates at full transmission power. It relies on time hopping to establish quasi orthogonal communication channels. A specific Time Hopping Sequence (THS) is associated to each network node, and there is also a broadcast THS. A node wanting to send a packet to another must use the THS of the destination. This allows robust concurrent transmissions, with minimal interference. The way these THS are distributed in the network is not specified.

In (UWB)² [37, 36], an approach opposite to DCC-MAC is adopted: instead of THS associated to destinations, the THS are associated to transmitters. A common signaling channel is used before the data exchange so that the destination node knows the THS sequence of the source and can receive the packet. The signaling channel is also used to allow new nodes to join the network.

Self-Balanced Receiver-Oriented MAC (SEBROMA) [10] also uses receiver-associated THS. Each node periodically broadcasts a ready to communicate message that includes its THS code. After sending this message, each node will listen for some duration, during which candidate senders can compete for channel access using an RTS-CTS exchange.

SDD [143, 142] makes the assumption that nodes can permanently monitor the channel at low power for control packets sent using a common THS. RTS-CTS exchanges are used to initiate communications that use transmitter-associated THS. The THS are uniquely associated to each transmitter, and can be derived from the MAC address as proposed in [57]. This scheme considers only single-hop networks.

Cuomo [32] models the MAC objectives as a joint power and rate assignment optimization problem, and proposes distributed algorithms to reach sub-optimal solutions.

CC-CDMA [185] targets WPAN and uses CDMA techniques to create orthogonal UWB communication channels. It requires a network coordinator. CDMA techniques are complex and lead to power-consuming transceivers.

3.3.2 Pulse Sensing and TDMA

In [11, 12], a pulse sensor device is proposed to overcome the lack of carrier. This should enable fast and efficient clear channel assessment, and allow the use of CSMA techniques. The feasibility and the performance characteristics of such a device in terms of power consumption, speed and accuracy are not clear, and thus this was not considered in this thesis. In the UWEN project, the nanomac [103] protocol was proposed. It is a p-non-persistent CSMA/CA-like protocol. The authors adopt the Energy Sense Multiple Access (ESMA) name since the carrier sensing is replaced by energy detection. It uses RTS-CTS exchanges and adopts the Virtual Carrier Sensing mechanism of IEEE 802.11 to further reduce collisions.

UCAN [87] is a MAC protocol designed for wireless personal area networks. It operates over an UWB-IR link, requires a device acting as network coordinator and all nodes must see each other. The network coordinator allows network-wide time synchronization and attributes time slots to each node to avoid collisions.

3.3.3 IEEE 802.15.4A MAC Features

The IEEE 802.15.4 protocol defines two modes of operation: beacon enabled and non beacon enabled. In the first mode, a coordinator, called the PicoNet Coordinator (PNC) sends periodic beacons. Each beacon is followed by a so-called Contention Access Period (CAP), during which all nodes can compete freely for channel access using a CSMA algorithm, and by a Collision Free Period (CFP), during which nodes communicate during time slots exclusively allocated by the PNC. In the non beacon enabled mode, all nodes use a CSMA protocol to communicate.

The introduction of an UWB-IR PHY layer in the IEEE 802.15.4A standard made this protocol unable to operate, since it relies on CCA (in both of its modes). Therefore, adaptations were defined in the standard. In particular, the CSMA mode is replaced by an ALOHA mode that does not rely on CCA for correct operation.

3.3.4 Remarks

While the idea of modifying the communication parameters, such as power level, transmission rate and time hopping sequences, to minimize interference is attractive and was extensively studied, it has a few drawbacks:

1. the strict regulations on radiated power levels mean that the UWB link budget is very low. Reducing the radiated power further will likely lead to high packet losses. As such, we favor the idea proposed in [86] to always transmit at the maximum power.

2. Varying time hopping sequences increases the receiver's complexity, and thereby its power consumption. For sensor networks, it may not be the best option. This thesis focuses instead on the IEEE 802.15.4A mandatory mode in the hope that this hardware will become available in the near future. This may not be the case for more sophisticated PHY layers, using for instance variable THS.
3. Variable communication parameters increase protocol complexity, and therefore the risk of incorrect implementation. There needs to be at least one set of base values so that all nodes can reach each other. In effect, this approach leads to multi-channel MAC protocols, even if the signal's center frequency remains the same. This is more or less clearly stated in the considered publications. While multi-channel MAC protocols have several advantages, notably concerning robustness to interferers, they require a careful and detailed study. This chapter focuses on how to get the best possible performance on a single communication channel.

Pulse sensing would be a welcome feature on UWB transceivers. However, there is a lack of evidence that this can be implemented at a low cost, operate at low power, quickly and reliably. The IEEE 802.15.4A standard does not consider it possible in its standard mode. Therefore this feature was not considered in this work either.

Many protocols presented here discuss energy efficiency. However, few consider specifically sensor networks and as such their main focus is on bandwidth efficiency and high data rates. And most solutions are based on custom PHY layers, which make them difficult or impossible to apply on a standard compliant IEEE 802.15.4A UWB PHY. Therefore, relatively few ideas can be reused in the design of a novel ultra low power UWB MAC protocol, which must instead be based on narrowband ULP MAC research. We believe that the following ideas provide significant performance improvements, at low complexity and implementation costs:

- transmission at maximum power levels [86] (see point 1 above for justification);
- Time hopping for interference robustness [150];
- Operation without relying on a pulse sensing device or any other CCA mechanism.

3.4 WideMac

UWB-IR differs in many ways from narrowband radios. Its impulsive nature prevents channel polling, and thus classical collision avoidance techniques are irrelevant (this is why the IEEE 802.15.4 MAC is adapted to use ALOHA instead of CSMA for the IEEE 802.15.4A UWB PHY). Its high processing gain facilitates simultaneous multiple access, and the radio power consumption is much higher in reception than in transmission [133]. These particular characteristics make the design of a novel protocol worth investigating.

Subsection 3.4.1 presents the key ideas behind our protocol design, while subsection 3.4.2 defines the operation mode of WideMac and introduces its key

parameters. Subsection 3.4.3 discusses some aspects related to signaling traffic and subsection 3.4.4 focuses on the backoff algorithm.

3.4.1 WideMac Design Rationale

The first thing to decide when discussing a novel medium access control protocol is whether the medium access should be determined by a scheduling algorithm or by a contention resolution algorithm. One of the claimed advantages of scheduled protocols in ultra low power applications is that they don't need acknowledgments as they guarantee collision-free medium access. However, the dynamic and open nature of the wireless channel makes, in our view, acknowledgments mandatory. If a packet loss is not detected at that layer then a higher layer will have to do it. Such a shuffling of duties between the different protocol layers does not appear to be the best way to optimize the global system performance. Another advantage of scheduled protocols is the potentially higher use of the available bandwidth. This seems to be a stronger point, and indeed random access protocols typically only achieve a fraction of the physical bandwidth. However, this reasoning assumes an interference free channel, smooth and predictable traffic and no coexistence problem. These two aspects of communication integrity through acknowledgments and of robustness to interference and scalability with coexisting networks are the main reasons that lead us to adopt a random access approach.

Scheduled access protocols all suffer from signaling overhead which increases the base power consumption. They lack scalability when considering network size and density, and adaptability to low data rate network traffic. In addition, the problem of fully distributed synchronization is not easy to solve, especially when considering all involved physical effects.

A robust and low power MAC protocol for UWB-IR should:

- use random access to cope with interferers and to deal with coexistence issues,
- take advantage of the UWB high processing gain,
- exploit the low energy cost of transmissions as compared to emissions,
- use acknowledgments to guarantee the integrity of communications in all circumstances,
- avoid dependency on a network-wide time reference to improve scalability and robustness,
- share medium access among nodes fairly,
- efficiently use the energy and bandwidth resources.

3.4.2 WideMac Description

The WideMac protocol makes all nodes periodically (period T_W , identical for all nodes) and asynchronously wake up, transmit a beacon message announcing their availability and listen for transmission attempts during a brief time T_{Listen} . Figure 3.13 illustrates a single beacon transmission. It starts with a

known and detectable synchronization preamble colored in grey and is followed by a white colored data sequence which announces the node address and potentially other information, such as a neighbor list or routing table information (for instance, cost of its known path to the sink). A small listening time follows T_{Listen} , during which the node stays in reception mode and that allows it to receive a message.




Figure 3.13: Detailed view of a WideMac beacon.

When a node has a message to transmit, it first listens to the channel until it receives the beacon message of the destination node. This beacon message contains a backoff exponent value that must be used by all nodes when trying to access this destination. If this value is equal to zero, the source node can transmit immediately. Otherwise, it waits a random backoff time, waits for the destination beacon, and transmits its data packet. Because of the unreliability of the wireless channel, packets are acknowledged. If a packet is not acknowledged, or if the destination beacon was not received a retransmission procedure using the backoff algorithm is initiated, until the maximum number of retransmissions \maxTxAttempts is reached. The details of the backoff algorithm are described in subsection 3.4.4. Figure 3.14 depicts a sender node listening to the channel, ignoring the beacon message of another node, and sending its message to the destination after receiving its beacon. The exchange ends with an acknowledgment message transmitted by the receiver node and addressed to the sender node.




Figure 3.14: An initial WideMac data transmission.

Once a node has received a beacon from a neighbor node, it can predict the times of the next beacon emissions of this node since it knows the common beacon period. This helps to reduce the power consumption, but the accuracy of this information decreases over time because of the imprecision of the quartz Θ : after a time l , the node must listen during a time $4\Theta l$ (see [45] for details) and when $4\Theta l > T_W$ the information is too old to be useful. This procedure is similar to the reduction of wake-up preamble performed in WiseMAC [44], but here the source node is listening instead of transmitting, and the timing information is acquired through the beacon rather than through the acknowledgment.

Listening to the channel during a beacon period allows fast and simple neighbor discovery. This is useful for distributed routing protocols, allows beacon collision avoidance, and if performed regularly, allows relative clock drift learning (a method that reduces the apparent value of Θ to around 1 ppm). Problems related to beacon collision are studied in subsection 3.4.3.

By keeping nodes in sleep mode most of the time, WideMac saves energy. By using asynchronous wake-ups, it avoids costly and signaling traffic for network-wide time synchronization that would be sensitive to interference, and avoids the creation of traffic bottlenecks such as in S-MAC. The beacon emission allows to overcome the lack of carrier sensing capability, by letting candidate senders know when a destination node is available. Despite increasing slightly the node power consumption, it also enables fast neighbor discovery. This reduces the signaling traffic overhead at the routing layer and can overall reduce the system's total power consumption.

Table 3.3 compares WideMac system assumptions with those of various other wireless MAC protocols. WideMac minimalistic requirements are partly by design (multihop operation, no network-wide time synchronization, no out of band signaling channel, no listen-while-talk, no wake-up radio, adaptation to interference and coexisting networks, no traffic prediction) and partly due to the UWB-IR PHY layer (no CCA).

The idea of a receiver-initiated exchange for low-power sensor networks was already introduced in RICER [88] but with a dedicated signaling channel. WideMac proposes instead to use only one communication channel by relying on the robustness to interference of UWB-IR and on backoff algorithms. Synchronization preambles of beacon messages can, and will, collide, because of relative clock drift between nodes. The backoff algorithm and retransmission procedures can, however, preserve correct operation at the price of an increased mean delay. In addition, the effect of beacon collisions can be minimized. This problem is considered in more detail in subsection 3.4.3 and the different options for the backoff algorithm are studied in subsection 3.4.4.

3.4.3 WideMac Beacon Collisions

The risk of beacon collisions is an obvious problem of WideMac. Since all nodes periodically send signaling information, and since they do so asynchronously, even if they start at non-overlapping times the beacon transmissions will eventually collide due to clock drift, or nodes mobility. Thankfully, there are several ways to address this problem, by reducing its occurrence and by minimizing its consequences when it happens. But before considering them, we begin by evaluating the significance of the problem by directly applying a well-known result of probability theory, the Birthday Paradox (this result gives the probability that at least two persons in a group share their birthday) to obtain an analytical upper bound on the beacon collision probability. The results obtained with this model are then discussed and compared with network simulation results.

Since we consider here packet collisions over continuous time, and not birthdays uniformly distributed over a discrete set of days, we must make some simplifying assumptions:

1. Each beacon is vulnerable to collisions during a time t_{vuln} ;

Table 3.3: System assumptions of wireless MAC protocols.

	Full connectivity	Network-wide time sync	OOB Channel	Listen-while-talk	wake-up radio	CCA/Listen-Before-Talk	No interferer	No coexistence	No power constraints	Traffic predictability
ALOHA [5, 124]									●	
CSMA [122]						●			●	
CSMA/CD [122]				●		●			●	
MACA [17]						●			●	
MACAW [17]						●			●	
IEEE 802.11 Ad Hoc (IBSS)						●				
Bluetooth	●	●				●				
IEEE 802.15.4 beacon enabled mode [151]	●	●				●	●	●		
PicoNode [64]			●		●	●				
S-MAC [48]						●				
SCP-MAC [177]						●				●
WiseMAC [45], X-MAC [20], B-MAC [113]						●				
TRAMA [120], TDMA-W, L-MAC		●				●	●	●		
ISOMAC [179]										
Crankshaft [66]		●				●	●	●		
RI-MAC [157]						●				
WideMac [130]										

2. All beacons are equal: we do not consider heterogeneous received power levels, and power levels are high enough to guarantee good reception;
3. Time is discretized in slots of duration t_{vuln} , and thus there are $\lfloor t_{vuln}/T_W \rfloor$ such slots during a wake-up interval T_W ;
4. Beacons (of duration T_B) arrive uniformly and independently between 0 and T_W , and use one and only one slot of duration t_{vuln} ;

Since the radios cannot perform channel polling, and since transmissions are in practice unslotted, the vulnerability window is the same as for ALOHA, which was found equal to $2t$ (see figure 3.5a) where t is a packet duration (assuming equal packet sizes). However, with UWB-IR, we can make several hypotheses

on t , related to the robustness of the UWB-IR synchronization algorithm and to the burst demodulation circuit:

- worst case: the synchronization preamble requires a jam-free preamble, and is not robust to interference from another synchronization preamble and neither to burst transmissions. Thus it only succeeds in situations similar to figure 3.15a. In that case $t_{vuln} = 2T_B$ where T_B is the beacon duration;
- average case: the synchronization preamble requires a jam-free preamble but is robust to burst interference, as illustrated on figure 3.15b (i.e. it can synchronize on a SYNC preamble when data bursts are being transmitted but not when another synchronization preamble is interfering). In that case $t_{vuln} = 2t_{SYNC}$ where t_{SYNC} is the duration of the SYNC preamble;
- best case: the synchronization preamble requires only a fraction of the SYNC preamble to achieve synchronization, and is robust to both burst and SYNC interference. In that case $t_{vuln} = 2t_{minSYNC}$ where $t_{minSYNC}$ is the minimum jam-free duration of the SYNC preamble such that the receiver can synchronize on it. Figure 3.15c illustrates this type of interference;




Figure 3.15: The three possible cases of beacon collision from the SYNC algorithm point of view (SYNC sequence in grey, burst data modulation in white).

Due to a lack of available hardware and difficulties to obtain IEEE 802.15.4A UWB MATLAB simulation results, we cannot make an informed choice among these scenarios. We therefore adopted a conservative approach and evaluated only the first two cases: worst and average. In addition, we believe that experimentation with a real system would lead to a combination of the results from the three scenarios, depending on the particular channel conditions.

We can now express the probability of having at least one collision during one wake-up interval when considering N nodes, as a function of the wake-up interval T_W and the vulnerability window t_{vuln} :

$$P[\text{BeaconCollision}|N \text{ nodes}] = 1 - \prod_{k=1}^{N-1} \left(1 - k \left\lfloor \frac{t_{vuln}}{T_W} \right\rfloor \right). \quad (3.1)$$

It is equal to 1 minus the probability of no collision. We considered the IEEE 802.15.4A mandatory bit rate of 0.85 Mbps, the default synchronization preamble duration of 71.5 μs , a beacon packet size of 10 bytes (94 μs) and a total PHY overhead, including the preamble, of 137 μs (see table 2.5). Thus in the worst case, we have $t_{vuln}^{worst} = 2 \cdot 231 = 462 \mu\text{s}$ and in the average case, $t_{vuln}^{avg} = 2 \cdot 71.5 = 143 \mu\text{s}$.

Figures 3.16 and 3.17 show with a continuous green line the probability of observing at least one beacon collision as a function of the number of directly

reachable nodes for various values of the wake-up interval T_W in respectively the worst and average case scenarios. Numerical simulation results obtained with the Ghassemzadeh LOS channel model are shown with a dashed red line. 95% of the simulation results were in an interval around $\pm 5\%$ of the average value displayed on the figures, or closer.

Comparing the two figures leads to the following observations:

1. A reduction of the vulnerability period t_{vuln} lowers the collision probabilities;
2. Even with high wake up rates ($T_W = 50$ or 100 ms), the collision probability remains low when considering low network densities (5 nodes and less);
3. Average wake up rates ($T_W = 100$ or 200 ms) may allow operation with up to 20 reachable nodes;
4. The problem almost disappears when considering low wake up rates ($T_W = 1$ or 2 s) combined to medium network densities (up to 10 nodes).

In practice, the densities considered here are very high and atypical in sensor networks, in which a node can rarely communicate with more than 10 nodes. The extremely low power limits imposed on UWB systems also make such densities rather improbable. Such high values were nevertheless included for completeness, and because future applications may lead to different situations, such as with body area networks and more generally highly mobile networks.

This analysis is limited because it only considers the network without data transmissions, which would increase the collision probability, and because it does not consider heterogeneous power levels. In practice if there are collisions with beacons of varying power levels, a capture effect similar to what is observed with narrowband FM systems could happen, and one of the beacons may nevertheless be received. Such an effect would reduce the collision probability. The figures show only average results and assume random beacon arrival times. In practice two nodes that send their beacons close to each other would do that repeatedly, leading to systematic collisions. And most importantly, the mathematical model used here largely overestimates the collision probability by assuming only $\lfloor t_{vuln}/T_W \rfloor$ available beacon positions. In practice beacons can be emitted at any time, leading to lower collision rates, as illustrated with the simulation results in dashed red line on figure 3.17.

Therefore, we can conclude from this analytical discussion that while beacon collisions are a real problem, in a wide majority of the considered configurations their probability is low enough to allow the system to operate provided some additional adequate techniques.

We considered the following technical approaches to deal with beacon collision:

1. Improvements at the PHY layer: more robust and faster synchronization algorithms, and faster radio bit rates (the standard suggests bit rates as high as 31 Mbps) would further reduce the beacon collision probability. While we cannot control the future development of UWB-IR systems, the inclusion of relevant improvements in the IEEE 802.15.4A UWB PHY as

optional modes indicate that these evolutions have already been seriously considered.

2. Reliance on clock drift: since all nodes schedule their beacon transmission by using their own quartz crystal as clock reference, they all use a slightly different time. Common quartz have an accuracy of ± 30 part per million (ppm), and additional imprecisions lead to a global clock accuracy of ± 100 ppm. Figure 3.18 shows the overlapping duration of two beacons as a function of the relative clock drift Θ_r between two nodes ($t_{overlap} = \frac{t_{width}}{\Theta_r}$). Assuming uniform clock drift distribution, in most cases the problem disappears in less than a minute, but in some cases it can become very large (several minutes).
3. Randomized beacon emission times : RI-MAC [157] proposes to systematically choose a random time in the interval $[0.5T_W, 1.5T_W]$. While simple and effective, this approach has the disadvantage of preventing the prediction of the beacon transmission time, thereby increasing the power consumption.
4. Collision Avoidance: nodes may listen for other beacons before sending their first beacon, and choose an emission time that would not conflict with the others. This approach leads to a distributed time synchronization problem, which we want to avoid for stability, scalability, robustness and coexistence reasons.
5. Adapted transmission procedure: when listening for the destination beacon, if the beacon is not received after a wake-up interval then the source node should continue to listen, especially if some other beacons were received. A backoff waiting time may be used to let clock drift separate the beacons.
6. Rely on the strict regulations concerning the radiated power levels: while these limits, developed to protect already-deployed communication systems, are generally understood as inconvenient for UWB, they also protect UWB receivers from interference of other UWB systems as a low radiated power level reaches the noise level faster. This reasoning lead us to the hypothesis that 20-30 nodes would be a reasonable maximum number of directly reachable neighbors.
7. Decoupling the beacon emission and listening frequencies: instead of sending a beacon at each wake-up, only send it once every 2,3...10 wake-ups. This allows to significantly reduce the beacon collision rate while keeping the average latency low, despite a higher latency for the first packet exchange.

We propose to adopt a solution built on points 2, 5 and 6: clock drift, randomized beacon emission times, adapted transmission procedure and reliance on side-effects of UWB regulations. We do not consider beacon collision avoidance due to potentially high complexity, instability and doubts on robustness. Decoupling the Beacon emission and listening rates is more attractive, but as this technique does not prevent beacon collisions it cannot solve the problem alone, while it clearly increases the system complexity. Further, such

a technique could be adopted at a later stage once the entire system has been developed and evaluated.

By slightly randomizing the beacon emission times, and sending the beacon at a time $T_W + \text{uniform}(-k_{dev}T_W, +k_{dev}T_w)$, where *uniform* takes a random number between its bounds, we can avoid the long-term disappearance of a node due to another beacon preceding its own beacon, and make the node at least intermittently reachable. By using a random time much smaller than what was proposed in RI-MAC, the power consumption can still be reduced by beacon emission prediction (at a slight increase depending on the exact value of k_{dev}). The transmission procedure should be aware that the destination beacon might be jammed and must listen for the beacon during several periods T_W before dropping the packet. We set this listening interval as an integer multiple of the wake-up interval: $k_{TxL}T_W$, where $k_{TxL} \in \mathbb{N}_{0,1}$ is an integer number larger than 1 and must be chosen as a function of the beacon collision probability. Finally, the strict radiated power limits combined to adequate wake-up intervals allow one to deploy a system without fearing high beacon collision rates that would render the system unusable.




Figure 3.16: Beacon collision probabilities for the worst case synchronization algorithm scenario (default $71.5 \mu s$ SYNC preamble duration, 0.85 Mbps mandatory bit rate).




Figure 3.17: Beacon collision probabilities for the average case synchronization algorithm scenario. The analytical model overestimates the error rate obtained with the network simulator (default $71.5 \mu s$ SYNC preamble duration, 0.85 Mbps mandatory bit rate, 95% confidence interval around 5% of the plotted mean values for the simulation results).




Figure 3.18: Clock drift makes WideMac beacon collisions unavoidable and non-persistent. The time during which successive beacon collisions happen depends on the relative clock drift between two nodes and on the vulnerability window t_{vuln} .

3.4.4 WideMac Backoff Algorithm

The backoff algorithm has a major effect on collision, latency and fairness. WideMac periodic beacons allow the sender nodes to get some information on the channel state at the destination. This can be used to reduce the hidden and exposed terminal problems.

The WideMac transmission procedure works as follows: a candidate sender node first listens for the receiver node's beacon. Once it finds it, it can either immediately attempt transmission (default for lightly loaded networks) or it can start a backoff timer before sending (this is activated by a flag *alwaysBackoff* in the beacon). In both cases, the sender node waits for an acknowledgment. If it does not arrive, a retransmission procedure begins. The sender node chooses a random time parametrized by the receiver node's Backoff Exponent (BE) which was broadcast in the beacon, using a binary exponential backoff:

$$T_{Backoff} = N_{Backoff} \cdot T_W, \text{ where } N_{Backoff} \in [0, 2^{BE_{Receiver}} - 1].$$

The backoff time is thus a function of the wake-up interval T_W and of the channel state at the receiver node, as captured by $BE_{Receiver}$. Such a receiver-based backoff parametrization was also proposed in RI-MAC [157]. The use of a slotted backoff time based on T_W is natural since all candidate sender nodes are synchronized on the receiver node's wake up times: using a fraction of T_W would not change anything as the node would not transmit before receiving the destination beacon. Using an integer multiple of T_W for the unit backoff duration would increase latency and spread the traffic, but this can also be achieved by adapting the value of $BE_{Receiver}$ to the traffic conditions.

The value of $BE_{Receiver}$ should reflect the current congestion at a node. Several algorithms exist:

- CSMA's binary exponential backoff (BEB) algorithm, which leads to rapid variations of the backoff exponent, since its value doubles at each unsuccessful transmission attempt and is reinitialized at its minimum for each new packet to transmit;
- MACAW's Multiplicative Increase and Linear Decrease (MILD), which multiplies the backoff window by a constant to increase it and subtracts another constant to the backoff window to decrease it, leading to smoother variations;
- Using a Sift distribution as in Crankshaft [66], which was shown to be optimal in collision reduction when the network size is known, but requires carrier sensing;
- Adapting the BEB algorithm so that the backoff exponent is doubled at every other transmission failure, leading to more stable values, as proposed in the MedWin proposal at the IEEE 802.15.6 task group on medical body area networks [155].

In RI-MAC, the authors chose to use the BEB algorithm. However, with IEEE 802.15.4A, it is more difficult for the destination node to detect collisions. Thus, we adopt a combination of the RI-MAC, MACAW and MedWin approaches:

the initial backoff exponent value for the transmission is taken from the destination beacon. Upon transmission failure (i.e. missing ACK), it is the source node that increases its backoff exponent. It adopts the MedWin strategy of doubling it at every other failure in order to minimize latency. The node backoff exponent (which we call local BE, or BE_{local}) is initially set to a minimum value BE_{min} which can be equal to zero. As in the MedWin proposal, at every other reception failure, thus when the radio synchronized on a frame but could not decode it (invalid checksum), BE_{local} is doubled. BE_{local} is decreased by 1 at each successful frame reception. Thus, the local backoff exponent evolves similarly to MACAW's MILD strategy.

3.5 Store and Forward Analytical Power Consumption Analysis

Recent ultra low power wireless MAC protocols tend to be complex, to the point that they cannot be completely described in a conference publication. And while they can be described in a few sentences, their performance can be greatly influenced by implementation-time design choices, parameter values, topology and traffic configurations. The few available protocol comparisons on hardware platforms have mostly used publicly available implementations in the TinyOS sensor network programming framework, to reduce the implementation effort. Even in this context, very few protocols are available.

While analytical comparisons make simplifying or restricting assumptions, most notably at the physical layer, they can include more protocols. This makes them suitable for a first comparative evaluation, before implementation in a network simulator or in a real sensor network. This section presents analytical models of power consumption for various scheduled and random access protocols in a store and forward scenario, as illustrated on figure 3.19. These models are used in subsection 3.5.3 to evaluate the effect of respectively traffic rate and network density on the power consumption. This section is followed by analytical models of latency in section 3.6.




Figure 3.19: Store and Forward traffic: node 2 receives a packet from node 1, stores it, and forwards it to node 3.

3.5.1 Models Objectives and Assumptions

This section focuses on simple store-and-forward functionality, typical of a sensor node in a sparsely populated network and part of a routing tree. The models developed in the following subsections all make the following four assumptions:

1. *Full Connectivity*: the network is composed of N devices with full connectivity. This allows to exclude routing problematics and to consider overhearing.

2. *Store and Forward Traffic*: each node periodically receives a packet and forwards it to another node. The packet arrival rate is Poisson distributed with parameter λ : $P[NbArrivals = k] = \frac{e^{-\lambda}\lambda^k}{k!}$. Therefore the mean time between two packet arrivals (from the same source) is $\lambda^{-1} = L$. The packets destinations are chosen in the network so that each node receives as much data as it is transmitting, and the destinations do not change during network operation.
3. *The radio is the dominant energy user*: while the MAC processing requires some energy from the microcontroller, we assume that it is negligible when compared to the radio power consumption. We can thus estimate the power consumption of the communication subsystem by calculating how much time the radio spends in each of its modes.
4. *Detailed radio state machine*: since the power consumption is derived from the time spent in each of the radio modes, it is important to model these accurately. We use the finite state machine illustrated on figure 3.20, with three steady states Sleep, Rx and Tx, and four transient states SetupRx, SetupTx, SwitchRxTx and SwitchTxRx. The radio can always leave any state (steady or transient) and immediately enter sleep mode. The time spent in a transient state is a constant $T_{TrState}$, the power consumption in each state is P_{State} and the energy cost of a transition from one steady state to another is $E_{TrState}$. The values used for evaluation in the next section were taken from [150] and from [133]. The radio state machine energy and timing parameters are regrouped in table 3.4.
5. *Packet-based radio*: the radio does not trigger interrupt requests to the microcontroller at each byte reception. Instead, it receives the whole frame, eventually applies error correction codes and then delivers it to the microcontroller.

Some protocol models make the following additional assumptions related to timing and modulation:

1. *Clock drift*: as each node uses its own time source, based on a quartz crystal, there are inevitable variations between the nodes. This can influence the communication protocol's operation. A quartz precision θ is given in parts per million (ppm), and typical quartz have a precision of 30 ppm. This effect is taken into account for random access protocols.
2. *Network-wide synchronization*: TDMA protocols assume that they can maintain accurate network-wide synchronization. We assume in this section that this hypothesis holds.
3. *CCA mode 6*: some Low Power Listening protocols that depend on listen-before-talk (LBT) where included in this comparison, assuming that an optional mode of the standard was used in that case to enable their operation. This mode interleaves detectable synchronization preamble sequences with undetectable data sequences, thereby increasing the frame length so that it may be detected. To our knowledge, no radio able to detect such frames has ever been realized. We included these protocols here for completeness, although we do not expect them to work on future UWB-IR systems.

Data packet size is set to 50 bytes ($T_M = T_{SYNC} + T_{Data-M} + T_{RS} = 0.542$ ms) and acknowledgment messages and beacon messages are 4 bytes long ($T_{Ack} = T_{SYNC} + T_{Data-Ack} + T_{RS} = 0.11$ ms).

Table 3.4: Radio power consumption and timing parameters.

Parameter	Value	Parameter	Value
P_{Rx}	30 mW	$T_{SetupRx}$	1 ms
P_{Tx}	1 mW	$T_{SetupTx}$	1 ms
P_{Sleep}	60 μ W	T_{SwTxRx}	22 μ s
$P_{SetupTx}$	1 mW	T_{SwRxTx}	10 μ s
$P_{SetupRx}$	30 mW	T_{CCA}	1 ms
P_{SwTxRx}	15 mW	T_{Sync}	71 μ s
P_{SwRxTx}	15 mW	θ	100 ppm
		Bit rate	0.85 Mbps

(a) Energy parameters.

(b) Timing parameters.




Figure 3.20: Detailed radio model including transient states.

3.5.2 Power Consumption Models

Using the previously defined theoretical models of radio and network traffic, the mean power consumption of a protocol can be evaluated by computing the mean time spent in each state of the radio. Protocol parameters are chosen so that latencies in low data rate conditions are similar: the periods of WiseMAC, WideMac and Crankshaft are all set to 500 ms, except for SCP-MAC which always uses an optimal period function of message rate. This last hypothesis was necessary to avoid choosing unfavorable values for SCP-MAC parameters.

Since all considered asynchronous random access protocols produced extremely similar results, and to improve readability of the figures, a close lower bound named Optimal Preamble Sampling is used instead.

Ideal Protocol

On average, during a time L , a radio using the ideal protocol must receive a packet, send a packet, and sleep the rest of the time. It costs an energy

$E_{SetupRx}$ to enter reception mode, an energy $T_M P_{Rx}$ to receive the incoming packet, an energy $E_{SetupTx}$ to enter transmission mode and an energy $T_M P_{Tx}$ to retransmit the packet. During the rest of the period $T_L - T_{SetupRx} - T_M - T_{SetupTx} - T_M$, the node is in sleep mode with a power consumption P_{Sleep} . This leads to the following expression of the average power consumption:

$$\begin{aligned} \overline{P_\lambda^{Opt}} &= \frac{1}{L} [T_M(P_{Tx} + P_{Rx}) \\ &+ E_{SetupTx} + E_{SetupRx} \\ &+ (L - 2T_M - T_{SetupRx} - T_{SetupTx})P_{Sleep}]. \end{aligned} \quad (3.2)$$

WideMac

During each wake-up interval T_W , a node must enter transmission mode (cost $E_{SetupTx}$), transmit its beacon (cost $T_{Beacon}P_{Tx}$), switch to reception mode (cost E_{SwRxTx}) and attempt a packet reception (cost $T_{Listen}P_{Rx}$). These costs are regrouped in the beacon energy E_{Beacon} . In addition, it must sometimes (on average $\frac{T_W}{L}$ times per wake-up interval) transmit a packet ($\overline{E_{Trans}}$) or receive one ($\overline{E_{Recv}}$), and sleep the rest of the time ($\overline{E_{Sleep}}$), giving the following average power consumption:

$$\overline{P_\lambda^{WideMac}} = \frac{1}{T_W} (E_{Beacon} + \overline{E_{Trans}} + \overline{E_{Recv}} + \overline{E_{Sleep}}), \quad (3.3)$$

where:

$$\begin{aligned} E_{Beacon} &= E_{SetupTx} + T_{Beacon}P_{Tx} + E_{SwTxRx} + T_{Listen}P_{Rx} \\ \overline{E_{Trans}} &= \frac{T_W}{L} [E_{SetupRx} + \overline{T_{BD}}P_{Rx} + E_{SwRxTx} \\ &+ T_M P_{Tx} + E_{SwTxRx} + T_{Ack}P_{Rx}] \\ \overline{E_{Recv}} &= \frac{T_W}{L} [(T_M - \overline{T_{Listen}})P_{Rx} + E_{SwRxTx} + T_{Ack}P_{Tx}] \\ \overline{E_{Sleep}} &= T_W - [T_{SetupTx} + T_{Beacon} + T_{SwTxRx} + T_{Listen} \\ &+ \frac{T_W}{L} (T_{SetupRx} + \overline{T_{BD}} + 2T_{SwRxTx} + 2T_M - \overline{T_{Listen}} \\ &+ T_{SwTxRx} + 2T_{Ack})]. \end{aligned}$$

$\overline{E_{Trans}}$ is the mean energy required to retransmit the received packets. To forward a packet, the node must first enter reception mode (cost $E_{SetupRx}$), listen for the destination beacon during a time $\overline{T_{BD}}$ (cost $\overline{T_{BD}}P_{Rx}$), switch to transmission mode (cost E_{SwRxTx}), send the packet (cost $T_M P_{Tx}$), switch back to reception mode (cost E_{SwTxRx}) and receive the acknowledgment ($T_{Ack}P_{Rx}$). The mean beacon detection time is given by a weighted sum, over all possible values of the time l between two packets (distributed as an exponential random variable), of the time required to receive the beacon. This time is upper bounded by T_W since this is the time between two beacons. It is also a function of the time l when the last packet was sent to the destination, and of the clock drift θ (if there is no clock drift, or perfect synchronization, $\theta = 0$ and $\overline{T_{BD}}$ is equal to zero since the next beacon emission time can be predicted with infinite accuracy) :

$$\overline{T_{BD}} = \int_0^{+\infty} \min(4\theta l, T_W) \frac{1}{L} e^{-l/L} dl = 4\theta L \left(1 - e^{-\frac{T_W}{4\theta L}}\right).$$

$\overline{E_{Recv}}$ is the mean energy required to receive the packets after transmitting the beacon. It consists of a term $(T_M - \overline{T_{Listen}}) P_{Rx}$ to account for the actual packet receptions (reduced to avoid counting the time T_{Listen} twice since it is already included in E_{Beacon}), and the terms E_{SwRxTx} and $T_{Ack} P_{Tx}$ to represent respectively the radio switching energy and the acknowledgment transmission.

Finally, the mean time spent in sleep mode $\overline{E_{Sleep}}$ is obtained by subtracting the times spent in all other states from the beacon period. This model makes the hypothesis that the wake-up interval T_W is sufficiently large so that we can neglect beacon collisions.

S-MAC / T-MAC

During a T-MAC frame T_W , if the periodic cost of synchronization is neglected, a radio using this protocol must listen for one empty RTS-CTS exchange (the last of the frame) and participate to $\frac{NT_W}{L}$ non-empty RTS-CTS exchanges. Among these exchanges, each node will send and receive $\frac{T_W}{L}$ packets. If, on average, the packet exchange frequency (N/L) is lower than the wake-up frequency ($1/T_W$), then there are “free” rts-cts periods. The expression $1 + NT_W/L 1_{\{NT_W/L > 1\}}$ accounts for those, with $1_{\{NT_W/L > 1\}}$ being the indicator function, equal to 1 when the indexed condition is true and equal to zero otherwise. This leads to the following equations:

$$\overline{P_{\lambda}^{S-MAC}} = \frac{(1 + NT_W/L 1_{\{NT_W/L > 1\}}) E_{rtscts} + (E_{send} + E_{recv}) T_W/L + E_Z}{T_W} \quad (3.4)$$

$$\begin{aligned} E_{rtscts} &= E_{SetupRx} + P_{Rx} T_{rtscts} \\ E_{send} &= E_{SwRxTx} + T_M P_{Tx} + E_{SwTxRx} + T_{Ack} P_{Rx} \\ E_{recv} &= (T_{SwRxTx} + T_M) P_{Rx} + E_{SwRxTx} + T_{Ack} P_{Tx} \\ E_Z &= [T_W - (1 + NT_W/L 1_{\{NT_W/L > 1\}}) T_{rtscts} \\ &\quad - 2(T_M + T_{Ack} + 2T_{SwTxRx}) T_W/L] P_{Sleep}. \end{aligned}$$

To keep the mathematical expressions simple, the power consumption during RTS-CTS exchange is always expressed as follows : $E_{rtscts} = P_{Rx} T_{rtscts}$. This means that we neglect the small power consumption variations due to the node transmitting an RTS or a CTS. This is true if a node does not participate to most RTS-CTS exchanges (large values of N, low traffic or both), or if the RTS-CTS period is significantly larger than the RTS-CTS messages ($T_{rtscts} \gg T_{RTS}$). To allow efficient operation of the protocol, this should be the case anyway, in order to avoid repeated collisions of the signaling traffic.

SCP-MAC

The following model of SCP-MAC’s power consumption makes three additional hypotheses:

1. the polling time is always set to its optimal value, supposing exact knowledge of the traffic rate ($T_{Poll} = \frac{L}{N}$),
2. the cost of network-wide time synchronization is neglected (see additional hypothesis number 2 in subsection 3.5.1),
3. failed channel accesses during contention windows are neglected.

A lower bound on SCP-MAC's average power consumption can be expressed as:

$$\overline{P_{\lambda}^{SCP}} = \frac{1}{L} \left[\overline{E_{Poll}^{SCP}} + \overline{E_{Recv}^{SCP}} + \overline{E_{Trans}^{SCP}} + \overline{E_Z^{SCP}} \right] \quad (3.5)$$

where

$$\overline{E_{Poll}^{SCP}} = (N-1)(E_{SetupRx} + T_{CCA6}P_{Rx}) + E_{SetupTx} + \frac{T_{CW}^{SCP}}{2}P_{Tx}$$

is the mean energy required to perform the periodic synchronous channel polling,

$$\overline{E_{Recv}^{SCP}} = T_M P_{Rx} + (N-1)\overline{T_O^{SCP}} P_{Rx}$$

is the mean energy required to receive the packets, including systematic overhearing of all other packets (due to the synchronous wake-up of SCP-MAC),

$$\overline{E_{Trans}^{SCP}} = T_M P_{Tx}$$

is the average cost of retransmitting the received packet and

$$\overline{E_Z^{SCP}} = P_Z \left\{ L - \left[(N-1)(T_{SetupRx} + T_{CCA6} + \overline{T_O^{SCP}}) + T_{SetupTx} + \frac{T_{CW}^{SCP}}{2} + 2T_M \right] \right\}$$

is the energy cost due to the time spent in sleep mode.

SCP-MAC authors considered that it is possible to stop receiving a frame as soon as the header has been decoded. In that case, we have $\overline{T_O^{SCP}} = T_{SYNC}$. This work considers instead a modern packet based radio with which the complete frame must be received before further processing. This leads to $\overline{T_O^{SCP}} = T_M$.

Crankshaft

During a Crankshaft frame T_W , a node must perform on average: one polling on its slot (with an energy cost E_{CCA}), transmit and receive data (respective average costs $\overline{E_{Recv}}$ and $\overline{E_{Send}}$), listen to broadcast slots (cost E_{BCast}) and sleep the rest of the time (average cost $\overline{E_Z}$). The expression of the Crankshaft average power consumption is:

$$\overline{P_{\lambda}^{CS}} = \frac{1}{T_W} \left[E_{CCA} + \frac{T_W}{L} (\overline{E_{Recv}} + \overline{E_{Send}}) + E_{BCast} + \overline{E_Z} \right] \quad (3.6)$$

where:

$$E_{CCA} = E_{SetupRx} + T_{CCA6}P_{Rx}$$

is composed of the energy to switch the radio into reception mode and of the energy required to perform the Clear Channel Assessment mode 6,

$$\overline{E_{Recv}} = T_M P_{Rx} + E_{SwRxTx} + T_{Ack} P_{Tx}$$

is composed of the energy of the frame reception $T_M P_{Rx}$, the energy to switch the radio into transmission mode E_{SwRxTx} and the acknowledgment transmission $T_{Ack} P_{Tx}$,

$$\overline{E_{Send}} = E_{SetupTx} + T_{CW} P_{Tx} + T_M P_{Tx} + E_{SwTxRx} + T_{Ack} P_{Rx}$$

regroups the energy to enter transmission mode, the energy spent contending for the medium $T_{CW} P_{Tx}$, the message transmission cost $T_M P_{Tx}$, the switching cost E_{SwTxRx} and the acknowledgment reception $T_{Ack} P_{Rx}$.

$$\overline{E_Z} = P_Z \left[T_W - (T_{SetupRx} + T_{CCA} + \frac{T_W}{L} (2T_M + T_{SwRxTx} + 2T_{Ack} + T_{SetupTx} + T_{CW} + T_{SwTxRx})) \right]$$

is the sleep mode power consumption multiplied by the time spent in that mode.

As the traffic scenario focuses on unicast traffic, we simply included one broadcast slot on which actual traffic is neglected. The energy cost of the broadcast slot is equal to the cost of setting the radio into reception mode and listening for a brief period of time to an empty channel: $E_{BCast} = E_{SetupRx} + T_{CCA} P_{Rx}$. The expected contention time T_{CW} is chosen proportionally to the values used in [66], in which $T_{CW} = 9.15$ ms with a radio bit rate of 61 kbps. As we consider a radio with a bit rate of 850 kbps, we divide T_{CW} proportionally and set it to 0.65 ms.

B-MAC

Since it is a Listen Before Talk protocol, B-MAC requires the multiplexing of synchronization preamble sequences with the data frame, and a radio that can accurately and quickly (in a time T_{CCA6}) detect these sequences.

B-MAC average power consumption is given by:

$$\overline{P_{\lambda}^{BMAC}} = \frac{1}{T_W} \left(E_{CCA}^{BMAC} + \overline{E_{Recept}^{BMAC}} + \overline{E_{Trans}^{BMAC}} + \overline{E_Z^{BMAC}} \right) \quad (3.7)$$

where E_{CCA}^{BMAC} is the energy required for the clear channel assessment in mode 6 and is equal to:

$$E_{CCA}^{BMAC} = E_{SetupRx} + T_{CCA6} P_{Rx},$$

$$\begin{aligned} \overline{E_{Recept}^{BMAC}} &= \frac{T_W}{L} \left[\left(T_M + \frac{T_W}{2} \right) P_{Rx} + E_{SwRxTx} + T_{Ack} P_{Tx} \right] \\ &+ (N - 2) \frac{T_W}{L} \left(\frac{T_W}{2} + T_M \right) P_{Rx} \end{aligned}$$

is the average energy required to listen to the wake-up preamble, receive the packets and acknowledge them, and to overhear $N - 2$ transmissions (from all nodes except our own transmissions and those addressed to us),

$$\overline{E_{Trans}^{BMAC}} = \frac{T_W}{L} [E_{SetupTx} + (T_W + T_M) P_{Tx} + E_{SwTxRx} + T_{Ack} P_{Rx}]$$

is the average energy needed for the transmissions of one packet, and

$$\begin{aligned} \overline{E_Z^{BMAC}} &= P_Z \left[T_W - T_{SetupRx} - T_{CCA6} - \frac{T_W}{L} \left(\frac{T_W}{2} + T_M + T_{SetupRx} \right. \right. \\ &\quad \left. \left. + T_{SwRxTx} + T_{Ack} + (N - 2) \left(\frac{T_W}{2} + T_M \right) \right) \right. \\ &\quad \left. + \frac{T_W}{L} (T_{SetupTx} + T_W + T_M + T_{SwTxRx} + T_{Ack}) \right] \end{aligned}$$

is the average energy lost in sleep mode.

WiseMAC, SyncWUF, CSMA-MPS, X-MAC and Optimal Preamble Sampling

For completeness, we also considered recent ultra low power random access protocols that depend on carrier sensing, even though this feature is not available with UWB-IR.

WiseMAC, SyncWUF, CSMA-MPS and X-MAC all make use of adaptive wake-up preambles. Separate models were developed for each one. Since they performed extremely closely in the evaluation [129], a unique lower bound on their power consumption is considered instead. It is established by considering a zero length wake-up preamble. This lower bound is called here Optimal Preamble Sampling.

A node using one of these protocols must, during each period T_W , perform one carrier sensing (E_{CCA}^{OptPS}), eventually spend some energy receiving data (E_{Recept}^{OptPS}), some more to transmit data (E_{Trans}^{OptPS}) and spend the rest of the time in sleep mode (E_Z^{OptPS}). This leads to the following expression:

$$\overline{P_\lambda^{OptPS}} = \frac{1}{T_W} \left(E_{CCA}^{PS} + \overline{E_{Recept}^{PS}} + \overline{E_{Trans}^{PS}} + \overline{E_Z^{PS}} \right)$$

where:

$$\begin{aligned}
E_{CCA}^{PS} &= T_{CS}P_{Rx} + E_{SetupRx} \\
\overline{E_{Recept}^{PS}} &= \frac{T_W}{L} \left[(\overline{T_{LP}^{PS}} + T_M)P_{Rx} + E_{SetupTx} + T_{Ack}P_{Tx} \right] \\
&\quad + (N-2)\frac{T_W}{L}\overline{T_O^{PS}}P_{Rx} \\
\overline{E_{Trans}^{PS}} &= \frac{T_W}{L} \left[E_{SetupTx} + (\overline{T_{MR}^{PS}} + \overline{T_{CDC}^{PS}} + T_M)P_{Tx} \right] \\
\overline{E_Z^{PS}} &= P_Z \left[T_W - T_{SetupRx} - T_{CS} - \frac{T_W}{L} (\overline{T_{LP}^{PS}} + T_M + T_{SetupTx} \right. \\
&\quad \left. + T_{Ack} + (N-2)\overline{T_O^{PS}}) - \frac{T_W}{L} (T_{SetupTx} + \overline{T_{MR}^{PS}} \right. \\
&\quad \left. + \overline{T_{CDC}^{PS}} + T_M + T_{SwTxRx} + T_{Ack}) \right]
\end{aligned}$$

where the expected time spent listening to the wake-up preamble $\overline{T_{LP}^{PS}}$ and the expected time spent in overhearing $\overline{T_O^{PS}}$ are both set to zero.

When using this model to evaluate WiseMAC, the wake-up preamble expected duration is set to $\overline{T_{CDC}^{PS}} = 2\theta L \left(1 - e^{-\frac{T_W}{4\theta L}}\right)$ [44, 94]. When evaluating optimal preamble sampling, it is set to zero.

3.5.3 Results Analysis

Figures 3.21 and 3.22 show the power consumption of the ideal protocol, WideMac, S-MAC, optimal preamble sampling and Crankshaft as a function of the mean time L between packets. The protocols duty-cycle periods (T_W) vary between 100 ms and 1 s, and the network size varies between 2 and 20 nodes. Clock drift is considered and equal to $\Theta = 100 \cdot 10^{-6}$. Radio parameter values can be found in table 3.4.

All protocols show the same qualitative behavior as the ideal protocol. For high traffic, the power consumption increases abruptly, and for very low traffic, the power consumption stabilizes around a minimum. In the case of the ideal protocol, the occasional cost of transmission and reception becomes negligible in low traffic compared to the sleep mode power consumption. The ideal power consumption therefore becomes equal to the sleep mode power consumption, at 60 μ W. The other protocols also reach a lower limit on their power consumption, always higher than the ideal.

For WideMac, the deviation from ideality is caused by the periodic beacon transmission and by the periodic idle listening. This deviation decreases for higher values of T_W since this decreases the frequency of the protocol operating overhead.

Optimal preamble sampling power consumption is just a bit lower than WideMac. The difference is mainly due to the periodic beacon emission cost and to the beacon detection time when attempting transmission. The gap between optimal preamble sensing and ideal power consumption is due to the periodic idle listening. As all ULP MAC protocols that we are aware of require devices to periodically listen to the channel, the performance of Optimal Preamble Sampling may be seen as a more realistic lower bound on a protocol power consumption. Real preamble sampling protocols must, in addition

to the cost of periodic listening, also account for preamble transmission and overhearing. As the results for WiseMAC were very close to OptPS, and as it depends on CCA, which is unavailable on UWB-IR transceivers, WiseMAC was omitted from the figures to improve readability.

Crankshaft leads to slightly higher power consumption levels. This is due to the periodic channel polling on the broadcast slot. This evaluation does not capture the energy cost of the time synchronization required for the correct operation of Crankshaft. A real implementation of Crankshaft would use more energy, and this time synchronization cost would probably increase with the number of nodes.

The worst performer of this study is T-MAC. It suffers from its synchronous operation, which leads to numerous idle listenings and overhearings.

We conclude this analytical study of the power consumption with the following key observations:

- recent ultra low power MAC protocols allow to reach close to ideal power consumption levels;
- in absence of CCA, WideMac reaches the best performance;
- synchronous operation can significantly degrade the power consumption;
- Scheduled protocols (TDMA) do not perform better from an energy viewpoint than random access protocols, even when ignoring the cost of time synchronization or coexistence and interference problems;
- periodic beacon transmissions can be performed at minimal energy cost with UWB-IR thanks to the extremely low cost of transmissions.

These findings validate our approach of using periodic beacons to enable ultra low power and scalable communications in absence of a carrier.




Figure 3.21: Analytical power consumption of various MAC protocols with an UWB-IR radio in a simple Store and Forward traffic scenario, as a function of the mean traffic inter-arrival time L , for various values of the duty-cycle period T_W and various number of nodes.




Figure 3.22: Analytical power consumption of various MAC protocols with an UWB-IR radio in a simple Store and Forward traffic scenario, as a function of the mean traffic inter-arrival time L , for various values of the duty-cycle period T_W and various number of nodes.

3.6 Star-Topology Analytical Latency Analysis

This section builds and uses analytical models of latency for several MAC protocols. This allows to compare their performance in various scenarios and develop a better understanding of the possible trade-offs. The drawbacks and advantages of analytical latency models, when compared to network simulation, are the same as for analytical power consumption models, as presented in section 3.5. Basically, they allow to compare a wider range of protocols at the price of sometimes over simplifying assumptions.

This section is organized as follows. Subsection 3.6.1 presents analytical models of latency for several MAC protocols, and states the assumptions behind those. Subsection 3.6.2 presents validation results for some of these models, compared with experimental measurements and network simulation models. Finally, subsection 3.6.3 discusses latency results obtained with these models.

3.6.1 Assumptions and Latency Models

An analysis of the average latency of these protocols can be made by using standard results of queuing theory [15], assuming some simplifying hypotheses.

Considering that:

- packets are generated at each sensor device according to a Poisson process of parameter λ_s ,
- and that MAC contention resolution mechanisms do not significantly modify this distribution,

the sum of these N independent and identically distributed random point processes also follows a Poisson distribution, of parameter $\lambda = N\lambda_s$. If packets are of constant size, the service time is deterministic. And if we assume infinite buffer capacity, the queuing model $M/D/1/\infty$ can be applied.

The average delay of a $M/D/1/\infty$ system is given by:

$$E[Delay] = \frac{1}{\mu} + \frac{\rho}{2\mu(1-\rho)} \quad (3.8)$$

where $\rho = \lambda/\mu$ is the traffic intensity, $\lambda = N\lambda_s$ is the packet rate and μ^{-1} is the service time.

This model allows to compute the average end-to-end packet latency with a variable number of traffic sources, provided that all traffic sources generate the same Poisson traffic.

The Ideal Protocol

The service time for the ideal protocol is given by:

$$\mu_{ideal}^{-1} = T_{SetupTx} + T_M$$

as the node only has to enter transmission mode and send the message. It does not have to wait before transmitting since the destination is immediately aware of the incoming packet, and it does not listen for an acknowledgment since transmissions are always successful.

CSMA

For CSMA, the service time is given by:

$$\mu_{CSMA}^{-1} = \overline{T_{Backoff}} + T_{CCA} + T_{SwRxTx} + T_M + T_{SIFS} + T_{Ack}$$

where $\overline{T_{Backoff}}$ is the mean time spent in backoff, T_{CCA} is the time needed to perform a clear channel assessment, T_M is the time required to transmit the message, T_{SIFS} is the time between the end of a packet transmission and the end of the transmission of an acknowledgment (it must be larger than the radio switching time) and T_{Ack} is the time needed to transmit the acknowledgment.

The mean time spent in backoff can be computed by assuming an always idle channel:

$$\overline{T_{Backoff}} = \frac{1}{2} (2^{\min BE} - 1) aUnitBackoffSlot.$$

This holds only if there is almost no interfering traffic.

WideMac

For WideMac and at low data rates,

$$\mu_{WideMac}^{-1} = \frac{T_W}{2} + \overline{T_{Backoff}^{WideMac}} + T_M + T_{SwTxRx} + T_{Ack}$$

since the packet can arrive at the MAC layer randomly at any time during the destination node sleep interval. This is a lower bound rather than the exact value for two reasons:

- the service time distribution is not deterministic but uniform in the interval $[0, T_W]$;
- a deterministic service time is the optimal distribution for minimizing the mean time spent in the system (for a constant service time μ).

The WideMac mean backoff time is given by:

$$\begin{aligned} \overline{T_{Backoff}^{WideMac}} &= \overline{N_{Backoff}} \cdot T_W \\ &= \frac{1}{2} (2^{\min BE} - 1) T_W. \end{aligned}$$

Crankshaft

The Crankshaft mean service time can be expressed as:

$$\mu_{Crankshaft}^{-1} = \frac{T_W}{2} + T_{Contention} + T_M + T_{SwTxRx} + T_{Ack}$$

which is very similar to the WideMac mean service time.

S-MAC / T-MAC

$$\mu_{S-MAC}^{-1} = \frac{T_W}{2} + T_{rtscts} + T_M + T_{SwTxRx} + T_{Ack}$$

Again, this expression is very similar to the one found for WideMac.

3.6.2 Latency Models Validation

Results obtained with the analytical latency model for the CSMA protocol are compared in this paragraph with an implementation of the IEEE 802.15.4 non-beacon enabled mode (CSMA) in the discrete event simulator Omnet++, and with experimental measurements by Philips Research Laboratories, Eindhoven made on the Philips Acquis Grain wireless sensor networking platform. The experiment considered the simple case of one sender and one receiver (see figures 3.23 and 3.24 for respectively unicast and broadcast results). The simulation and the experiment used the parameter values given in table 3.5. Traffic generating nodes recorded the time at which each packet was given to the MAC layer for transmission (this time is also recorded in the packet itself), and the time at which the MAC informed it of the transmission success or failure. We call the difference between these two times the *sender service time*. Similarly, we call *receiver service time* the difference between the moment the sender gives the packet to the MAC and the current reception time at the receiver. In addition, the number of channel access failures, received data frames, duplicates and acknowledgments are also recorded. In the simulator, the receiver and service time values were extracted using the OMNeT++ API.

Table 3.5: IEEE 802.15.4 CSMA model parameters.

Parameter	Description	Value
minBE	Minimum Backoff Exponent.	3
maxBE	Maximum Backoff Exponent.	5
maxCSMABackoffs	Maximum number of backoffs.	4
maxFrameRetries	Maximum number of transmission trials.	3
AckWaitDuration	Acknowledgment wait interval.	864 μ s
SIFS	Minimum time between a frame and its acknowledgment.	192 μ s
aUnitBackoffPeriod	Backoff slot duration.	320 μ s
CCADetectionTime	Time required to perform a clear channel assessment.	128 μ s
Data packet size		50 bytes
Data packet duration	Time to transmit the whole data frame.	1920 μ s
ACK packet size		11 bytes
ACK packet duration	Time to transmit the whole acknowledgment frame.	352 μ s
Inter-arrival time		100 ms.
Number of packets		10000

In the case of broadcast traffic (figure 3.24), the simulated service time at the receiver and at the sender are more or less equal. This is due to the fact that the network simulator provides a single, network-wide time reference, and that processing times are not taken into account: while the source code of the simulation model is executed, the simulation time is stopped at a precise value, until all events scheduled at that simulation time have been processed. The simulation model does not attempt to capture processing times.




Figure 3.23: Validation of the CSMA analytical latency model with one sender, unicast traffic and low traffic intensity ($\rho = 0.039$). The lines show the cumulative distribution function (CDF) of the service time as measured respectively from left to right, in simulation at the sender and at the receiver (superposed lines), and on the Philips testbed at the sender and at the receiver. The analytical prediction of the mean service time is shown with a red dot, at $T=3.98$ ms. It is in line with simulation results. Experimental results deviate because of the transfer times between the microcontroller and the radio transceiver on the Philips WSN platform.

We also observe that eight values are possible for the simulated service time: approximately 2.25, 2.5, 2.9, 3.2, 3.5, 3.8, 4.2 and 4.5 ms. This is clearly due to the minBE parameter set to 3: the number of backoff slots is between 0 and $2^3 - 1 = 7$, thus 8 possible values with a slot size of 320 μ s. The minimum service time of about 2.25 ms (when the backoff is equal to zero) can be explained as follows: an initial time to perform the clear channel assessment (128 μ s), the time for the sender radio to switch from reception to transmission (192 μ s) and the frame transmission at 250 kbps (1920 μ s): $0.128 + 0.192 + 1.920 = 2.24$ ms.

The testbed measurements lead to larger service times, around 1.5 ms more for the receiver and 0.9 ms for the sender. This can be attributed to a number of factors: software interrupts handling, MAC processing time, data exchanges between the radio chip and the microcontroller (through the SPI bus), system




Figure 3.24: Validation of the CSMA analytical latency model with one sender, broadcast traffic and low traffic intensity ($\rho = 0.0336$). The lines show the cumulative distribution function (CDF) of the service time as measured respectively from left to right, in simulation at the sender and at the receiver (superposed lines), and on the Philips testbed at the sender and at the receiver. The analytical prediction of the mean service time is shown with a red dot, at $T=3.42$ ms. It is in line with simulation results. Experimental results deviate because of the transfer times between the microcontroller and the radio transceiver on the Philips WSN platform.

task scheduler, and software processing time of the measure itself. After the radio finishes receiving a frame, it informs the microcontroller through an interrupt request. This triggers the execution of an interrupt handling routine, which reads some control data from the radio concerning the frame such as its size and configures the SPI bus to send the frame to the microcontroller. Transferring a 60 bytes packet on a 2 MHz SPI bus alone takes more than 200 μ s. After the MAC processing time, the timestamping of the packet at the application takes around 50 μ s. While we cannot estimate the costs of each software routine, we attribute the remaining deviations to software issues. For instance, when the MAC sends the data frame to the application, it does so by creating a task in the system. The task handling code is executed every 320 μ s and will cause here on average an additional delay of 160 μ s. The variability of these causes also explains why the curve is smooth and why the discrete backoff values are barely noticeable. This delay can be reduced by using a system on chip platform, in which the radio directly writes the received bytes in central memory [107], and by optimizing the inter process communications (such as between the MAC and the application). The smaller service time measured at the sender is due to the fact that the sender does not have to transmit the whole data frame back from its radio to the microcontroller. The discrete

values of the backoff windows can also be seen on the measured sender service time.

We now explain how the analytical models can be used to predict these values. Equation 3.8 gives us the average delay as a function of the traffic intensity ρ and of the mean service time μ^{-1} . For the unicast case, we obtain $\mu^{-1} = 128 + 192 + 1920 + 192 + 352 = 3904 \mu\text{s}$, and $\rho = 0.039$. This leads to

$$E[\text{delay}|ucast] = 3.98 \text{ ms.}$$

This value is plotted as a red dot on figure 3.23. It is equal to the median value as measured in the simulator. It is also equal to the average value computed from the raw simulation results.

Similar results are obtained in case of broadcast traffic. In this case, the SIFS and Acknowledgment values must not be accounted in the service time, leading to $\mu^{-1} = 128 + 192 + 1920 = 3360 \mu\text{s}$, and $\rho = 0.0336$. We obtain:

$$E[\text{delay}|bcast] = 3.42 \text{ ms.}$$

This value is plotted as a red dot on figure 3.24. It is equal to the median value as measured in the simulator. It is also equal to the average value computed from the raw simulation results.

For higher values of ρ , the model was validated through simulations in a star topology. Figures 3.25 and 3.26 show the simulation results obtained for star networks with respectively five and ten sensors. The solid line shows the latency as measured in the simulator and the dashed line shows the latency computed with the queuing model. The grayed part of the figure corresponds to the saturation zone (queuing model validity limit), in which traffic offer is higher than the system's capacity: $\rho > 1$. The two lines match closely for all values of traffic intensity, excepted for ρ values close to 1. These results confirm that the analytical model can be used for a wide range of traffic intensities ρ .

3.6.3 Latency Results

Figure 3.27 shows the analytical latencies for the ideal, CSMA, WideMac, S-MAC and Crankshaft protocols. Five sensors are periodically sending packets to a sink with a Poisson arrival rate λ_s . Results are limited to congestion-free operation ($\rho < 1$), which correspond to the validity range of the models. The ideal average latency is below 5 ms and consists mainly of the time required to set up the radio into transmission mode at the sender, the packet transmission time T_M being smaller. CSMA deviates from ideality because of the backoff time and acknowledgment transmissions.

All ULP MAC protocols results (including WideMac) were obtained with a zero backoff time. This simplifying assumption, which holds in case of low traffic), leads to identical results for all of these protocols. In low traffic conditions (more than five seconds between two packets from the same sensor on average), we observe a common asymptotic limit equal to $\frac{T_M}{2}$, or half the radio duty-cycle period. For relatively higher traffic (less than 2 seconds between two packets from the same sensor on average), the service time quickly increases and the system saturates.

These results, while based on simplifying assumptions, highlight the fundamental trade-offs that must be made to reach ultra low power consumption,




Figure 3.25: Latency model validation (CSMA, 5 nodes).

whatever the strategy. To reduce the power consumption, ULP MAC protocols duty-cycle the radio. This duty-cycle leads to two performance limits. The first performance limit consists of the expected latency in low traffic conditions, which cannot be lower than half the duty-cycle period (assuming uniform packet generation, i.e. unpredictable traffic). The second performance limit consists of the maximum throughput: as a device duty cycles its transceiver, other devices have fewer opportunities to get access to it. While there are ways to increase a device availability (for instance the T-MAC scheme that increases the maximum throughput of S-MAC, and the more bit mechanism of WiseMAC), the associated latency will likely increase significantly, potentially leading to timeouts or buffer overflows. This is especially true when there are several candidate transmitters, i.e. contention situations. Besides, the effect of the backoff mechanism is difficult to model analytically. It obviously has a great effect on latency. While sensible mathematical models can be developed using Markov chains, we believe that using a network simulator allows to reduce the number of assumptions and obtain more detailed results. This approach is adopted in the following chapter.

3.7 Observations

This chapter presented WideMac, a novel ultra low power asynchronous random access receiver-initiated medium access control protocol that, unlike most protocols, does not depend on any carrier sensing mechanism. This key property makes it especially suitable for Ultra Wideband Impulse Radio technology,




Figure 3.26: Latency model validation (CSMA, 10 nodes).

which operates without carrier and whose ongoing transmissions are notoriously difficult to detect [137].

WideMac includes several mechanisms from the existing literature on wireless MAC protocols (acknowledged transmissions [122], automated retransmissions, receiver-based backoff parametrization [157]) to address well-known problems such as the hidden terminal, the exposed terminal, collisions and interference. It reaches ultra low power consumption levels by duty-cycling the radio transceiver, adopting an asynchronous random access approach [129] that offers high robustness to interference, high scalability with network traffic and network size, enables a simple implementation and offers fast neighbor discovery. This last feature is of high interest from a systems perspective, as this information can be useful as is for the application in the case of mobile network, can be exploited further by also providing high precision UWB ranging and can definitely accelerate the initialization of a routing algorithm.

Two aspects of the protocol have been identified as critical for its correct operation: the beacon collision problem and the backoff algorithm. Several solutions were proposed in section 3.4.3 to address the beacon collision problem. In particular, we showed that unavoidable clock drifts between devices, slightly randomized beacon emission times, adapted transmission procedures and strict UWB radiated power limits would allow correct protocol operation. The backoff algorithm, studied in section 3.4.4, is a combination of several existing approaches: it is parametrized by the receiver node as in RI-MAC [157], it is increased at every other transmission failure as in the MedWin IEEE 802.15.6 proposal [155], and the local backoff exponent varies using the MILD strategy as defined in MACAW [17].




Figure 3.27: Average latencies for the ideal, CSMA and WideMac protocols (analytical model, 5 sensors sending packets to a sink).

We have developed analytical models to study the power consumption and the latency of WideMac. This has enabled us to compare it with several other ULP and wireless MAC protocols. The power consumption models has led to two key observations. First, nowadays ULP MAC protocols are close to the ideal minimal low power consumption. Second, the beacon emission cost of WideMac does not penalize it when compared to the other protocols. Contrarily, its reliance on the beacon emission to synchronize for communications allows WideMac not to depend on a complex rendezvous phase based on reservation mechanisms (e.g. RTS-CTS) or on a time synchronization protocol.

The latency models lead to strikingly similar results for all ULP MAC protocols when the backoff time was neglected. It was found that the duty-cycle approach, required to save energy, implies a fundamental latency limit equal to half of the duty-cycling period, at least when the traffic is not predictable (in particular cases of more periodic traffic, nodes may eventually synchronize to reduce this latency). The radio transceiver duty-cycling was also found to pose severe practical limits on latency when the traffic increases for all ULP MAC protocols. This is due to the limited availability of the destination node, and we observed that increasing the number of candidate transceivers was likely to raise the importance of this problem. In contrast, the CSMA protocol (as defined in IEEE 802.15.4) led to results much closer to ideality.

In summary, WideMac is similar to previous ULP MAC protocols in terms of power efficiency and latency. However, it offers two advantages. First, its embedded discovery capability, and second, the absence of reliance on carrier sensing.

This qualitative and analytical study of WideMac enabled us to develop our

understanding of wireless MAC protocols and their design. While more detailed mathematical models could be developed, we believe that further evaluation would benefit more from the use of a network simulator. This is the subject of the next two chapters. Chapter 4 studies the use of WideMac in a small body area network scenario and chapter 5 considers a multihop network.

WideMac in a Star Topology Medical Body Area Network

This chapter evaluates WideMac from a systems level perspective in the context of medical body area networks, using discrete event simulations. It is shown that WideMac in combination with the IEEE 802.15.4A UWB PHY exceeds the application requirements in terms of reliability, latency and power consumption, and that it is competitive to the following technical alternatives: a narrowband ultra low power solution (WiseMAC and Texas Instruments CC 1100 868 MHz radio transceiver), the IEEE 802.15.4 non beacon enabled mode (CSMA and Texas Instruments CC 2420 2.4 GHz radio transceiver), and an FM-UWB ULP (WiseMAC) solution. As MBAN are highly susceptible to operate in high density environments, the performance degradation caused by such a configuration is also considered.

This chapter is structured as follows. Section 4.1 introduces the reader to MBAN applications and identifies key performance objectives. Section 4.2 describes the four wireless communication solutions compared in this work. Section 4.3 explains how the simulator was configured in order to obtain meaningful and comparable results, and identifies parameters of the protocols stack susceptible to influence the performance. Section 4.4 presents the simulation results in the context of no interference, and confirms the importance of some parameters. Optimal values are identified for the considered application, and performance limits are observed. Section 4.5 studies what happens to the system performance when several body area networks must coexist. Finally, section 4.6 concludes this chapter with general remarks on WideMac in the context of medical body area networks.

4.1 Medical Body Area Networks

While some areas of medicine are nowadays heavily dependent on expensive high technology analysis and visualization equipment such as magnetic resonance imaging, large parts of this activity field are still making relatively little use of today's computer and communications technology. However, this situation is evolving with the introduction of various forms of electronic medical records [2, 1].

Going one step further, the widespread adoption of miniaturized sensors to

monitor various physical parameters during extended periods of time is now seriously considered by the industry. Indeed, the IEEE Standards Association (IEEE-SA) has recently created the 802.15.6 task group with the goal of agreeing on a worldwide standard for such networks. This group is currently evaluating several radio technologies, among which we can find a classical low power narrowband solution (MedWin consortium, backed by Texas Instruments and GE Healthcare), several ultra wideband impulse radio systems (NICT, France Telecom, IMEC, Texas Instruments) and other ultra wideband systems: FM-UWB (CSEM) and chaotic UWB (NICT). Several MAC proposals complement these PHY layer technologies, most of them very similar to the IEEE 802.15.4 standard and some inspired by wireless sensor networks research in ultra low power MAC protocol.

Monitoring several health parameters continuously and simultaneously is useful for telemedicine, home care, rehabilitation and detection of life threatening environments. In addition to better monitoring of patients, it will allow the development of preventive health care by giving continuous and early feedback to potentially everybody. To be successful, these systems must be convenient to the end user: as they are permanently carried, they must be light, unobtrusive and comfortable. To reduce obtrusiveness, CSEM has developed a system using only four contact points with the body for a European Space Agency project on continuous multi-parameter health monitoring system [24]. Multiple physiological parameters are measured at each point of contact, and signal processing algorithms running on the sensors can estimate oxygen saturation (SpO₂), blood pressure (without cuff), core body temperature, respiration, electrocardiogram and activity. One of these multiparameter sensors is shown on figure 4.1, and figure 4.2 illustrates the complete system.



Figure 4.1: Combined SpO₂, core body temperature and activity sensors with ECG and respiration electrode.

To increase user comfort and to eliminate the risk of wire damage, the use of wireless data links between the sensors and a data collecting device connected to a smart phone is considered. The radio system should be small, inexpensive, ultra low power so that the sensors do not need to be recharged every day, allow coexistence with simultaneously operating networks and be robust to interferers. The smart phone would act as a gateway device between the medical body area network and an overlay data network (3G or IEEE 802.11) to allow monitoring by a human operator when needed.

The requirements of MBAN differ significantly from those of wireless sensor

networks: the network size is reduced and the network is mobile. As a consequence of the first point, multiple hop routing should be less of a problem. The second point leads to new challenges: the wireless environment can vary faster and thus the system should be robust to fading, and since humans tend to meet regularly the system should allow some degree of coexistence.



Figure 4.2: CSEM multi-parameter health monitoring system developed for the European Space Agency Long Term Medical Survey project at the Concordia Base Station in the Antarctica.

This work studies the feasibility of replacing wires from CSEM MBAN system with a wireless solution. The MBAN system operates as follows: four sensors periodically generate data to be collected at the sink. Three of them only generate activity and respiration information, while the fourth sensor generates core body temperature and SpO₂ information additionally.

Activity and respiration are encoded using 5 bytes, and the additional data generated by the fourth sensor takes up another 5 bytes. Thus three of the four sensors generate 5 bytes of application payload per measurement while the fourth one generate 10 bytes.

All sensors perform their measurement simultaneously, once per second. There are no hard latency constraints as the goal is to record the long term evolution of the user. The main evaluation criteria are the packet transmission success rate, which should be close to 1, and the power consumption at the sensors and at the sink, which should allow operation without recharging a battery at least for a week, and preferably much more (from a month to a year). The latency is also evaluated.

4.2 Compared Technologies

This work compares four wireless communication solutions. Two of them are based on a narrowband IEEE 802.15.4 compatible transceiver and the other two are based on UWB: the UWB-IR PHY as specified in IEEE 802.15.4A and an FM-UWB PHY. These four solutions were chosen because comparable

solutions are currently considered for inclusion in the upcoming IEEE 802.15.6 BAN standard. They are described below.

Several ultra low power narrowband transceivers are commercially available today. This work considers in particular the Texas Instruments CC1100 radio transceiver [160] operating at 868 MHz due to the use of this chip in the CSEM WiseNode sensor network platform and extensive practical experience. This type of solution has the advantage of being already widely available and cheap; however it is susceptible to frequency-dependent propagation effects, and the legally accessible spectrum is already used by other systems. In this work, this radio is associated to the IEEE 802.15.4 non beacon enabled mode (a CSMA protocol) under the name Nic802154_TI_CC1100. This radio is also associated to WiseMAC, the ULP MAC protocol in this evaluation under the name NicWiseMAC_TI_CC1100.

FM-UWB [58] is a low power, low data rate Ultra Wideband technology offering high robustness to jammers and fast synchronization. It is considered here in association with the WiseMAC ULP MAC protocol, under the name NicWiseMAC_FMUWB. This solution combines the advantage of clear channel assessment capability typical of narrowband radios, with the robustness to propagation effects of UWB. However, this comes at the cost of a relatively low bit rate when compared to UWB-IR.

Finally, WideMac is considered in association with the IEEE 802.15.4A UWB PHY.

Table 4.1: Radio transceiver characteristics.

Parameter	TI CC1100	UWB-IR	FM-UWB
Bit rate	250 kbps	0.85 Mbps	250 kbps
Sleep Power	60 μ W	60 μ W	60 μ W
Tx Power	51 mW	1 mW	5.5 mW
Rx Power	49.2 mW	30 mW	15 mW
SYNC Power	49.2 mW	50 mW	15 mW
Rx Setup Time	1 ms	1 ms	1 ms
Tx Setup Time	1 ms	1 ms	1 ms
Turnaround Time	22 μ s	50 μ s	50 μ s
Synchronization Time	200 μ s	72 μ s	200 μ s

Table 4.1 summarizes the key parameters used for all radio transceivers. The parameters for FM-UWB are conservative: for instance, the power consumption estimates in Rx and Tx are taken from a prototype built with discrete components. An integrated solution would consume closer to 8 mW in reception and 4 mW in transmission. The parameters of the TI CC1100 radio were measured on a system used in wireless sensor network applications (see also section 4.3). The parameters for the UWB-IR radio were taken from publications [133] and from the IEEE 802.15.4A standard when available and estimated from existing narrowband systems otherwise.

WiseMAC and WideMac wake-up intervals influence power consumption, latency and throughput. It is thus important to choose these values adequately. The following section discusses in detail the system model and the considered parameter ranges.

4.3 Methodology

This section provides information on how the four wireless solutions were compared. Subsection 4.3.1 describes the general Omnet++ simulation setup. Subsection 4.3.2 selects simulation parameters that may influence system performance, and identifies ranges of interest for each of these parameters. Finally, subsection 4.3.3 discusses the implementation of the metrics in the simulation models, and the calibration of the simulation models with realistic values.

4.3.1 Overview

The four considered solutions are compared through network simulations, using the discrete event simulator Omnet++ 4.1 and the modeling framework MiXiM 1.2. The network is defined as a compound Omnet++ module composed of four MBANHost instances, which are themselves compound modules. Figure 4.3 represents graphically the definition of a host module. It is composed of application, transport, routing and NIC (Network Interface Card) modules. The other modules that can be seen on the figure are utility modules to obtain the host address (arp), exchange data between modules using a publish-subscribe architecture (utility), model node mobility (mobility), estimate the power consumption (battery) and generate statistics (batteryStats and stats). The application layer is configured to generate 1 packet per second, with the first packet sent randomly within a time interval $[0, 1 s]$ to avoid systematic collisions.




Figure 4.3: A graphical representation of the MBANHost module.

These packets are delivered by the simulation engine to the Aggregation layer, which acts as a simple transport layer for sensor networks. It regroups application packets in a single, larger packet to minimize the energy waste caused by the per-packet fixed energy consumption overhead at the MAC layer, and to reduce the traffic load. Its main parameter, T_{Aggr} , defines the minimum time between two packets emissions by this layer. Aggregation packets are delivered to the DummyRoute networking layer, which simply adds a network ID field to the packet and forwards it to the NIC module.

The NIC module is specified as an Omnet++ module interface, so that it is possible to iterate over various NIC modules implementing the interface. In addition to simplifying the work, compared to defining as many hosts modules as there are NIC modules to evaluate, this approach also guarantees consistency: all other simulation parameters (for the application, transport and networking modules) will take the same values since they are specified from the same configuration file. A MiXiM NIC module is composed of a MAC layer and of a PHY layer. All three non UWB-IR modules use the same cumulative noise + interference model described in chapter 2 subsection 2.6.7. This model can be applied to FM-UWB radios because of the similarity of the modulation technique. The NICWideMac module uses the PHY layer presented in chapter 2 section 2.7. The IEEE 802.15.4 CSMA, WiseMAC and WideMac were all implemented in the MiXiM framework. In addition, both PHY layers and the CSMA module were released under the GPL open source license and have been integrated into MiXiM 1.2 as reference models, replacing the initial simpler modules.

Packets sent to the wireless channel are represented as AirFrames. The MiXiM ConnectionManager module computes which hosts are under the maximum interference range and delivers a copy of the AirFrame message to each such host. Propagation models are applied at reception (called AnalogueModel in MiXiM). For narrowband and FM-UWB radios, a simple pathloss attenuation depending deterministically on distance is applied (no fading). The UWB-IR model uses a Ghassemzadeh NLOS stochastic channel model (see chapter 2).

4.3.2 Parameters

The following parameters of interest were identified in the various modules defining a host:

- At the application layer, the time between two packets and the packet size. Both are defined by the application. Three of the four sensors generate 5 bytes per application packet, and one generates 10 bytes (see section 4.1).
- At the aggregation layer, the minimum time $T_{Aggregation}$ between two packets.
- At the MAC layer, the wake-up interval T_W for the ULP MAC protocols, the $minBE$ and $maxBE$ values of the backoff algorithm, the number of frame transmission attempts and the maximum number of messages awaiting transmission $maxTxQueueLength$.
- At the PHY layer, the transmit power.

As the BAN application generates packets at a relatively high rate for a wireless sensor network (1 per second), we expect significant energy savings from the Aggregation module. We consider the following value for the aggregation time $T_{Aggregation}$: 0.99 s (to deactivate it), 5 s, 10 s, 30 s and 60 s. At the MAC layer, the T_W value is expected to effect largely the latency and power consumption figures, as explained in chapter 3 sections 3.5 and 3.6. We consider the following values: 100, 250, 500, 750 and 1000 ms. The smallest value 100 ms is chosen

using eq. 3.1 from subsection 3.1, so that $P[\text{Beacon Collision}|T_W = X] < 5\%$ ($P[\text{Beacon Collision}|T_W = 100\text{ms}] = 3.7\%$), and the upper value 1000 ms is selected because larger values are not expected to lead to significant power consumption reductions. Numerous combinations are possible for the minBE and maxBE parameters. They take integer values larger than 1. A lower minimum backoff exponent minimizes the minimum latency, while a larger value helps to smooth the traffic and avoid temporary congestion when several nodes access the same destination. A large maximum backoff exponent can lead to very large latencies, which may not be problematic for the application itself but can cause buffer overflow at the source node. We considered various values of minBE , maxBE : (1, 3), (1, 5), (2, 4), (3, 5). Simulations tended to oscillate between the minBE and maxBE values, without stabilizing on intermediate values. The number of frame retransmission attempts after a missed acknowledgment is equal to 3 for all simulations. The maximum queue length is set to 10. The MAC header length is set to 32 bits for WideMac and WiseMAC, and to 72 bits for CSMA. Acknowledgment packets are 80 and 40 bits long for respectively WideMac/WiseMAC and CSMA.

The effect of clock drift is taken into account at the MAC layer, with a clock drift of 30 ppm. To simplify the study, the radiated power is fixed. The narrowband solutions transmit at 1 mW, and the UWB solutions use the maximum allowed E.I.R.P. of 37 μW (see chapter 2).

4.3.3 Metrics and Calibration




Figure 4.4: Power consumption in simulation without traffic for various values of T_W .

In addition to those parameters that can be varied (see previous subsection),

many parameters that effect significantly the performance must be fixed as realistically as possible. These concern mainly the physical layer: the radio bit rate, the setup times, transition times, power consumption values, etc... We parametrized the WiseMAC TI CC 1100 and the WiseMAC FM-UWB solutions from experimental measurements [128], and the WideMac solution from publications and educated guesses.

These parameter values influence directly the considered metrics: packet success rate, average latency, and power consumption. Packet success rate is estimated by counting each time an application layer in one of the sensor modules generates a packet and delivers it to its transport layer, and dividing this number by the number of packets actually delivered at the application layer of the sink module. Average latency is straightforward to obtain at the sink module application layer, as all Omnet++ messages are timestamped at their creation. Power consumption estimates are based on the Energy Framework for Omnet++, integrated into MiXiM[52]. This framework allows a number of devices to declare each to a Battery module, a number of states that they may enter, and define associated current draws. In our case, the physical layer declares one device (representing the radio), with four (five for the UWB radio) states: Sleep, Tx, Rx, Switch (and SYNC for the UWB radio). Depending on the modes between which the radio is switching, the Switch current draw is reconfigured. As we could unfortunately not evaluate the power consumption of an UWB-IR device, because of the unavailability of the hardware, we validated our model with measurements on the WiseMAC TI CC 1100.

Figure 4.4 illustrates the power consumption of the three ULP solutions obtained from the simulator as a function of the duty-cycle period T_W , without traffic. As expected, all power consumptions decrease with T_W , with an asymptotic value equal to the sleep mode power consumption (60 μ W). The FM-UWB power consumption is lower than the other solutions in all cases. This can be explained by considering its power consumption in transmission and reception modes, lower than for the two other radios. These results validate the protocols implementation and our usage of the energy framework.

Power consumption simulation estimates from the WiseMAC TI CC 1100 solution are represented as triangles on figure 4.5, where they can be compared to an analytical prediction of the simulation output (dashed line) and to analytical extrapolations made directly from the experimental measurements (solid line). The analytical prediction of the simulation output is obtained with the following equation :

$$P_{idle}^{WiseMAC} = \frac{I_{Sleep}T_WV + I_{SetupRx}T_{SetupRx}V + I_{Rx}T_{CCA}V}{T_W}.$$

It is expressed as a function of the sleep, reception setup and reception mode currents I_{Sleep} , $I_{SetupRx}$ and I_{Rx} and of the duty-cycle, reception setup and clear channel assessment times T_W , $T_{SetupRx}$ and T_{CCA} , as well as of the battery voltage V . Currents are defined as increments from the sleep current to simplify the expressions. The analytical extrapolation based from measurements is expressed as follows:

$$P_{idle}^{WiseMAC*} = \frac{I_{Sleep}T_WV + E_{Sampling}}{T_W},$$

where $E_{Sampling}$ is the energy required for a clear channel assessment. The parameter values can be found in table 4.2.

Table 4.2: Power consumption parameters for the WiseMAC TI CC 1100 radio transceiver analytical models.

Parameter	Value
V	3.0 V
I_{Sleep}	20 μ A
$I_{Rx} = I_{SetupRx}$	16.4 mA
$E_{Sampling}$	109.8 μ J
T_{CCA}	1 ms
$T_{SetupRx}$	1 ms

These results validate the implementation of the MAC protocol and of the radio power consumption estimates. Further, they confirm that the approach used to estimate the power consumption by considering the instantaneous current consumptions in various radio states (dashed lines and triangles) allows to closely estimate the measured power consumption (solid line).




Figure 4.5: Power consumption of the WiseMAC TI CC 1100 solution without traffic for various values of T_W , compared to two analytical estimates (one based directly on experimental measurements, and the other using the same approach of adding current consumptions in each mode as in the simulator).

4.4 Results without coexistence

This section presents extensive simulation results showing the effect of transport (AggregationTime), MAC (T_W , $minBE$, $maxBE$) and application ($nbHosts$, equal to the number of sensors + 1 to account for the sink device) parameters on performance. Subsections 4.4.1, 4.4.2 and 4.4.3 focus respectively on the packet transmission success rate, the average latency and the average power consumption at the sink and at the sensors. Subsection 4.4.4 summarizes the key observations made in this section and compares the simulation results with analytical lower bounds for the four considered solutions.

4.4.1 Packet transmission success rate

First, we observe on figure 4.6 that the case without aggregation leads to the most intense traffic load in terms of channel access requests. Whatever the technology, all solutions lead to better results when using the aggregation layer (i.e. when $aggregationTime > 1$ s), as can be seen on figures 4.7, 4.8, 4.9 and 4.10 for aggregationTime values of respectively 5, 10, 30 and 60 seconds.

Still in the case of no aggregation, we observe congestion problems for all ultra low power solutions. Reducing the wake-up interval leads to better results. Even in the worst case, with $nbSensors=4$, ULP solutions can handle the traffic load with a wake-up interval of 100 ms. We also observe a small but not negligible performance degradation between WideMac and the two WiseMAC based solutions. This is due to the lack of clear channel assessment capabilities, which constrain WideMac to use larger backoff unit slot durations. In turn, these larger backoff times lead to congestion at the entrance of the MAC layer: packets are dropped because the MAC entry queue is full (100 messages capacity, see below). Contrarily to ULP solutions, IEEE 802.15.4 CSMA leads to 100% PSR in all cases, thanks to its greater accessibility.

Simulation results with other backoff exponent values ($(minBE=1, maxBE=3)$, $(minBE=1, maxBE=5)$, $(minBE=2, maxBE=4)$) are not included here. Smaller values of $minBE$ lead to better results in the case $nbSensors=1$, confirming the explanation that the packet losses are due to capacity limits in the MAC layer rather than intrinsic design issues of WideMac. However, performance in the case of multiple access was severely affected, leading us to favor a $minBE$ value of 3. With large $maxBE$ values ($maxBE=5$) we initially observed large packet drops. This was caused by a too small buffer size at the MAC layer: $maxTxQueueLength$ was set to 10. Increasing it to 100 reduced the scope of the problem. This hypothesis on memory capacity is realistic: 100 10 byte packets need only 1 kB of memory.

When considering larger values of $T_{Aggregation}$, the packet success rate quickly increases. All ULP solutions can handle the traffic with $T_{Aggregation} = 5$ s.

Some performance anomalies can be observed, for instance on figure 4.10 for FM-UWB: in the case $nbSensors=2$, the PSR is degraded to a 50% level, despite it being close to 100% in the case $nbSensors=1$ and noticeably higher than 50% in the case $nbSensors=3$. We attribute such anomalies to a simulation artifact: as discrete event simulations are essentially Monte-Carlo simulations, whose behavior is driven by pseudo random number generators (pRNG), it can happen that a particular simulation run leads to an improbable sequence of events. Two

solutions exists. The first is to increase the simulated time duration, and the second is to use other seed values for the pRNG and average the results. This was however not done as it would have required to resimulate all configurations for a computation time 5 times longer (if considering 5 different seed values), and the time required to obtain the results presented here was already close to 10 hours without this additional complexity. Further, the error concerned FM-UWB+WiseMAC and not WideMac, which is the main subject of study. The fact that these anomalies happen with large values of $T_{Aggregation}$ (30 and 60 s) confirms this intuition: simulation runs were stopped after 20 minutes of simulated time, which means that only 40 aggregation packets were generated per sensor in the 60 s case.

We identify duty-cycle limits to allow correct operation of WideMac in the considered MBAN applications, assuming four sensors. They are regrouped in table 4.3. We observe that WiseMAC solutions seem to perform slightly better, and can use larger values of T_W . The following sections study how this influences the average latency and the power consumption.

Table 4.3: WideMac duty-cycle limits to guarantee a correct operation in the MBAN LTMS application.

For $T_{Aggregation}$ [s]	WideMac T_W [ms] period must be \leq :
0.99	100
5	500
10	750
30	1000
60	1000




Figure 4.6: Packet Success Rate without aggregation.




Figure 4.7: PSR with an aggregation packet every 5 seconds.




Figure 4.8: PSR with an aggregation packet every 10 seconds.




Figure 4.9: PSR with an aggregation packet every 30 seconds.




Figure 4.10: PSR with an aggregation packet every minute.

4.4.2 Latency

Figures 4.11, 4.12, 4.13, 4.14 and 4.15 show the average latency for $T_{Aggregation}$ values of respectively 0.99 (deactivated), 5, 10, 30 and 60 s. The figures only show the results for the case $\text{minBE}=3$, $\text{maxBE}=5$. The first two figures are of particular interest because they correspond to congestion cases for the ULP solutions, especially WideMac (see also PSR figures 4.6 and 4.7).

First, we consider the no aggregation case. With only one sensor ($\text{nbSensors}=1$), the latency is proportional to the wake-up interval T_W , excepted for WideMac. After further inspection, the very large latencies obtained in these cases are attributed to a combination of two related factors: WideMac backoff implementation differs from WiseMAC, and packets get dropped at the entrance of the MAC module because the transmit queue is full (100 packets are waiting). While both backoff algorithms use backoff exponents, and a slotted backoff mechanism, with WiseMAC (and CSMA) the backoff unit slot duration is very short: its duration is 1 ms. This is possible thanks to the CCA mechanism: the first node to get access to the channel will most likely be detected by a later node, avoiding collisions. This is not possible with UWB-IR, and hence the backoff unit slot duration is set to T_W , the duty-cycle period. Therefore, WideMac backoff algorithm leads (in the case $T_W = 750$ ms) to an average backoff of 1.15 s, which is enough to cause an accumulation of packets at the MAC layer and fill the transmit buffers. Using smaller values of minBE lead to better results in the case $\text{nbSensors}=2$ and $T_{Aggregation} = 0.99$, but did not scale with the number of sensors. Another approach, consisting of modifying the backoff unit slot duration, was evaluated. Values of 0.25, 0.5, 1 and 1.5 T_W were considered. In all cases, the packet success rate was significantly lower (by 10% or more). For values lower than T_W , this is attributed to higher collisions, and for the higher value, to under-usage of the channel, congestion and packet drops at the MAC entrance.

When considering the case $\text{nbSensors}=2$, always in the case without aggregation ($T_{Aggregation} = 0.99$ s), we observe a similar behavior for the WiseMAC-based solutions, starting with $T_W = 500$ ms. It can be seen that the more “aggressive” behavior of the WiseMAC backoff algorithm allows it to somewhat reduce the average latency compared to WideMac.

When considering $T_{Aggregation} = 5$ s, the average latency is approximately equal to $T_{Aggregation}$, except when the network becomes congested. This is the case for WideMac with $T_W = 750$ and 1000 ms. This can be compared to the relevant PSR results on figure 4.7: these two cases also lead to packet losses. Inspection of simulation result logs prove that these packets were dropped at the MAC entrance (transmission buffers full).

For larger values of $T_{Aggregation}$ (10, 30 and 60 s), no congestion is observed and the average latency is equal to $T_{Aggregation}$ in all cases.




Figure 4.11: Average latency without aggregation.




Figure 4.12: Average latency with an aggregation packet every 5 seconds.




Figure 4.13: Average latency with an aggregation packet every 10 seconds.




Figure 4.14: Average latency with an aggregation packet every 30 seconds.




Figure 4.15: Average latency with an aggregation packet every minute.

4.4.3 Power consumption

Figures 4.16, 4.17, 4.18, 4.19, 4.20 and 4.21, 4.22, 4.23, 4.24 and 4.25 show the power consumption for the sensors and for the sink, again as a function of the wake-up interval.

We observe that all ULP MAC protocols considered reach power consumption levels below 500 μ W. This is two orders of magnitude lower than the energy consumption of IEEE 802.15.4 non beacon enabled mode: 50 mW. This high power consumption level is hence not represented on the figures.

For both the sink and the sensor, the FM-UWB transceiver uses significantly less energy than the TI CC 1100 narrow-band solution. This is due to its lower power consumption, both in reception and in transmission: respectively 15 against 49.2 mW and 5.5 against 51 mW.

The WideMac UWB-IR sink power consumption is high when considering small duty-cycle intervals. This is due to the high energy cost of trying to synchronize with such a signal, combined to the relatively small but additional cost of the beacon emission.

The sensor power consumption of the UWB-IR solution follows a similar pattern, albeit with higher values and more variability. Smaller backoff exponent values sometimes lead to more regular and lower power consumption, especially for long duty-cycle values T_W . This is due to some specificities of the considered protocol implementation: when the sensor is in backoff mode, it must still send its periodic beacons. After sending such a beacon, the MAC protocol checks if there are awaiting packets. When trying to transmit (or in this case, resume the transmission) the first packet, it first checks if the locally cached destination beacon information is still valid. With longer backoffs, the probability that this information becomes invalid increases. If this happens, the protocol listens for the destination beacon. As the synchronization mode of the radio uses large amounts of energy, this is inefficient. And the longer the duty-cycle period, the higher the energy penalty. Several solutions exist: first, the beacon information validity could be extended. Second, the beacon rediscovery procedure could predict the next destination beacon emission and wake up just before. Third, smaller values of T_W can be used, so that the penalty is not too high.




Figure 4.16: Average sensor power consumption without aggregation.




Figure 4.17: Average sensor power consumption with an aggregation packet every 5 seconds.




Figure 4.18: Average sensor power consumption with an aggregation packet every 10 seconds.




Figure 4.19: Average sensor power consumption with an aggregation packet every 30 seconds.




Figure 4.20: Average sensor power consumption with an aggregation packet every minute.




Figure 4.21: Average sink power consumption without aggregation.




Figure 4.22: Average sink power consumption with an aggregation packet every 5 seconds.




Figure 4.23: Average sink power consumption with an aggregation packet every 10 seconds.




Figure 4.24: Average sink power consumption with an aggregation packet every 30 seconds.




Figure 4.25: Average sink power consumption with an aggregation packet every minute.

4.4.4 Summary

Figures 4.26 and 4.27 represent the sensor power consumption as a function of the average latency, for all considered values of T_W , as points (only results for which the packet success rate was higher than 95% are included on this graph). Blue squares, pink circles, red diamonds, and green triangles identify respectively the IEEE 802.15.4 CSMA, WideMac UWB-IR, WiseMAC CC 1100 and WiseMAC FM-UWB solutions. For each solution, an associated lower bound on performance is drawn in a dashed line of the same color. For CSMA, the line is simply equal to the reception mode power consumption. For the WiseMAC solutions, the line represents the power consumption of the Optimal Preamble Sampling protocol (periodic wake-ups + packet emissions and acknowledgment receptions). See section 3.5 for more information on analytical power consumption models. For WideMac, the Optimal Preamble Sampling model is reused, by considering in addition the cost of the periodic beacon emission. Finally, a vertical dotted line shows the lower bound on the average latency for all ultra low power of duty-cycle period T_W , equal to $T_w/2$ (when considering random packet arrival times). See section 3.6 for more information on analytical models of latency. A few points associated to ULP MAC protocols indicate average latencies lower than this lower bound. This is due to the periodicity of the packet generation process (1 second exactly between two packets), combined to uncongested networks: if the packet generated by the application is forwarded quickly by the aggregation layer (case $T_{Aggregation} = 0.99$ s) to the MAC layer, and if this MAC layer generates a short backoff time and if the packet transmission succeeds with the first transmission attempts, the packet latency will be shorter than the average. Repeating the simulations and changing the pRNG seed values, and averaging the results over these repetitions would eliminate these deviations.

The CSMA results confirm our expectations. The power consumption is always equal to the reception power consumption, and this solution clearly allows the lowest latency.

All ULP solutions quickly reach a lower bound on their power consumption when increasing the acceptable average latency. This lower bound depends on the duty-cycle period value T_W and on the transceiver characteristics. It can be seen that increasing T_W to values higher than 500 ms does not lead to significant power savings, while reducing it quickly increases the mean power consumption.

The WiseMAC results, for both the narrowband and the FM-UWB radios, confirm that the dynamic wake-up preamble of WiseMAC is a very efficient technique for CCA enabled radios: the simulation points are very close to the optimal preamble sampling analytical results. This also highlights the maturity of this solution and of the evaluated implementation, which has been developed and used in real sensor network deployments for a few years at CSEM.

The WideMac results are not as close to ideality. This can be attributed to two factors: the first is the lack of clear channel assessment capability, causing collisions and therefore energy waste and increased latency, and the second is the novelty of the protocol, which may allow further improvements and fine-tuning. Despite these potentials for improvements, the observed results are already close to the WiseMAC TI CC 1100 solution and not dramatically far from the FM-UWB solution. This confirms the validity of the UWB-IR solution

for the envisioned application.




Figure 4.26: Power consumption vs. average latency for the four solutions (for T_W values of 100, 250 and 500 ms), with simulation results and analytical lower bounds represented with respectively points and dashed lines. The vertical black line indicates the lower bound on the average latency.




Figure 4.27: Power consumption vs. average latency for the four solutions (for T_W values of 750 and 1000 ms), with simulation results and analytical lower bounds represented with respectively points and dashed lines. The vertical black line indicates the lower bound on the average latency.

4.5 Results including coexistence

This section presents simulation results showing the influence of coexistence on performance. The simulation setup introduced in section 4.3 and which was studied in detail in section 4.4 is reused. This time, the number of sensors is fixed to 4. Two, three and four networks are configured to operate simultaneously in the same simulation setup. Each network is placed in a plane, and planes are separated by a distance of 1 meter. The network of interest is located at the edge. Again, performance results are obtained for a variety of simulation parameter values. The duty-cycle period T_W takes the following values: 100, 250, 500, 750 and 1000 ms. The aggregation time $T_{Aggregation}$ takes the values 0.99, 5, 10, 30 and 60 s. The backoff exponent values are set to $\text{minBE}=3$ and $\text{maxBE}=5$, following the observations made in section 4.4.

Subsections 4.5.1, 4.5.2 and 4.5.3 focus respectively on the packet transmission success rate, the average latency and the average power consumption at the sink and at the sensors. Subsection 4.5.4 summarizes the key observations made in this section and compares the simulation results with analytical lower bounds for the four considered solutions.

4.5.1 Packet Success Rate

Figures 4.28, 4.29, 4.30, 4.31 and 4.32 show the PSR for aggregation times of 0.99, 5, 10, 30 and 60 seconds.

Considering 4.28, we observe that the IEEE 802.15.4 non beacon enabled mode solution is the only one to cope with the traffic, with the exception of the ULP solutions with $T_W = 100$ ms and $\text{nbNetworks} = 2$.

Increasing the aggregation time to 5 s reduces by a factor 5 the number of MAC frames to transmit (at the expense of the frame length, multiplied by the same factor). This improves the ULP protocol results significantly. They can cope with two networks with all values of T_W , and maintain relatively high PSR in the case $\text{nbNetworks}=3$ and 4 for the WiseMAC based solutions. The WideMac UWB-IR solution falls behind, once again penalized by its lack of clear channel assessment capability. Values of $T_{Aggregation} = 10$ s or 30 s slightly improve the results.

Finally, the case $T_{Aggregation} = 60$ s leads to a performance degradation for WiseMAC, the cause of which is unclear. WideMac maintains a PSR close to 100%, as does the IEEE 802.15.4 solution.




Figure 4.28: Packet Success Rate without aggregation.




Figure 4.29: Packet Success Rate with an aggregation packet every 5 seconds per sensor.




Figure 4.30: Packet Success Rate with an aggregation packet every 10 seconds per sensor.




Figure 4.31: Packet Success Rate with an aggregation packet every 30 seconds per sensor.




Figure 4.32: Packet Success Rate with an aggregation packet every minute per sensor.

4.5.2 Latency

Figures 4.33, 4.34, 4.35, 4.36 and 4.37 represent the average latency as a function of the aggregation time $T_{Aggregation}$, with values of respectively 0.99, 5, 10, 30 and 60 seconds.

In the case of no aggregation, congestion is observed for all technologies expected CSMA / IEEE 802.15.4 non beacon enabled mode, and the configuration $T_W = 100$ ms, NbNetworks=2. In this case we observe a service time of about $T_W/2$ for WiseMAC (the backoff term being negligible). For WideMac, with minBE=3 we would obtain an average backoff time of $\frac{2^3}{2}T_W = 4T_W = 0.4$ s. The larger observed value 1 s suggests retransmission attempts, an hypothesis confirmed by inspection of simulation logs showing numerous missing acknowledgments.

In all other cases, the high latency leads to packet rejections at the entrance of the MAC. This explains the packet losses observed on figure 4.28, and demonstrates the reliability of acknowledged transmissions. Despite difficult channel conditions, all packets considered for transmission are eventually delivered.

The results are significantly better in the case $T_{Aggregation} = 5$ s. The latency is often close to the expected value $T_{Aggr}/2$, expected for WideMac which suffers from congestion for T_W values of 750 and 1000 ms. Higher values of $T_{Aggregation}$ lead to situations without congestion.




Figure 4.33: Average latency without aggregation.




Figure 4.34: Average latency with an aggregation packet every 5 seconds per sensor.




Figure 4.35: Average latency with an aggregation packet every 10 seconds per sensor.




Figure 4.36: Average latency with an aggregation packet every 30 seconds per sensor.




Figure 4.37: Average latency with an aggregation packet every minute per sensor.

4.5.3 Power Consumption

Figures 4.38, 4.39, 4.40, 4.41 and 4.42 represent the average sensor power consumption for aggregation times $T_{Aggregation}$ of respectively 0.99, 5, 10, 30 and 60 seconds. Figures 4.43, 4.44, 4.45, 4.46 and 4.47 represent the same information for the sink power consumption.

First, we observe that generally, the power consumption of both the sink and the sensors increase with the number of coexisting networks, for all ULP solutions. This expected result is due mainly to collisions (and therefore re-transmissions) and overhearing.

Second, the power consumption decreases with longer duty-cycle periods T_W , for the three ULP solutions, independently of the coexistence situation. While this result was expected, the complexity of the considered random access protocols could have lead to unexpected side effects.

Finally, we observe that the WideMac solution leads to power consumption results similar to the WiseMAC narrowband solution. The WiseMAC FM-UWB solution is always noticeably more power efficient. This result highlights the importance of the underlying radio characteristics on the potential performance of a wireless solution.




Figure 4.38: Average sensor power consumption, without aggregation.




Figure 4.39: Average sensor power consumption with an aggregation packet every 5 seconds.




Figure 4.40: Average sensor power consumption with an aggregation packet every 10 seconds.




Figure 4.41: Average sensor power consumption with an aggregation packet every 30 seconds.




Figure 4.42: Average sensor power consumption with an aggregation packet every minute.




Figure 4.43: Average sink power consumption, without aggregation.




Figure 4.44: Average sink power consumption with an aggregation packet every 5 seconds.




Figure 4.45: Average sink power consumption with an aggregation packet every 10 seconds.




Figure 4.46: Average sink power consumption with an aggregation packet every 30 seconds.




Figure 4.47: Average sink power consumption with an aggregation packet every minute.

4.5.4 Summary

Figure 4.48 represents each considered simulation configuration (technology, T_W , $T_{Aggregation}$) as a point whose coordinates are determined by the latency and sensor power consumption results obtained with this configuration. Points with PSR lower than 75% were excluded.

The power consumption of ULP solutions barely increases with the number of coexisting networks. However, fewer configurations lead to acceptable performance results with nbNetworks=3 and 4. In this last case, very few configurations of WideMac are kept. Simulation logs show that this is mainly due to synchronization errors. Further work and calibration of the UWB-IR simulation model with future hardware implementations or MATLAB synchronization logics may improve the results. For instance, the synchronization logic simulation model requires the synchronization preamble to be free from any interference during 60 μ s, without any consideration for the power level of the interference, or for the type of interference (preamble or burst data symbol).

Results are spread between 100 μ W and 1 mW. However, all solutions present points below 500 μ W, which is clearly a satisfactory power consumption level.




Figure 4.48: Power consumption - average latency graph representing each simulated configuration as a point. Configuration for which the packet success rate was lower than 75% are excluded.

4.6 Observations

This chapter evaluated four wireless solutions for medical body area networks. The first solution combined the IEEE 802.15.4 non beacon enabled mode MAC protocol to an IEEE 802.15.4 narrowband PHY compatible transceiver (Texas Instruments CC 1100). Two other solutions combined the ultra low power WiseMAC protocol with a narrowband radio transceiver and an FM-UWB transceiver. The fourth solution combined an IEEE 802.15.4A UWB-IR PHY compatible radio with the WideMac protocol. The performance in terms of reliability (packet success rate), power consumption and latency was evaluated with network simulations implemented in the Omnet++ discrete event simulator [165] and the MiXiM simulation framework [81].

The analysis of simulation results excluding coexistence problems led to the identification of an optimal backoff algorithm parametrization, and highlighted the importance of aggregation techniques to let Ultra Low Power MAC protocols cope with high intensity traffic. It was shown that WiseMAC is a mature solution close to the optimum, and that WideMac enables IEEE 802.15.4A UWB PHY transceivers to reach comparable power consumption levels for similar latency.

Introducing simultaneously operating networks in the simulation setup significantly degraded the packet success rate, most notably in cases of no or little aggregation ($T_{aggregation} = 0.99$ or 5 s). All solutions, however, showed good robustness: power consumption was barely affected, and the reduced channel capacity lead to increased backoff times and packet rejections at the entrance of the MAC layer. This is much better than undetected packet losses as can happen with a scheduled protocol, as coexistence situations may evolve over time: therefore, a temporary network congestion would lead to packets buffering, with correct transmission at a later time.

The radio parameters had a critical effect on the power consumption. In particular, the current draw of the UWB-IR radio synchronization mode negatively impacted the WideMac solution. Conservative estimates were made, and the model could certainly benefit from calibrations with actual UWB-IR transceiver implementations.

We conclude that despite some performance differences between the three ULP solutions, all can answer the application requirements. The lack of clear channel assessment capability of the UWB-IR radio was compensated by the receiver-initiated design of the WideMac MAC protocol, without impeding too much on the system power consumption. Solutions to further reduce its power consumption and increase its transmission success rate were identified.

WideMac in a static multihop wireless sensor network

This chapter evaluates the performance of WideMac and UWB-IR in the context of a static multihop wireless sensor network, to determine the potential of this technology for this type of applications. The evaluation is performed using the Omnet++ network simulator, and considers the effect of topology, network size, network density and routing layer configuration on four metrics (connectivity, packet success rate, latency and power consumption).

Two other wireless communication systems have also been considered: IEEE 802.15.4 non beacon enabled mode on a Texas Instruments TI CC1100, and WiseMAC on the same radio. This provides useful points of comparison, as the two ultra low power solutions (WideMac and WiseMAC) can be compared to CSMA, to highlight the influence of the ULP MAC protocols. Similarly, comparing WideMac + UWB-IR with the narrowband solutions singles out the effect of the PHY layer (including channel effects). WiseRoute, the routing protocol used in this work was selected because of our prior experience with it in real short and long duration WSN deployments, and because its simplicity makes its behavior easier to understand.

This chapter is structured as follows. Section 5.1 introduces the reader to the routing problem in wireless sensor networks, and describes several well-known routing protocols. The section ends with a presentation of the protocol selected for simulations. Section 5.2 explains the methodology adopted for this evaluation, defines the simulation setup, and justifies the chosen simulation parameters. Sections 5.3 discusses the simulation results and section 5.4 concludes the chapter.

5.1 Routing in wireless sensor networks

In its most general definition, routing addresses the problem of finding the best path between two points in a network. The network can be wired or wireless, static or mobile, combine hosts of various capabilities (heterogeneous/mixed network) or composed only of identical hosts (homogeneous network), lasts for years or only a few minutes. Depending on the considered application, the definition of “best path” can vary largely. This expression can mean lowest latency for audio and video applications, highest throughput for transferring

large quantities of data, lowest energy for battery operated hosts or large data centers... Most often, the best path must combine some of these properties. Sometimes, any path will suffice.

The type of path itself varies: in its simplest form, routing connects two points of a graph. This is called a point to point, peer to peer or unicast link. Sometimes, a node wants to reach all other nodes that are part of the network, and this is called network-level broadcast. In multicast, a node wishes to reach a set of nodes. In anycast traffic, a node wants to reach one node among a set of nodes. Finally, convergecast considers the case when a set of nodes attempt to reach a common destination. This is especially typical of wireless sensor networks for periodic data collection. These communication patterns are illustrated on Figure 5.1.

Various strategies exist on how to find the best path: it can be proactive or reactive, cluster-based, fully distributed or infrastructure based, position assisted or link quality aware.




Figure 5.1: Network-level traffic patterns.

This section limits itself to the problematic of routing in the context of wireless sensor networks. It reuses the protocols classification proposed by Al-Karaki et al. [7], according to the network structure, the protocol operation, and how the source finds a route to the destination.

Al-Karaki et al. identify three possible network structures:

Flat-based: nodes are assigned equal roles and functionalities (see subsections 5.1.2 and 5.1.3);

Hierarchical-based: nodes play different roles in the network (see subsection 5.1.4);

Location-based: the positions of the nodes are known and used to route packets (see subsection 5.1.5).

They find five modes of operation:

Multipath-based: these protocols aim at increasing the network reliability by discovering and maintaining multiple paths to the destination. This increased reliability is obtained at the expense of energy consumption and signaling traffic.

Query-based: destination nodes generate queries for the data, and nodes that generate the type of requested data send it back to the destination (see subsection 5.1.5).

Negotiation-based: these protocols aim to disseminate all collected information in the network to every node, so that a user can access all of the data by interacting with any single node. They assume that nodes geographically close generate similar data, and attempt to reduce the flows of messages by exchanging metadata information between nodes so that similar data is not propagated several times. See subsection 5.1.3.

QoS-based: these protocols attempt to deliver packets within certain metrics: energy, delay, bandwidth.

Coherent-based: since sensor networks can generate vast amounts of data, it is tempting to perform in-network data processing. This can be done in a non-coherent or in a coherent way. With non-coherent data processing, data is processed locally before being forwarded to other nodes for further processing. With coherent data processing, data is forwarded almost as is to a local *aggregator* that will process coherently the data from several nodes.

Finally, they observe three strategies for a source to find a route to a destination:

Proactive: all routes are established before the generation of data packets (see subsections 5.1.4 and 5.1.9).

Reactive: routes are searched for on demand (see subsection 5.1.2).

Hybrid: these protocols combine the two previous approaches. For instance, short routes are discovered proactively, while longer routes are found reactively.

This section describes various protocols and classifies them according to the Al-Karaki taxonomy. Some considered protocols were developed for Ad Hoc networks (see subsections 5.1.1 and 5.1.2); they are included because they introduce algorithms also relevant for WSN routing. Due to the large amount

of existing protocols, this section is not exhaustive. Rather, we focus on the most well-known protocols to illustrate their basic concepts and to highlight the differences between the various routing protocol categories.

Subsection 5.1.1 presents Destination Sequenced Distance Vector (DSDV) and Optimized Link State Routing (OLSR), two routing algorithms developed in the context of mobile ad hoc networks (MANET) research that are representative of the two fundamental approaches: distance-vector based and link-state based. Subsection 5.1.2 follows with a description of two flat topology reactive algorithms: Ad hoc On demand Distance Vector Routing (AODV) and Dynamic Source Routing (DSR). Subsection 5.1.3 continues with two flat topology proactive algorithms: SPIN and TORA. Hierarchical topology protocols, which attribute different roles to the network nodes, are described in subsection 5.1.4, and subsection 5.1.5 gives a short overview of location-based protocols. Subsection 5.1.6 describe protocols specifically designed for network level broadcast traffic, a functionality often required to enable the operation of other routing protocols. Subsection 5.1.7 presents recent research on routing protocols for UWB networks and subsection 5.1.8 discusses the various types of metrics used by routing algorithms. Finally, subsection 5.1.9 describes WiseRoute, the CSEM protocol used in this evaluation.

5.1.1 Mobile Ad Hoc Networks Routing Basics

Research in Mobile Ad hoc Networks produced several routing protocols. This subsection describes Destination Sequenced Distance Vector (DSDV) and Optimized Link State Routing (OLSR), as they illustrate clearly the conceptual differences between two approaches: distance vector based and link state based. The first approach is local: nodes inform their neighbors (and only them) of local topological changes, i.e. nodes locally exchange the routes they know. The other approach is network-wide: each node maintains its own network connectivity map, from which it can independently derive optimal routes to any other node. This requires each node to communicate to the whole network, the list of its neighbors.

Therefore, the distance vector approach is much favored in wireless sensor networks, as they tend to generate less control traffic and scale better with network size.

Destination Sequenced Distance Vector (DSDV)

DSDV [68, 109] aims to build a routing table that associate to each destination, a next hop information and the number of hops to reach the destination. It is a modified version of the distributed Bellman-Ford algorithm [16] used in wired networks, so that it does not form loops, even in a rapidly changing topologies such as encountered in wireless networks.

Each routing table entry contains the destination, the neighbor to which the message must be forwarded in order to eventually reach the destination, and the distance to the destination (in hops). During the initialization, each node broadcasts a message in order to discover all its direct neighbors. The routing table is updated to include a route to each discovered neighbor. Then, each node broadcasts its routing table so that changes in the network are notified to every node. Each such broadcast includes a sequence number to

evaluate the freshness (or obsolescence) of broadcasted routing tables (and to avoid the counting-to-infinity distance vector routing problem). When a node receives a routing table for which the sequence number is greater than the sequence number of its own routing table, it updates the modified entries with the “next hop” and “hop number” fields. It then forwards to its neighbors the received routing table, after having increased the field “hop numbers” of every destination. When the same table is received several times on a given node through different routes, the node will select the routing table that came through the shortest route (in number of hops) from the table source. In order to reduce the message size, it is also possible to forward only the modifications between two versions of the routing table.

Optimized Link State Routing (OLSR)

The Optimized Link State Routing protocol (OLSR) [26] is a proactive non-uniform Link State routing approach. In the original Link State algorithm, each node propagates its link state information to all other nodes in the network. OLSR reduces the overhead so that fewer nodes re-broadcast link state information. In OLSR, every node transmits its neighbor list using periodic beacons. This enables all nodes to build a list of their two-hop neighbors. OLSR uses an extraction algorithm for Multipoint relay (MPR) selection. The multipoint relay set of a node N is the minimal (or near minimal) set of N’s one-hop neighbors such that each of N’s two-hop neighbors has at least one of N’s multipoint relays as its one-hop neighbor. In OLSR, each node selects its MPR independently and only the knowledge of its two-hop neighbors is needed. When a node broadcasts a message, all of its neighbors will receive the message. Only the MPRs, which have not seen the message before, rebroadcast the message. Therefore, the overhead for message flooding can be greatly reduced. Using OLSR, each node periodically floods the link state information of its MPR set through the network. The frequency of link state updates is adjusted according to whether a change of the MPR set has been detected. If the MPR set has been changed, the period of link state exchange is set to a minimum value. If the MPR set remains stable, the period is increased until it reaches a refresh interval value. Each node obtains network topology information and constructs its routing table through link state messages. Routes used in OLSR only include multipoint relays as intermediate nodes.

5.1.2 Flat topology reactive protocols: Ad hoc On demand Distance Vector Routing (AODV) and Dynamic Source Routing (DSR)

Developed for Ad Hoc networks, these two protocols are reactive flat network based to reduce the control messages overhead of DSDV and OLSR. Source nodes attempt to find a route by flooding a route request. The flood request is repeated by all nodes until a node that knows a route to destination (possibly the destination itself) receives the message. At this point, this node sends a route reply to the source node, reversing the path followed by the flood request to reach that node. In the case of AODV [110], the source node only stores the next hop information in its routing table, while in the case of DSR [75], the source node stores the complete route and sends it with the message. In their

initial versions, these protocols use hop count information as metric. Further work introduced various metrics to take into account energy consumption, delay and throughput. With DSR, the overhead of storing the full route in the message header can be non-negligible in large networks. This is even more the case in sensor networks, where data packet sizes are typically small. And the increased memory requirements at each node are another disadvantage of this approach.

Both techniques are limited in their scalability as they require network-wide flooding (see subsection 5.1.6 for information on network broadcast protocols). The overhead of this flooding can be bounded by introducing a maximum number of route request message repetitions, at the expense of a lower reliability. If all nodes start sending route request simultaneously, the network is at risk of collapsing.

5.1.3 Flat topology proactive protocols

The protocols described in this subsection generate some signaling traffic before the transmission of data. SPIN protocols attempt to identify sources of identical or similar data to avoid transmitting repeatedly the same information on the network. TORA aims to keep signaling messages local (and avoid network-wide broadcasts).

Negotiation-based protocols: Sensor Protocols for Information via Negotiation (SPIN)

SPIN [69, 84] is a family of adaptive protocols designed to propagate all the information collected by the network to each node. Since this can clearly overload the network, these protocols assume some redundancy of measurements among the nodes, and use a negotiation phase to exchange metadata and identify these redundancies. This mechanism enables SPIN protocols to disseminate identical or similar data packets only once.

Temporally Ordered Routing Algorithm (TORA)

This is a proactive routing algorithm based on the concept of link reversal. TORA [105] improves the partial link reversal method by detecting partitions and stopping non-productive link reversals. In TORA, the network topology is regarded as a directed graph. A Directional Acyclical Graph (DAG) is accomplished for the network by assigning each node I a height metric h_I . A link directional from I to J means $h_I > h_J$. In TORA, the height of a node is defined as a quintuple, which includes the logical time of a link failure, the unique address of the node that defines the new reference level, a reflection indicator bit, a propagation ordering parameter and an unique address of the node. The first three elements collectively represent the reference level. The last two values define an offset with respect to the reference level. Like water flowing, a packet goes from upstream to downstream according the height difference between nodes. DAG provides TORA with the capability that many nodes can send packets to a given destination and guarantees that all routes are loop-free. TORA has three basic operations: route creation, route maintenance and route erasure. A route creation operation starts with setting the height

(propagation ordering parameter in the quintuple) of the destination to 0 and heights of all other nodes to “null” (i.e., undefined). The source broadcasts a query packet containing the destination’s address. A node with a non-null height responds by broadcasting an “update” packet containing the height of its own. On receiving an update packet, a node sets its height to one more than that of the update generator. A node with higher height is considered as upstream and the node with lower height is considered as downstream. In this way, a directed acyclic graph is constructed from the source to the destination and multiple paths route may exist. The DAG in TORA may be disconnected because of node mobility. So, route maintenance operation is an important part of TORA. TORA has the unique feature that control messages are localized into a small set of nodes near the occurrence of topology changes. After a node loses its last downstream link, it generates a new reference level and broadcasts the reference to its neighbors. Therefore, links are reversed to reflect the topology change and adapt to the new reference level. The erase operation in TORA floods “clear” packets through the network and erase invalid routes.

5.1.4 Hierarchical topology protocols

This subsection describes several protocols that separate nodes in two or more categories, with different roles. Typically, a large part of the nodes fall in a data source category, and the other are selected as routers (often called cluster heads). Nodes around a cluster head or router forward all of their traffic to this node, which then takes care of delivering this data by communicating with other cluster heads and routers. These protocols are generally designed for convergecast traffic.

With LEACH, it is assumed that all cluster heads can directly communicate with the sink. TEEN extends LEACH by defining a method for cluster heads to communicate, and enables operation in a true multihop context. It also defines thresholds below which measured data is not forwarded to the sink. CBRP extends the concept of cluster based routing by defining a method for the cluster heads to find a destination node. This allows unicast traffic between any two nodes in the network. Finally, the ZigBee tree routing protocol is described because of its practical interest.

Low-Energy Adaptive Clustering Hierarchy (LEACH)

LEACH [69, 84] is one of the first hierarchical routing algorithms for sensor networks. The idea is to form clusters of the sensor nodes based on the received signal strength and to use local cluster heads (CH) as routers to the sink. To reduce the amount of information that must be transmitted to the base station, CHs are supposed to aggregate data arriving from nodes that belong to the respective cluster, sending an aggregated packet to the base station. LEACH assumes that all nodes can transmit with enough power to reach the base station if needed. Therefore, it is not applicable to networks deployed in large regions. It also assumes that nodes always have data to send, and nodes located close to each other have correlated data (for aggregation to be possible). The operation of LEACH is separated into two phases, the setup phase and the steady state phase. During the setup phase, a predetermined fraction of nodes, p , elect themselves as CHs using a probabilistic function that

results in every node to be periodically elected as a cluster head and in electing roughly the same number of CHs at each round. Each elected CH broadcast an advertisement message to the rest of the nodes in the network to inform other nodes of its presence as a CH. All the non-cluster head nodes, after receiving this advertisement, decide on the cluster to which they want to belong to. This decision is based on the signal strength of the advertisement. The non cluster head nodes inform the appropriate cluster heads that they will be a member of the cluster. After receiving all the messages from the nodes that would like to be included in the cluster and based on the number of nodes in the cluster, the cluster head node creates a TDMA schedule and assigns each node a time slot when it can transmit. This schedule is broadcast to all the nodes in the cluster. Moreover, each cluster communicates using different CDMA codes to reduce interference from nodes belonging to other clusters. During the steady state phase, the sensor nodes transmit data to the CHs, who aggregate them and forwards them directly to the base station. After a certain time, which is determined a priori, the network goes back into the setup phase again and enters another round of selecting new CHs.

Threshold sensitive Energy Efficient sensor Network protocol (TEEN)

TEEN [95] is a hierarchical routing algorithm designed to be responsive to sudden changes in the sensed data. Moving from the assumption of LEACH about the reachability of the base station, TEEN builds a hierarchical grouping where neighbor nodes form clusters and first-level cluster heads (CHs) form second-level clusters. Routing goes from each node to its first-level CH, to the second-level CH, to the base station. Each CH performs data aggregation to reduce the cost of transmission. To reduce power consumption still remaining responsive to critical changes of the sensed data, TEEN uses a reactive reporting mechanism. After the clusters are formed, the cluster head broadcasts two thresholds to the nodes: The hard threshold (HT) and the soft threshold (ST). The former identifies the range of interest for the sensed attribute, while the latter determines the minimum relevant change for the sensed attribute. Nodes transmit only when the sensed attribute is in the range of interest identified by the HT and when the change with respect to the previously value sent is greater than the ST. One can adjust both hard and soft threshold values (which are broadcast at cluster-formation time) in order to control the number of packet transmissions. At the same time, the TEEN approach is not suitable for applications that require periodic reports since the base station may not get any data at all if the thresholds are not reached.

Cluster Based Routing Protocol (CBRP)

The CBRP algorithm [93] structures the network in clusters of a diameter of two hops. The cluster head is the node that has the lowest identifier in the cluster. As a consequence, the nodes must know the network topology at a distance of two hops, in order to form the clusters. A node, which has an identifier lower than all its neighbors, decides to become a cluster head and advertises itself by sending messages to its neighbors. Before making such a decision, every node waits for the nodes that are in the neighborhood at

two hops and that have a lower identifier to make their decision. If one of its neighbors at two hops is a cluster head, then it decides to join the corresponding cluster or to become a new cluster head. In all cases, it notifies its neighbors at two hops so that they are aware of its status and that they can make a decision. A node can participate to several clusters if several cluster heads are present in its neighborhood at two hops. Every node then notifies the cluster head of the close clusters. The cluster head manages a table in which every neighboring cluster is present along the local node through which messages must go to reach it. The path to a destination is built by looking into the cluster whether the destination is present or not; if not, the request is forwarded to all cluster heads of the neighboring clusters. This process is repeated in each of the neighboring clusters by their respective cluster heads: if the destination belongs to its cluster, the cluster head forwards the request to it. Otherwise, it sends the route request to all adjacent cluster heads. All cluster heads add their addresses in the request packet in order to avoid loops. When the destination receives the route request, it sends a response with the route recorded in the packet to the source.

ZigBee Tree Routing

IEEE 802.15.4 and ZigBee are standards-based protocols that provide the network infrastructure required for wireless multi-hop network (including WSN) applications. IEEE 802.15.4 itself defines the physical and MAC layers, whereas ZigBee defines the network and application layers. ZigBee defines three types of devices: ZigBee coordinator devices, ZigBee router devices and ZigBee end devices. Every network must contain only one ZigBee coordinator, whose primary responsibility is to set up the parameters for building a network and to start that process. ZigBee routers can be used to extend the range of a network by acting as relays between devices that are too far apart to communicate directly. ZigBee end devices do not participate in routing. ZigBee specifies an algorithm (called Cskip) that provides address ranges to routers and coordinators, to be assigned to joining devices (i.e. child nodes) in a systematic manner. As result of this process, a tree structure spanning the whole network is created: the coordinator is designated as the root of the tree and the end devices become the leaves of the tree. In such a network, a node can communicate with a remote node by sending frames along the tree, called Tree Routing. The basis of tree routing is that each node can determine if it needs to forward a packet, destined to a particular node, up to its parent node or down to one of its child nodes, by simply looking at the destination address: if it belongs to a descendant, the packet is passed down to the child node leading to the destination, otherwise the packet is sent upward. Since tree routing follows the structure of a tree rather than taking the shortest path, routes may be longer than necessary (thus generating extra traffic) and are more likely to fail. To improve routing efficiency, the ZigBee algorithm also lets routers discover shortcuts by using AODV (see subsection 5.1.2).

5.1.5 Location-based protocols

Many WSN routing protocols assume that some of all nodes in the network know their position. This enables the concept of geocasting, which consists

of addressing nodes by geographical region rather than network address. The problem of obtaining the position information is often ignored, or under-estimated.

LAR attempts to propagate the nodes positions and speeds by piggybacking this information in DSR-like signaling packets. This geographical information is then used to reduce flooding, as unicast traffic is handled by a DSR-like algorithm. GEAR combines position and energy information to achieve low power routing, and GPSR presents a solution to solve the “empty region” problem, which occur when the shortest network path towards a node is not a straight geographical line. Finally, Directed Diffusion combines geocasting with convergecast traffic.

Forward Zone & Location Aided Routing (LAR)

Forward Zone [80] is a technique that can be adopted in route discovery processes to limit route request message flooding to a region encompassing the route source’s position and the route destination’s position, so saving network resources like energy and bandwidth. It consists in allowing only nodes inside the region defined by the source to become relayers. Possible optimizations of the basic scheme include re-computation of the forward zone at relayers and use of different region’s shapes (e.g. rectangle, cone). Location Aided Routing (LAR) [80] is an on-demand source routing scheme (based on Dynamic Source Routing, see subsection 5.1.2) that uses Forward Zone as a solution improving the flooding efficiency. Using LAR, every node is supposed to know its actual position and speed. Sources and destinations (and intermediate relayers) of route discovery processes put their own information in route requests and in route replies, respectively, thus every node receiving those messages learns position and speed of the messages’ originators at times when the messages are produced. Hence, using Forward Zone technique, when a node X wants to start at time t_2 a route discovery process towards a node Y , whose position and speed at time $t_1 < t_2$ are known, X defines a circular “expected zone” for Y , centered in Y , with radius equal to Y ’s speed multiplied for $(t_2 - t_1)$. Then, X defines a forward zone including its actual position and Y ’s expected zone and sends a route request message with this information.

Geographical and Energy Aware Routing (GEAR)

GEAR [180] is a routing scheme for battery-powered WSNs. GEAR is based on position and energy awareness (through the use of localization systems, like GPS, and beacon messages carrying their sender’s information) and uses such knowledge to perform low power routing. It deals with routing of queries (not data yet) from sinks to all sources inside target regions specified by the sinks. The forwarding technique is split into: 1) “greedy” routing up to the first reachable node in the target region, 2) restricted flooding or recursive greedy routing inside the target region. In the first step, every query forwarder (including the sink) selects one of its neighbors as “next hop” for the packet’s target region. To take its decision, the forwarder assigns an “estimated cost” to each of its neighbors related to the actual target region, and then chooses the neighbor with the minimum cost. Neighbor’s cost is a function (e.g. linear combination) of the known neighbor’s position and residual energy and of the energy cost to traverse the link with the neighbor. Nodes closest to the target

region's centre (i.e. greedy approach) and provided with the highest residual energy level are likely to be chosen. After selecting its next hop, the forwarder assigns itself a "learned cost" on the basis of the selected neighbor's estimated cost and communicates it backward to the previous hop. This kind of communication is repeated hop-by-hop by every node belonging to the paths traveled by queries, replacing estimated costs with more accurate learned costs. In the second step, only nodes inside the target region are involved in (re) transmission of the received query. By using restricted flooding, the packet is diffused through broadcasts. By using recursive greedy routing, every relay partitions the intended target region in four sub-regions, produces one copy of the received packet for each sub-region, and sends the four copies by using the greedy routing technique of the first step. In this case, potential relayers stop a query forwarding when they recognize that none of their neighbors are inside the query's target region.

Greedy Perimeter Stateless Routing (GPSR)

GPSR [78] is a position-based routing algorithm that uses a "greedy routing" approach whenever possible and switches to "planar graph traversal" to circumnavigate network's void regions, where the greedy approach is failing. Using greedy packet forwarding, the sender of a packet includes the approximate position of the recipient in the packet. This information is gathered by an appropriate location service. When an intermediate node receives a packet, it forwards the packet to a neighbor lying in the general direction of the recipient. Ideally, this process can be repeated until the recipient has been reached. Unfortunately, greedy routing may fail to find a path between sender and destination, even though one does exist. For example, a relay may find itself closer to the destination than any other node within transmission range. Greedy routing therefore has reached a local maximum from which it cannot recover. In this case, GPSR switches to the planar graph traversal approach. It is performed on a per-packet basis and does not require nodes to store any additional information. A packet enters the recovery mode when it arrives at a local maximum. It returns to greedy mode when it reaches a node closer to the destination than the node where the packet entered the recovery mode. Planar graphs are graphs with no intersecting edges. GPSR first constructs a connected planar sub graph of the graph formed by the nodes in the network. Then, based on this planar sub graph, GPSR forwards the packet on faces of the graph progressively closer to the destination. On each face, the packet is forwarded along the interior of the face by using the right hand rule: forward the packet on the next edge counter clockwise from the edge on which it arrived. Whenever the line between source and destination intersects the edge along which a packet is about to be forwarded, check if this intersection is closer to the destination than any other intersection previously encountered. If this is true, switch to the new face bordering on the edge the packet was about to traverse. The packet is then forwarded on the next edge it was about to be forwarded along before switching faces. The header of a packet contains additional information such as the position of the node where it entered recovery mode, the position of the last intersection that caused a face change, and the first edge traversed on the current face. Therefore, each node can make all routing decisions based only on the information about its local neighbors. This

includes the detection of an unreachable destination, when a packet traverses the first edge on the current face for the second time.

Directed Diffusion

With Direct Diffusion[72], sinks ask for data about certain phenomena from specific target regions by flooding the network with query packets called “interests”, that specify the requested attribute, threshold values and desired update interval. Every node receiving an interest caches the address of the neighbor (together with other information) the packet is coming from, so establishing a “gradient” towards the sink having originally expressed the interest. After that a query reaches the nodes belonging to the associated target region, they start sensing their environment and produce reports as required. Data flow from sources to sinks along multiple paths formed by following gradients. To reduce the number of traveled routes (and, consequently, the waste of resources), every node “reinforces” only a way upward (towards sources), by sending queries at higher rates to the selected neighbor. Reinforcement process is used also for repairing broken routes. Directed Diffusion does not explicitly deal with the energy conserving issue, except for some actions adopted to limit packet diffusion (e.g. reinforcement).

5.1.6 Network-level broadcast protocols

This subsection present protocols that address the network-level broadcast problem. Plain flooding, which requires each node to rebroadcast the packet once, is the simplest solution, but it is energy expensive and efficient. Probabilistic broadcast algorithms extend the flooding algorithm by introducing a rebroadcast probability (if this probability is 1, we obtain flooding).

Flooding

This protocol may be used to directly propagate the user data or during the route-establishment phase of another protocol. At its creation, a packet is uniquely identified by the address of the originating node and a sequence number that is incremented for each flooding initiation. The initiator locally broadcasts the packet. Each neighbor who correctly receives it records the packet initial source address and sequence number in a list of recently seen floodings. If this reference is already present in the list, the packet is dropped. Otherwise, the packet is forwarded using a MAC broadcast again. This procedure is applied by every node in the network. When used as the data forwarding scheme, this protocol is compatible with mobility. However, it is very inefficient and it does not scale with network density and at higher traffic load. As such, it is used as an upper bound for energy efficiency in the comparisons with other algorithms.

Probabilistic Broadcast

This protocol [28, 135] implements a smart flooding in which not all nodes relay the information. Compared to flooding, it exhibits the following differences: Each message is given a time-to-live (TTL) value at origin. The unique ID property is kept. All nodes keep in internal RAM received messages which

TTL has not expired yet. This applies to all messages, when the node is the final destination of the message and when it is a potential relay. Messages which TTL has expired are deleted by the nodes. At message reception, the address of the source of the local broadcast is saved in a neighbors table. A node which is not the final destination of the message forwards it with probability p . The probability of forwarding p depends on the number of entries present in the neighbors table by means of a look-up table. The protocol allows a relay to try several probabilistic forwarding attempts uniformly distributed during the remaining TTL time of the message. This number of attempts can be a protocol parameter or be automatically set for each message according to network conditions.

5.1.7 UWB Routing

UWB-IR wireless sensor networks are subject to the same problems as narrow-band wireless sensor networks. Hence, all proposed algorithms for NB WSN also apply to UWB-IR WSN. Their performance, however, may differ dramatically because of the changes in propagation, range, power consumption and high precision ranging capabilities. This section briefly presents publications on routing that explicitly target UWB.

Xu et al. [167] study several routing metrics for UWB WSN by simulation. They focus on throughput and power consumption. Their power consumption model is highly simplified: they assume that the sleep mode power consumption is equal to zero, they neglect the cost of transmitting control packets required by their MAC (RTS-CTS scheme), and assume that the power consumption in reception mode is equal to the transmit power consumption. They use a simple additive white Gaussian noise approach to multiple access interference modeling, and an idealized version of CSMA/CA for their MAC layer. Therefore, many causes of energy waste are neglected in that work. In UCAN [87], a novel routing metric is proposed, that depends on several link quality estimations (RSSI, interference, delay). It is not clear how these estimations should be made by the sensors. They also identify LAR as a relevant routing algorithm for UWB. In [35], De Nardis et al. study routing in a low rate IEEE 802.15.4A UWB network. They focus on multiple user interference, which they model with the Pulse Collision Model in Omnet++. The authors do not explain how they adapt the PCM model to burst modulation, nor how they model the synchronization preamble. They assume that frames interfere either completely or not at all. They consider a so-called cognitive routing protocol, with the aim of introducing learning capabilities in the routing algorithm. The metric, or cost function, of their routing algorithm, depends on delay and multiple user interference. This MUI information is evaluated globally, and the routing algorithm is actually run in a unique simulation module that distributes route information to each node.

Radunovic et al. [118] study the routing problem in an ad hoc wireless multihop UWB network, with the aim of maximizing throughput. They do not consider the problem of energy consumption, specific to wireless sensor networks. [145] focuses on cooperation between the MAC and routing layers to minimize the number of transmissions required from source to node. As [118], they also optimize throughput and do not consider power consumption.

5.1.8 Routing Metrics for WSN

Several metrics [173, 172, 34, 167, 87] have been proposed in the literature. We classify these broadly in five groups:

1. Hop count: this selection criterion is often used because it offers a good compromise between transmission delay and power consumption. But this criterion does not balance the transmission load over the network: some nodes will be heavily involved in the transmission (in general as relays), whereas others will not. Proactive algorithms easily build paths based on this criterion, because they know the entire network topology. In reactive algorithms, the hop number is computed during the path establishment. As a matter of fact, every node is waiting for a short period of time, so that all nodes at the same distance (in terms of hop number) can broadcast at approximately the same time. The destination node can then decide through which path the response must be sent. It has been shown in [34] that hop count may not be the most adequate metric when addressing routing for wireless sensor networks.
2. Time-based: the reactive algorithms easily compute the transmission delay, because it corresponds to the path of the request that has been received first by the destination (in the case the relaying nodes did not wait before forwarding the request). This criterion is more difficult to compute by proactive algorithms, because every node must add the mean time a message is spending at this node to the control messages it sends. This criterion can reduce bottlenecks, because nodes with a high traffic load may be less likely to receive further messages, as routes through this node will suffer from increased latency.
3. Geographic: these protocols assume that the nodes are location aware. This knowledge comes from a GPS or Galileo device or from a distributed positioning algorithm. The messages are then either addressed directly to geographic coordinates (geocasting) or there exists a mechanism to translate a node address into its geographic position. These protocols are interesting in the sense that they enable new kind of approaches for the applications by making these inherently more context-aware. More importantly, this approach can also eliminate the need for a routing table if addressing is made directly through geographical coordinates, and reduce or eliminate the need of flooding. The main problem with georouting is the backtracking procedure: when a node has a message to route, it will send it to its neighbor that is closest to the direction of the destination. But if this node cannot forward it in this direction any further, the message must be sent back and rerouted to another node. This kind of obstacle or hole avoidance procedure is critical, because assuming full network connectivity and high, uniform node density is not realistic.
4. Link quality: this groups all the metrics that estimate link quality [173], often through radio signal strength measurements or by packet error rate. The key idea here is to select links with low packet error rate. Since on wireless sensor network, nodes the most power-hungry component is the radio, the objective is to reduce its use. With radio transmissions, a message is considered successfully transmitted when the sender has received

an acknowledgment from the receiver (positive ACK). If this acknowledgment is not received, the sender will usually retry transmission several times (this is dependent of the MAC layer). With bad links, it is not uncommon to have average retry rates between 3 and 5. Compared with a high quality link where the retry rate is very close to 1, the bad link has an energy cost several times higher. Trying to use the best quality links should enable both a reduction in power consumption and shorter delays since it may reduce the total amount of frames transmissions. The difficulty of this approach is to find an estimator of the link quality, which gives exploitable results while at the same time not requiring huge amounts of resources (like traffic generation, idle listening, memory requirements, etc.). Another problem of this approach lies in the asymmetry of links: a link may be of high quality for packets sent from A to B, but packets traveling in the other direction could be lost. This makes link quality estimation more complex.

5. Energy based: this approach is based directly on energy consumption. These metrics try to identify either the least power consuming routes, or routes that go through nodes with high remaining energy levels.

5.1.9 WiseRoute

WiseRoute denotes the automatic route establishment protocol over the WiseMAC Medium Access Control, both currently part of WiseStack, the communication protocols stack used at CSEM for wireless sensor network deployments. It is a proactive, flat topology protocol based on network flooding for the routing tree initialization. The routes are established thanks to a dedicated flooding (see 5.1.6) originated at the sink. When the flooding packet is received for the first time, the address of the last node which forwarded it is recorded as the next hop, if the received signal strength of this packet is higher than a predefined threshold (named *rssiThreshold*). Because flooded packets which are already known to a node are discarded, so are the routes they carry. This property ensures that the resulting tree is free of loops.

An optimization of WiseRoute is considered in this work: because WiseRoute selects the first link strong enough, it can lead to suboptimal choices. Further, this metric only considers the quality of the last hop, neglecting the global cost of the complete route. Our improved solution operates as WiseRoute for the initial route selection, but does not neglect other route flood packets. Instead, these packets are used to evaluate if a better route towards the sink is available. This is performed as follows. An additional field *propagatedRSSI* is added to each packet, that stores the minimum (lowest quality) RSSI observed on the whole route advertised by this packet. The sink sets this field to its minimum possible value, and the one-hop neighbors of the sink (at a distance such that the received RSSI is higher than *rssiThreshold*) update this field with the RSSI value that each of them observes on the received packet. When this message is rebroadcast by each of these one-hop neighbors, the nodes that receive it update the field as follows: $propagatedRSSI = \min(propagatedRSSI, packetRSSI)$, where *packetRSSI* is the RSSI information associated to the message reception at the node. And if the node's current route *propagatedRSSI* value is lower than the *propagatedRSSI* value of the newly discovered link, then

the next hop information is changed. This maintains the loop-free property of the network, as the metric monotonically increasing, and allows the network to select better routes as they are discovered.

5.2 Methodology




Figure 5.2: UWB-IR and narrowband systems have comparable ranges (about 60 meters) when using the same EIRP transmit power (-14.3 dB or $37 \mu\text{W}$). The performance of the narrowband system degrades less quickly than for UWB.

To compare UWB and narrowband technologies fairly, the narrowband radio is configured to transmit at the same power level as the UWB-IR transceiver: $37 \mu\text{W}$ (from the UWB peak EIRP limit of -14.3 dBm / 500 MHz). Figure 5.2 illustrates the packet success rate as a function of the distance between the source and the destination, using unicast acknowledged traffic obtained through simulations (WiseMAC for the narrowband radio, WideMac for UWB-IR). The narrowband model uses a simple free space propagation model (pathloss exponent $\alpha=2.5$) combined with a shadow fading model lognormal distributed ($\mu = 0$ and $\sigma = 6$ dB, see [31, 9, 14, 138] for narrowband channel model parametrization) and the UWB-IR model uses the Ghassemzadeh LOS channel model (see chapter 2). Both the UWB and the NB solutions are shown to operate well up to 60 meters. We guess from this figure that a good routing algorithm should choose links no longer than 60 meters. However, we do not assume that the distance information is known. Therefore, the routing algorithm must evaluate a route quality from the link information provided by the radio, and from the information provided by the source node. We assume that the UWB-IR transceiver provides Signal to Noise Ratio (SNR) information to the upper layers when receiving a packet, similarly to existing NB radios (in prac-




Figure 5.3: Mean SNR values as a function of distance (200 packets at each position), for the narrowband and UWB channels using the MiXiM framework and Omnet++.

tice the radio provides Received Signal Strength Information, or RSSI, which can be converted to SNR by deducing the thermal noise value from it). Figure 5.3 shows how the mean SNR varies with distance for the two technologies, and figure 5.4 illustrates all SNR values from the simulator.

Two topologies are considered: a line and a grid. In the case of the line, the sink is located at one extremity of the line, and in the case of the grid, the sink is located close to an edge. For each topology, the number of sensors and the network density vary. For the line topology, the distance d between two nodes takes the following values: 10, 20, 30, 40 and 50 meters, and for the grid topology d takes the values 10, 25 and 50 (the maximum value is chosen in accordance with the maximum link length deduced from figures 5.2, 5.3 and 5.4). For the line topology, the following number of nodes are considered: 10, 20 and 30. And for the grid topologies: 9 (3^2), 16 (4^2) and 25 (5^2).

Three radio technologies are compared: one (WideMac) is based on UWB and two on narrowband (WiseMAC and IEEE 802.15.4 CSMA). This work initially also included UWB + ALOHA as specified in the IEEE 802.15.4A standard. However, the long simulation time of UWB-IR combined to ALOHA, made this technology too computationally intensive to evaluate in all the cases considered here. Indeed, since this protocol keeps the radio in reception mode all the time, the complex reception process is started at each time that a node within interference range begins transmitting a radio. The WiseMAC and WideMac T_W duty-cycle periods are both set to 250 ms, and the narrowband radio bit rate is set at its maximum, 250 kbps. All MAC layers attempt up to three times to transmit a frame in case of missing acknowledgment. Backoff




Figure 5.4: RSSI values as a function of distance (200 packets at each position), for the narrowband (left) and UWB (right) systems using the MiXiM framework and Omnet++.

exponents limits minBE and maxBE are set to respectively 3 and 5.

The WiseRoute snrThreshold value, that defines the minimum quality expected from a link to consider it as a route towards the sink, is varied between 1, 10 and 20 dB.

The application models a simple environmental data monitoring deployment. Each sensor generates one packet with 10 bytes of application data every minute, which is immediately transferred to the routing layer for delivery at the sink. The simulation is configured to generate 20 packets at each node (1 packet per minute per node). The application starts generating packets after 10 (simulated) minutes, to allow the network to initialize itself and to avoid generating application traffic before the end of the route initialization procedure. This initialization time is estimated by counting the time taken for each node to rebroadcast the route flooding messages.

Four metrics of interest are identified, to evaluate respectively the route selection algorithm, the proportion of application packets that successfully reach the sink, the mean latency of those packets and the mean power consumption at the sink and at the sensors. The route selection algorithm performance is measured by counting the proportion of nodes that are part of the routing tree (i.e., those which have found a route towards the sink). The application packet success rate is evaluated by counting the number of packets generated at each

node, divided by the number of packets received by the application layer at the sink. Retransmitted packets at the MAC layer are correctly identified using sequence numbers to ensure that no application packet is counted several times. Similarly, the latency is recorded by observing at the sink the packet simulated creation time and the simulation time at which the packet is received. The power consumption is estimated using the energy framework as described in section 4.3.


For each combination of topology, network size, network density, radio technology, and value of `rsiThreshold`, three simulations are run with different pseudo random generator seed values, and the results are aggregated. This reduces the risk of improbable results due to a particular run.

5.3 Results


5.3.1 Packet Error Rate and Connectivity

Figure 5.5 illustrates the connectivity obtained for all considered network configurations. In case of high density (minimum distance = 10 or 20 for the line topology, 10 for the grid topology), the connectivity is close to 1. When the density decreases (minimum distance = 30, 40, 50) and considering the line topology, we observe that a larger network size decreases the connectivity. As the low density leads to each node having only a limited number of neighbours (typically two), when one of the nodes in the line cannot find a route towards the sink, all nodes further than this one the line will not find a route either. These results are in line with the PSR represented on figure 5.2. When considering the grid topology, connectivity is much higher as nodes often have several possible routes towards the sink. For both topologies, the `snrThreshold` parameter, which is used to ignore low quality links, does not seem to be very effective. Requiring high quality links leads instead to disconnected nodes. This parameter does not seem to be relevant.

The packet success rates (PSR), represented on figure 5.6, are clearly proportional to the network connectivity. Low connectivity cases lead to low PSR. Even in the case of high connectivity, the PSR sometimes reach relatively low values. The grid topology performs much better than the line topology, because of the increased route diversity. Indeed, whereas for the line topology a single bad link can dramatically effect the network PSR, for the grid topology it is much less likely that most traffic goes through a same bad link.




(a) Line topologies.




(b) grid topologies.

Figure 5.5: Percentage of nodes that have selected a route towards the sink after the routing tree initialization.



(a) Line topologies.



(b) grid topologies.


Figure 5.6: Proportion of application packets that successfully reached the sink.

5.3.2 Latency and average number of hops


Figure 5.8 represents the average number of hops, for all packets received at the sink. With the line topology, we observe globally short routes for small networks and high densities (top left part of the figure), longer routes for lower densities and large networks (middle), and again short routes for large networks and low densities (bottom right). This is due to low connectivity: the nodes most likely to use a long route are not connected, their packets never arrive and the average number of hops is close to 1 as the packets reaching the sink come mainly from direct neighbors. With the grid topology, this effect is not noticeable (except for higher values of `snrThreshold`). Instead, routes lengths increase. Lower densities and larger networks would lead to the same effect as observed for the line topology.

Figure 5.7 represents the average latency. A network level backoff timer was introduced, taking random values between 0 and 5 s. Therefore, the average network backoff duration is 2.5 s, and the average latency should be equal to 2.5 times the average number of hops. While this holds for the narrowband solutions, the UWB WideMac results are significantly higher. This is due to a more conservative backoff algorithm, more likely to take large values.

This network level backoff was introduced to reduce congestion during the protocol initialization. It makes however less sense in the case of unicast data traffic, and latency would be greatly reduced if it were removed.




(a) Line topologies.




(b) grid topologies.

Figure 5.7: Mean latencies measured on the packets that reached the sink.



(a) Line topologies.



(b) grid topologies.

Figure 5.8: Average number of hops, based on the packets that reached the sink.

5.3.3 Power Consumption

Figures 5.9 and 5.10 illustrate respectively the mean sensor power consumption and the sink power consumption. Results from CSMA were, as expected, always equal to 50 mW (reception mode power consumption) and were excluded from the graph to improve readability.


The two ultra low power solutions lead to a consistent offset between the narrowband and UWB solutions, at the advantage of the narrowband. In some cases, the UWB sensor power consumption spiked to between 5 and 10 mW. This occurred only with `snrThreshold` values of 10 and 20 dB, and confirms the idea that this parameter is not useful, as the 1 dB value always leads to the best results.

We attribute this relatively high sensor power consumption (almost never below 500 μ W) to the energy cost of the routing tree initialization, which lasts a few minutes and generates many broadcast packets. Each simulation lasted 30 simulation minutes; therefore this initialization time is not negligible. The time required to run all the simulations prevented us to increase significantly the simulation duration. Despite this penalty, the observed values for ULP solutions remain close to 100 times lower than CSMA.


The sink power consumption is systematically slightly lower for WideMac, and below 500 μ W. This is explained by the following characteristics:

- a high UWB-IR radio bit rate (0.85 Mbps instead of 250 kbps for the narrowband radio) that shortens communications, allowing therefore the transceiver to go back to sleep faster;
- a minimal energy cost of collisions and reception at the receiver side, as the periodic beacon emission is followed by a periodic synchronization time whether or not a frame is arriving or not.

Therefore the cost of the receiver role using WideMac should be minimal.




(a) Line topologies.




(b) grid topologies.

Figure 5.9: Mean sensor power consumption.



(a) Line topologies.



(b) grid topologies.

Figure 5.10: Sink power consumption.

5.4 Observations

This chapter studied the performance of WideMac in the context of a multihop wireless sensor network. It was shown that WideMac can reliably transmit packets in this type of environment.

While the objective of this chapter was not to study routing protocols themselves, we made a few observations on this topic:

- When relying on SNR / RSSI based metrics, introducing a minimum value to discard bad routes may be counter-productive, as this can increase the average number of hops, the network load and the probability of collision and packet loss.
- A backoff timer at the network level can greatly effect the latency as the route length increases. While such a feature can be useful to minimize congestion in flooding, its usefulness for unicast traffic is questionable.
- Despite numerous publications describing sophisticated routing algorithms, a simple algorithm such as WiseRoute can satisfy the requirements of a periodic data collection application.

This evaluation of WideMac allowed to identify a range of parameters in which an UWB-IR sensor network can operate well, at ultra low power consumption levels. Significantly, this work used the exact same routing code for both the narrowband and ultra wideband radios. Further performance gains may be made by optimizing the WiseRoute implementation. For instance, the routing tree initialization that uses broadcast packets may be replaced by piggybacking routing information in the periodic WideMac beacon. The routing layer would ask WideMac to perform a neighbor discovery operation, by listening for its neighbors' beacons. This could accelerate the initialization, reduce the power consumption and the associated traffic load.

WideMac-High Availability

This chapter presents WideMac - High Availability, an Ultra low power, high throughput and low latency MAC protocol that is robust to interference and coexistence aware.

We have seen in the previous chapter that ultra low power consumption levels can be reached in wireless communication systems by using adequate MAC protocols adapted to the selected PHY layer such as WiseMAC and WideMac, with minimal latency penalty. These protocols can cope with unreliable links and intermittent interference thanks to acknowledgments and retransmission mechanisms. Their backoff algorithms allow them to scale with network density and traffic in a fair manner. While these desirable properties enable the use of these protocols in many applications, from periodic environmental information collection to safety-critical building fire detection systems and body area networks, their robustness is sensitive to the unpredictability of the wireless medium. Their maximum throughput and minimum latency are constrained by their periodic duty-cycling.

On the other hand, there exist widely deployed wireless MAC protocols that allow significantly higher bandwidth usage and low latencies, at the expense of power consumption: IEEE 802.11, Bluetooth or IEEE 802.15.4. They do not duty-cycle the radio transceiver, and allow operation over multiple frequencies to avoid interference problems (in a static manner for the 802.11 and 802.15.4, and using frequency hopping for Bluetooth). It is therefore tempting to combine the ultra low power characteristics of WideMac with one of those higher performance protocols by defining a dual mode protocol. Further, introducing a multichannel extension to WideMac can reinforce its robustness to interferers and improve its scalability with coexisting networks.

This chapter begins with a description in section 6.1 of applications that reach the limits of current ultra low power MAC protocols. Section 6.2 compares several wireless communication systems such as IEEE 802.11b / WiFi, IEEE 802.15.1 / Bluetooth and IEEE 802.15.4 / ZigBee, studies which communication performance aspects they optimize and how this was achieved. Based on these findings, section 6.3 proposes two key improvements to WideMac. The first consists of introducing a dual mode mechanism to increase its throughput and reduce its latency. The second makes it more robust to interferers and improve its coexistence capability by defining a multichannel extension of the protocol. This improved version is named WideMac-High Availability, or

WideMac-HA. Section 6.4 presents mathematical models to evaluate the protocol ideas described in section 6.3, and section 6.5 uses these models to quantify these improvements. Section 6.6 summarizes the findings of this chapter.

6.1 Applications for reliable wireless communications

The reliability of communications over wireless channels can be greatly degraded by the fluctuations of these channels over time. These variations can be caused by changes in the propagation environment (fading, multipath) and by interferers. The use of UWB technology is an answer to fading and multipath. Interferers are particularly problematic when considering industrial environments: crane monitoring, for instance, illustrates this clearly. These heavy work machines are permanently monitored so that they are operated within their safety limits: the load must be lower than a maximum threshold, and compressive and tensile stresses are continuously estimated so that in case of danger, the crane stops (see Figure 6.1). Up to now, wired connections have been used because of their higher reliability. However, wires have their own problems: connection issues, mechanical failures, deployment difficulties, cost... and therefore machine manufacturing companies are interested in novel solutions. These applications have the following requirements:

- Very low packet error rate;
- Low latency;
- Fast sensor failure detection;
- Long battery life;
- Operation in industrial environments.




Figure 6.1: Crane communication and control system (source: Hirschmann Automation and Control GmbH).

Industrial applications are not the only ones to pose new challenges. Active Noise Control (ANC) systems consist of multiple secondary acoustic sources which cancel out the acoustic field created by a primary disturbance source. The control signals for the secondary sources are computed from processing the

input signals from multiple microphone signals in real time. The efficiency and performance of such a system is basically governed by the intelligence of the underlying control algorithms. Typically, control algorithms for ANC systems are implemented as model-based adaptive feed-forward or feedback algorithms. Wireless sensor networks allow to increase the points of measurements and the number of actuators in the system. These algorithms impose real-time communication constraints ; therefore the latency must be minimal. To control stationary harmonic noise processes, it is only required to transmit 50 samples “every once in a while”, and therefore existing ULP MAC protocols are sufficient. When considering transient processes, however, sensor values must be transmitted at a 1 kHz sampling rate and with as little latency as possible. Thus, these applications also require high throughput requirements. An ANC system for building automation is illustrated on Figure 6.2.




Figure 6.2: Active Noise Control for building automation (source: Profaktor GmbH).

Medical Body Area Networks (MBAN) are another emerging trend in the world of wireless sensor networks. The possibility to continuously monitor vital health parameters over long periods has the potential to greatly improve the quality of service in health care while at the same time reducing costs:

- patients who must now spend days in observation at the hospital could go home sooner;
- the positive feedback provided by these systems to healthy persons could encourage them to adopt and maintain good lifestyle habits;
- elderly people could maintain an autonomous lifestyle longer thanks to efficient autonomous alert systems (like fall detectors) and assistance for drugs intake;

- the security of persons in hazardous environments (for instance firefighters) could be assessed in real-time, allowing faster reaction to dangerous situations.

The growing industrial interest for these systems has led to the creation of the Task Group 6 (TG6) on Body Area Networks at the IEEE 802.15 Working Group on Wireless Personal Area Networks, with the objective of developing a standard optimized for low power devices and operation on, in or around the human body to serve a variety of applications including medical, consumer electronics and personal entertainment [149]. Figure 6.3 illustrates such a system. While no strict technical criteria were defined for the selection of standardization proposals, the following key desirable properties were identified:

- protection of body through low Specific Absorption Rate (SAR);
- coexistence or operation in presence of interference (UMTS, MICS, microwave oven, Bluetooth, IEEE 802.11b & g, IEEE 802.15.3, IEEE 802.15.4a, UWB), and mitigation interference ingress (coming into the PHY) and interference egress (coming from the PHY);
- link reliability, with a residual packet error rate below 1% at a distance separation of 1 meter;
- scalability, to reconfigure the system for ultra low power consumption, high reliability, low latency;
- mobility;
- technical feasibility;
- regulatory effect.

The aspects of scalability with traffic intensity and ultra low power operation are particularly difficult to meet at the same time, and the inherent mobile nature of these systems make them sensitive to coexistence and interference.

While the example applications described in this section are not exhaustive, their diversity and their usefulness make them worth investigating. The WSN protocol limits that they reach are common:

- limited throughput,
- high latency,
- low tolerance to coexistence,
- sensitivity to interference.

The following section compares WSN with other wireless communication systems to identify the relations between protocol design choices and various performance metrics.

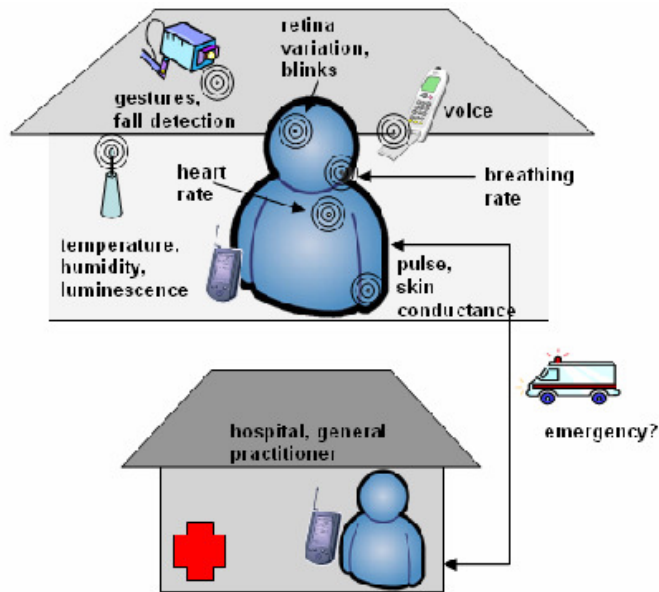


Figure 6.3: A medical body area network combined with a health status monitoring infrastructure (source: IST FP6 ESENSE).

6.2 Wireless MAC protocol design trade-offs

The performance of wireless communication systems can be evaluated through several metrics: throughput, latency, fairness, energy per bit transmission, spectral efficiency, long term power consumption, reliability... Some of these objectives are conflicting, but large application domains can be addressed by considering only a subset of these metrics.

6.2.1 Other wireless communication systems

The following list describes a variety of well-known protocols, identifies the objectives behind their design and explains how they were reached:

The IEEE 802.11 standard [154] offers high throughput, reliable data delivery and continuous network connections [102]. There is limited support for coexistence at the PHY layer by allowing operation over three to four (depending on local regulations) non-overlapping 22 MHz-wide channels (when considering the widely used IEEE 802.11b High Rate Direct Sequence Spread Spectrum (HR-DSSS) PHY layer). The spread spectrum technique offers some robustness to narrowband interference (processing gain). At the MAC layer, retransmission mechanisms, link adaptation procedures, and the introduction of the so-called virtual carrier sensing mechanism improve the robustness to intermittent interferers and address the well-known hidden node problem. The power consumption of the stations in infrastructure mode can be reduced at the expense of latency, and by relying on an always-on access point. A characteristic feature of IEEE 802.11 is its extensive support for station mobility [102], so that

it appears like a wired medium to higher layers (particularly the Logical Link Control layer). The typical raw bit rate is 11 Mbps [154], typical network size is less than 15-25 devices per access point [158] and typical device autonomy of a few hours.

The IEEE 802.15.4 standard [151] allows communications over relatively short distances. Connections should involve no or minimum infrastructure, to allow small, power-efficient and inexpensive low data rate communication solutions. The beacon enabled mode of the IEEE 802.15.4 standard [151] reduces the power consumption compared to IEEE 802.11. By introducing Guaranteed Time Slots (GTS), it allows to reach relatively high bandwidth usage. Its scalability is limited: the typical network size is 1 PAN coordinator and 7 slaves in its GTS mode. Even if the standard allows some nodes to be part of more than one network, this feature can quickly decrease the network performance because of the hidden node problems and collisions. The standard [151] supports coexistence at the PHY layer by defining several channels, spread over three frequency bands (one channel between 868 and 868.6 MHz in Europe, thirty channels between 902 and 928 MHz in North America (initially ten) and sixteen channels worldwide between 2400 and 2483.5 MHz) and using a DSSS modulation technique that offers some processing gain against interferers. Unfortunately most of the channels around 2450 MHz overlap with the three non-overlapping IEEE 802.11 channels, leaving only four channels relatively immune from this type of interference. A typical IEEE 802.15.4 network can operate on battery during weeks or months (depending on traffic type and network size). The raw bit rate varies from 10 to 250 kbps.

Bluetooth (later standardized as IEEE 802.15.1-2005 [153]) aims at low power consumption for short range communications. The physical layer uses frequency hopping over up to 79 1 MHz channels on the 2.402-2.480 GHz ISM band to avoid interference with other systems, with a raw bit rate of 3 Mbps. When interfering with an IEEE 802.11 network, a Bluetooth device will retry its transmission (Automated Repeat Request) on another hopping channel, maintaining the communication link at a price of an increased latency. As Bluetooth networks are primarily aimed at personal devices, the typical network size is small: a master device can be actively associated with up to 7 slaves, forming a piconet. The typical autonomy on battery is a few hours of active use.

Wireless Sensor Networks MAC protocols [177, 144, 45, 20, 157] minimize the power consumption at the expense of throughput and latency. They consider network sizes from tens to thousands of nodes, and aim at operation over several years on a single battery. The typical raw bit rate is 250 kbps, as they often use radio transceivers implementing the IEEE 802.15.4 PHY. As such they are sensitive to the same interference problems with IEEE 802.11. Most of these protocols do not deal well with interference and coexistence issues. Scheduled protocols are particularly sensitive to this as they usually do not implement Automated Repeat Request, and therefore their performance do not degrade gracefully. Further, they are sensitive to timing errors. The low robustness

to interference is however not always a problem: when considering low data rate applications, their traffic can be low enough that they can use other ISM bands which are forbidden to high traffic systems. Coexistence issues (with other WSN) are limited when considering static WSN, but become a problem when mobility is taken into account.

Table 6.1: Wireless communication systems optimize subsets of the possible objectives, usually at the expense of the others.

	PHY bit rate	Robustness	Coexistence	Network size	Autonomy
802.11b	2-11 Mbps	PHY (DSSS, rate adapt.), MAC (ARQ)	PHY (channel selection), MAC (ARQ)	< 25 nodes per access point [158]	hours
802.15.4	10-250 kbps	PHY (DSSS), MAC (ARQ)	PHY (channel selection)	Up to 7 devices (GTS mode)	weeks or months
Bluetooth	1-3 Mbps	PHY (channel hopping), MAC (ARQ)	PHY (channel hopping)	Up to 7 slaves per coordinator	hours or days
WSN - random access	10-250 kbps	PHY (DSSS), MAC (ARQ)	PHY (channel selection), MAC (ARQ, backoffs)	10-1000 nodes	months or years
WSN - scheduled access	10-250 kbps	PHY (DSSS)	PHY (channel selection)	10-1000 nodes	months or years

These observations are summarized in Table 6.1. This short overview of wireless communication systems highlights the trade-off between power consumption (or autonomy) and throughput and latency: 802.11b and Bluetooth provide higher throughput and low latency but offer only limited autonomy, despite 802.11 MAC protocol features that enable a device to periodically enter sleep mode and rely on the infrastructure to buffer incoming messages. The autonomy is increased in sensor networks (including 802.15.4) at the physical layer by using simpler and less power hungry radio transceivers, at the expense of throughput. At the MAC layer, further energy savings are reached by duty-cycling the radio so as to keep it in sleep mode as much as possible. This implies of course an increase of the minimum latency. Robustness to interference is addressed through several techniques: spread spectrum to offer

a processing gain against narrowband interferers, channel hopping combined to automated retransmission requests. We have found few techniques for co-existence: manual channel selection in IEEE 802.11 and IEEE 802.15.4, and reliance on MAC backoff algorithms for graceful performance degradation and fair resource allocation. We observe that the CSMA protocol is often used in these standard, either standalone (IEEE 802.15.4 non beacon enabled mode) or integrated in a more complex protocol (IEEE 802.11 and IEEE 802.15.4).

6.2.2 Multichannel MAC protocols

Many publications ([90, 98, 148, 186, 174, 55, 91, 114, 166] and references therein) have considered using several orthogonal (i.e. non interfering) channels to increase the network throughput. According to [90, 98], two main issues have been identified:

1. The *channel negotiation strategy* concerns the problem of exchanging information between two nodes, possibly listening on two different channels, so that they can exchange data afterward.
2. The *channel selection strategy* concerns how nodes are distributed on all channels.

Additionally, three problems specific to Multichannel-MAC (MMAC) have been identified:

The multichannel hidden-terminal problem: if a node is not informed of an ongoing communication on another channel, it may switch to it and cause a collision [148].

The deafness problem: a sending node attempts to reach a destination on the wrong channel.

The broadcast problem: a broadcast packet on one channel is not enough to reach all nodes.

Again according to [90, 98], the channel negotiation strategies can be divided in two categories:

1. *see-all* or *single Rendezvous*, in which case all nodes meet on the same channel to exchange control information. This strategy can be implemented by using a common dedicated control channel, a common control period (also called split-phase), or a common hopping sequence (McMAC [98]).
2. *visit-you* or *Parallel Rendezvous*, in which case the signaling traffic can be exchanged over multiple channels simultaneously. This is generally implemented with each node choosing a *private home channel*, or a *private hopping sequence*.

Single Rendezvous techniques such as Dedicated Control Channel or Common Hopping are more susceptible to congestion of control traffic, but can address the multichannel hidden terminal, deafness and broadcast problems provided that all nodes can hear each other. This assumption is often false in wireless

sensor networks, and therefore parallel rendezvous techniques are preferred in that case. Besides, Common Hopping also require some degree of time synchronization, which can be difficult to scale with the network size. And its frequent hopping can lead to performance degradation due to the repeated cost of channel switching.

The channel selection strategies can similarly be divided in two categories:

1. *global schedule*: all nodes know the channel activities of all other nodes;
2. *local schedule*: each node knows only about the activities of its neighbor nodes.

Almost all protocols are based on local scheduling because of its lower overhead, and better scalability. According to [90], the design of a local schedule is based on a selection criteria (first idle channel, least used channel, random) and the choice of the decision maker (receiver, sender, or negotiated between both). Regarding the selection criteria, choosing the first idle channel or the least used channel implies that this knowledge must be available, and this information collection has a cost. Choosing a channel at random is the simplest criteria but could lead to poor performance.

In most protocols, the channel selection decision is taken after one or more rounds of information exchange about the current channel states as seen by the sender and by the receiver. The overhead of this mechanism can be acceptable in high throughput networks, with large data packet sizes, but this is unlikely to hold for wireless sensor networks. As we have seen, most MMAC protocols focus on throughput optimization. Comparatively little results can be found for ULP or WSN MMAC protocols. We describe below all that we could find on the subject, and attempt to classify them using the previously described categories. Bluetooth can be included in the Common Hopping family with the exception that nodes do not leave the hopping sequence to exchange data packets. TMCP [174] aims at optimizing throughput and reducing latency, by spreading traffic over a small number of channels. It identifies the difficulty of maintaining coarse time synchronization in a WSN and proposes a scheme that does not require it. Channel assignment is performed globally and statically, at network deployment time. MMSN [184] also focuses on network throughput maximization. It proposes four distributed static assignment schemes (nodes do not change of channel after initial the assignment) the first two are better suited when a large number of channels is available, while the two others should lead to good performance even with a low number of channels. TMMAC [183] adopts a split-phase / common control period for the channel negotiation strategy. It requires network-wide time synchronization as it is a TDMA protocol. McMAC [98] and CAM-MAC [91] are parallel Rendezvous protocols that optimize throughput using respectively a CSMA with RTS/CTS and an asynchronous approach. Thus, even multichannel MAC protocols developed specifically for WSN applications do not address ultra low power consumption. Instead, they consider multichannel operation mainly to increase the network throughput, similarly to other MMAC protocols designed for Ad Hoc networks. The main difference between these two groups of MMAC protocols is the underlying hardware and traffic assumptions: protocols for Ad Hoc networks are most of the time based on IEEE 802.11 radio transceivers, while protocols for wireless sensor networks often assume an IEEE 802.15.4 radio transceiver such

as the Texas Instruments CC2420 [159] and smaller packet sizes (around 50 bytes).

The next section presents our ideas on how to combine the ultra low power consumption of WideMac with lower latency and higher throughput.

6.3 The WideMac-HA Protocol

This section describes WideMac - High Availability (WideMac-HA), a set of extensions to the WideMac protocol to improve its reliability, coexistence capabilities and peak throughput, and to reduce its latency.

WideMac-HA uses WideMac power saving mechanisms: asynchronous sleep, periodic beacon emissions and opportunistic local synchronizations. In addition to ultra low power consumption, it also offers two new features:

- an ALOHA-based low latency and high throughput interoperable mode to improve the performance as required and when energy levels allow it (this mode is directly inspired from the ALOHA protocol selected in IEEE 802.15.4A for the non beacon enabled mode, replacing CSMA as carrier sensing is unavailable);
- higher robustness to interference, better support for coexistence and closer to ideal protocol operating point (due to lower channel usage) by operating over multiple channels.

These two optimizations both increase the availability of the communication link, while preserving a low power consumption. Subsection 6.3.1 describes the ALOHA-based interoperable high performance mode and subsection 6.3.2 explains how WideMac can be extended to operate over multiple channels.

6.3.1 High Availability interoperable ALOHA based mode

Although WideMac can reach relatively high throughput by making use of a more bit (allowing a node to send more than one packet per wake-up interval), this solution is not satisfactory because it decreases fairness, and does not address the access bottleneck: since a node is only reachable once per wake-up interval, when there are many competing transmitters there is a high risk of collisions. An efficient backoff algorithm can minimize these collisions, but this leads to under-utilized bandwidth. Furthermore, system performance becomes unpredictable, as some nodes can be completely prevented from accessing the channel.

A high throughput and low latency high availability mode is therefore needed. Interoperability with WideMac is an important requirement to guarantee robustness:

- a node switching to this new mode should not prevent other nodes from communicating between each other using the WideMac power saving mode on the same channel;
- conversely, communications between nodes using the WideMac power saving mode should not prevent other nodes to use the high availability mode;

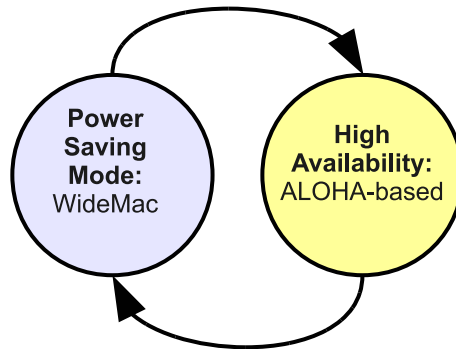


Figure 6.4: The dual-mode nature of WideMac-HA allows it to reach ultra low power consumption without sacrificing performance when and where it is needed.

- a node in low power mode should still be accessible to a node in the high availability mode;
- a node in high availability mode should still be accessible to a node in low power mode.

This is especially important in the context of multihop wireless networks as it might be impossible to make all network nodes switch from one mode to the other, due to the unreliability of wireless links. Avoiding the introduction of a difficult to maintain network-wide global state should make our solution inherently more robust. Therefore, the choice of a TDMA protocol for the High Availability mode is unsuitable as this technique attempts to guarantee collision-free time slots. Indeed, the periodic asynchronous beacon emissions of WideMac would sooner or later collide with the TDMA traffic. Instead, a random access ALOHA scheme (similar to the one defined in IEEE 802.15.4A), combined with periodic beacon emissions, is preferred. Intuitively, this protocol can be seen as a limit case of the WideMac scheme when the wake-up interval tends to zero, and by decoupling the beacon emission times and the wake-up times.

A node switching to the High Availability mode for reception can still receive packets from another node attempting to reach it using WideMac, provided that it sends periodic WideMac beacons. A flag is set in the beacon packet to announce that the node is in HA mode. This informs the source node that it can now use this mode the next time that it wants to communicate with the destination.

By simply combining two well-defined MAC protocols, we can obtain a variety of network configurations. If we consider a small BAN scenario, with a central node collecting packets from four sensor nodes, four cases are possible as illustrated on figure 6.5:

- (a) All nodes use the low power mode for maximum autonomy.
- (b) The central node, more powerful and possibly integrated into a smartphone, is in High Availability mode. The sensors can deliver their data faster, with less collision risks and with minimum latency. The sink can

reach the sensors, for instance to reconfigure their sampling or transmission rates, by using the low power mode.

- (c) All nodes are powerful enough to stay in the High Availability mode, maybe because battery recharging is not an issue, or because the application traffic needs require it. All exchanges use High Availability.
- (d) The sink and two sensors have enough energy to stay in High Availability mode, while two resource-constrained nodes are in low power mode. All communication links use High Availability excepted the down-link from the sink to the resource-constrained devices.

This description only considers static configuration of the operating mode. It should also be possible for the nodes to switch dynamically from one mode to the other. This can be done as follows: consider a network of nodes in low power mode. If one of them switches to the HA ALOHA mode, its neighbors still assume it to be in low power (LP) mode. This is not problematic as the node now in HA mode continues to send its periodic beacons and therefore remains accessible in LP mode. After such a transmission, the source node becomes aware of the destination node HA state, and can use this mode for further transmissions. If later on, the HA node switches back to LP mode, the sender node may still assume HA mode. This will lead to unsuccessful transmission attempts. But whereas ALOHA would drop the frame after a certain number of transmission attempts followed by a missing acknowledgment, here the node assumes that the destination switched back to WideMac and starts listening for the destination beacon. This guarantees that a mode switch will not lead to packet loss, since all nodes are always reachable in LP mode.

Deciding when to switch from one mode to the other is a decision that can greatly influence the system's performance, particularly if the switching frequency is too high. As this decision effects both the communication's performance and the system's autonomy, it is difficult to elaborate a generic solution to this problem. We believe that this aspect would be more appropriately handled at the system level in an energy management control plane.

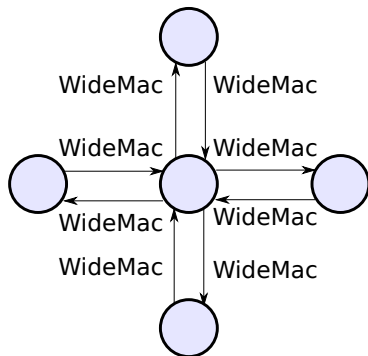
6.3.2 Multiple Channels

While operating over multiple channels leads to novel problems at the MAC layer, this approach offers significant advantages:


- by spreading traffic over multiple channels, each channel is less used. This implies less collisions, less overhearing and less interference between nodes (and thus facilitated coexistence).
- improved reliability when facing interferers. Nodes can scan for the channel state at initialization time to avoid bad channels. Furthermore, during operation, Detect and Avoid (DAA) techniques can be implemented at the MAC layer to maintain communications when a strong interferer degrades a previously usable channel.

As explained in section 6.2.2, when designing a multichannel MAC protocol we must address the following choices:


- channel negotiation strategy;



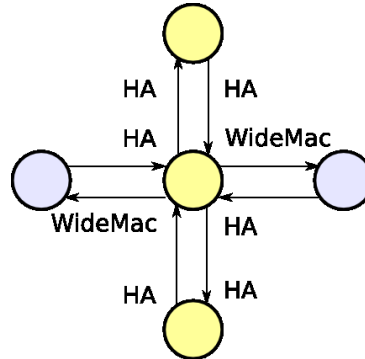
(a) All nodes save energy and use WideMac both for transmission and for reception.



(b) The sensors save energy by being accessible only in WideMac mode, and the sink is in high availability mode and listens permanently to the channel.



(c) The whole network is in high availability mode and all nodes can be reached extremely quickly using ALOHA, at the expense of power consumption.



(d) Two sensors are in high availability mode and can be accessed with the ALOHA protocol while two others, with less resources or monitoring less critical parameters, are in energy-saving mode and can be accessed only using WideMac.

Figure 6.5: Possible combinations of the WideMac-HA Low Power and High Availability modes in a BAN-like one-hop star topology network. Node colors depict their access mode (see Figure 6.4 for legend).

- channel selection strategy;
- channel selection criteria;
- channel selection decision making.

The channel negotiation strategy can be either single Rendezvous (common Dedicated Control Channel, common control period / split-phase or common hopping sequence) or parallel Rendezvous. Single Rendezvous techniques can lead to congestion on the control channel. The DCC technique is highly sensitive to interference on this channel, and a split-phase approach requires

network-wide time synchronization. We therefore favor the parallel Rendezvous technique, with each node selecting a private home channel.

The channel selection strategy, which concerns how a channel is chosen for a specific data exchange can be either global or local. We favor a local approach, as most protocols found in the literature, because of its lower overhead and better scalability. And contrarily to throughput maximization MMAC protocols, our primary aim is to increase the reliability of our solution at a minimum energy cost. Therefore we favor a very simple channel selection strategy.

The channel selection criteria can either depend on prior channel knowledge or randomized. The randomized approach seems attractive to us as it can minimize the energy consumption by avoiding the cost of channel estimation. Smarter decisions may however be taken later on, based on information captured opportunistically (i.e. overhearing, missing acknowledgments).

The channel selection decision making should similarly be minimal, and we want to avoid a complex data exchange between the sender and the transmitter. As in WideMac all nodes periodically send a beacon, an adaptation of WideMac to multichannel operation suggests an unconditional receiver-based channel selection.




Figure 6.6: WideMac Multichannel operation: node 2 must send a packet to node 1. It must first discover on which channel node 1 is sending its periodic beacon.

Based on this reflexion, we propose the following mode of operation: at initialization time, each node chooses a channel on which to perform its periodic beacon emissions. This can be done either completely at random or based on some prior channel state estimation. Nodes already operating could optionally transmit relevant information in their beacon. After the channel selection, the node would stay on the same channel, to minimize the cost of channel switching. Figure 6.6 illustrates a first data exchange between two nodes. Node 1 is sending its periodic beacon on channel 1 and node 2 does the same on channel 2. When node 2 has a message to send to node 1, it must first discover on which channel is node 1. It begins listening on channel 1 (for a maximum duration equal to the beacon period), and if it does not find node 1 there it will switch to another channel until all channels have been scanned. In the case illustrated on the figure, node 2 finds node 1 on channel 1 and delivers

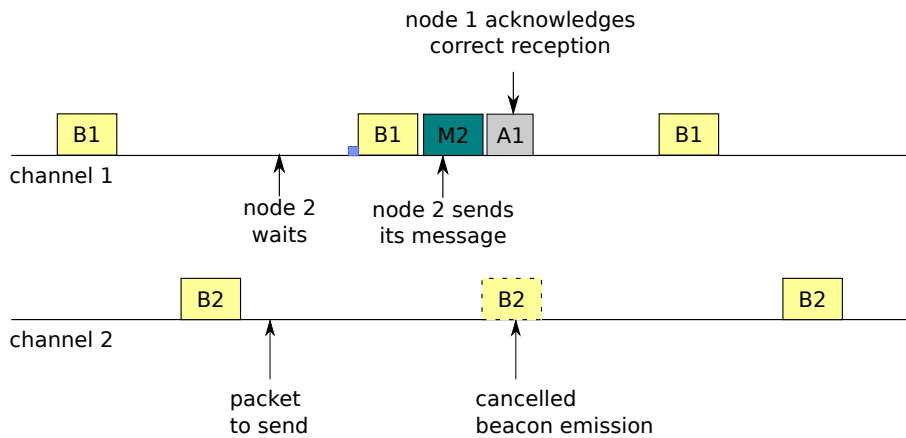


Figure 6.7: WideMac Multichannel operation: node 2 must send a packet to node 1. As it already knows on which channel node 1 is sending its beacon, and as it can predict its emission time, node 2 can save energy by keeping its radio in sleep mode.

its message. The correct reception of the message is acknowledged by node 1 on the same channel. The next time that node 2 has a message for node 1, it does not have to scan all channel. Instead, it can sleep until just before node 1 will send its beacon, after which it will deliver its message. This is illustrated on figure 6.7.

This static operation over multiple channels is enough to spread the traffic on all channels, increasing the performance and allowing better coexistence with other networks. It does not, however, provide robustness against a degraded channel (due to interference). To enable this, additional mechanisms are needed. We propose to implement a MAC-level Detect-and-Avoid (DAA) scheme, combined to a Rediscovery procedure. When a node frequently detects incoming transmissions after its beacon emissions, but always fails to receive a correct packet (idle listening caused by an interferer), or frequently receives packets addressed to other nodes (overhearing caused by high traffic), it can decide to switch to another channel. We call this procedure Detect-and-Avoid. The exact threshold at which the channel switch is decided depends on the system properties and on the characteristics of the underlying physical layer. After a node switches to another channel, any other node trying to contact it will begin by listening for its beacon message as before on the initial channel. After failing to receive the beacon, the source node will revert to the discovery procedure (as illustrated on figure 6.6) and scan all channels until it finds the destination node again. This mechanism allows to maintain communications even in presence of interferers. The channel switch and node rediscovery decisions must be taken adequately in order to maximize performance and minimize energy cost. Too frequent channel switches may lead to system's instability, by introducing race conditions.

Figure 6.8 shows WideMac operating on three channels, with node 1 and node 2 respectively using channel 1 and channel 2 for their periodic beacon emissions. After some time, node 1 detects interference on its channel and decides to switch to channel 3. Node 2, when trying to contact node 1, fails

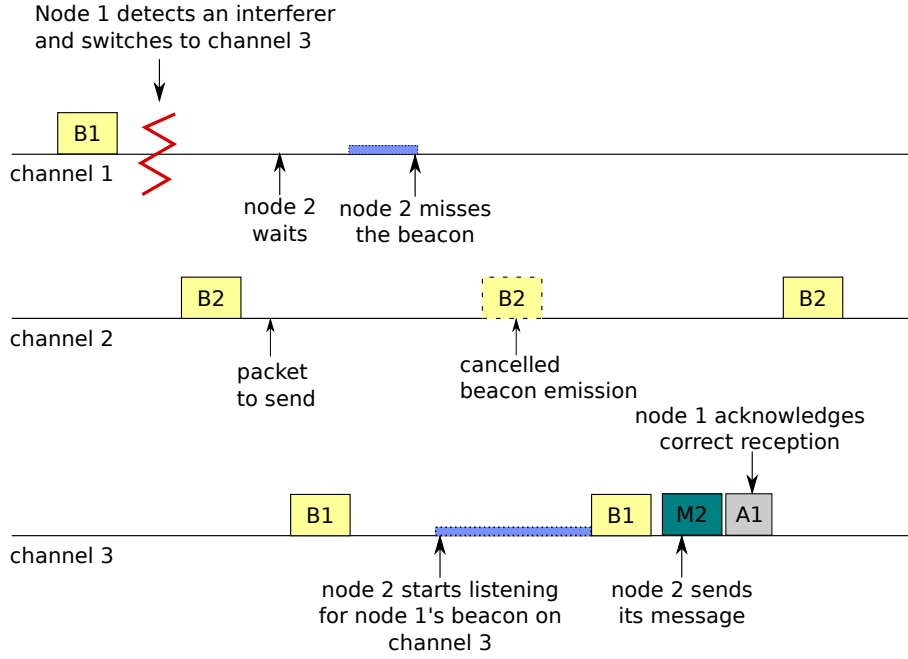


Figure 6.8: WideMac MultiChannel operation: Detect-And-Avoid by node 1 and rediscovery procedure by node 2.

to receive the beacon on channel 1, starts a rediscovery procedure and finally finds node 1 on its new channel.

6.4 Mathematical Models

This section presents analytical models of power consumption and average latency for the ideal protocol (see subsection 3.2.1), WideMac and ALOHA in the context of a medical body area network (see chapter 4), to enable the evaluation of the potential of WideMac-HA in the next section. Subsections 6.4.1 and 6.4.2 describe respectively the power consumption and average latency models.

6.4.1 Power consumption in a one-hop star topology network

First, we establish the power consumption of the ideal MAC protocol as introduced in subsection 3.2.1, for a sensor node and for the sink:

$$P_{Ideal}^{Sensor} = \frac{E_{SetupTx} + T_M P_{Tx} + (L - T_{SetupTx} - T_M) P_{Sleep}}{L}$$

$$P_{Ideal}^{Sink} = \frac{N(E_{SetupRx} + T_M P_{Rx}) + (L - NT_{SetupRx} - NT_M) P_{Sleep}}{L}.$$

The definition of the parameters not defined here can be found in table 3.4 and in section 3.5. We consider now the power consumption in ALOHA mode.

As the MAC protocol keeps the radio in reception mode, and the transmission mode power consumption is lower than the reception mode power consumption, an upper bound can be used:

$$P_{ALPHA}^{Sensor} = P_{ALPHA}^{Sink} = P_{SYNC}. \quad (6.1)$$

An analytical model of WideMac power consumption can be established through stochastic calculus [129], by following the approach presented in 3.5. It is equal to WideMac-HA power consumption in low power mode and in steady state. WideMac power consumption for the sensor node in a star topology network is given by:

$$\overline{P_{sensor}^{WideMac}} = \frac{1}{T_W} (E_{Beacon} + \overline{E_{Trans}} + \overline{E_{Recv}} + \overline{E_{Sleep}}), \quad (6.2)$$

where:

$$\begin{aligned} E_{Beacon} &= E_{SetupTx} + T_{Beacon}P_{Tx} + E_{SwTxRx} + T_{Listen}P_{Rx} \\ \overline{E_{Trans}} &= \frac{T_W}{L} [E_{SetupRx} + \overline{T_{BD}}P_{Rx} + E_{SwRxTx} \\ &\quad + T_M P_{Tx} + E_{SwTxRx} + T_{Ack}P_{Rx}] \\ \overline{E_{Recv}} &= 0 \\ \overline{E_{Sleep}} &= T_W - [T_{SetupTx} + T_{Beacon} + T_{SwTxRx} + T_{Listen} + \\ &\quad \frac{T_W}{L} (T_{SetupRx} + \overline{T_{BD}} + T_{SwRxTx} + T_M + T_{SwTxRx} + T_{Ack})]. \end{aligned}$$

Where the mean beacon detection time $\overline{T_{BD}}$ is defined as in section 3.5.2. Similarly, we obtain for the sink node using WideMac:

$$\overline{P_{Sink}^{WideMac}} = \frac{1}{T_W} (E_{Beacon} + \overline{E_{Trans}} + \overline{E_{Recv}} + \overline{E_{Sleep}}), \quad (6.3)$$

where:

$$\begin{aligned} E_{Beacon} &= E_{SetupTx} + T_{Beacon}P_{Tx} + E_{SwTxRx} + T_{Listen}P_{Rx} \\ \overline{E_{Trans}} &= 0 \\ E_{Recv} &= N \frac{T_W}{L} [T_M P_{Rx} + E_{SwTxRx} + T_{Ack}P_{Rx}] \\ \overline{E_{Sleep}} &= T_W - [T_{SetupTx} + T_{Beacon} + T_{SwTxRx} + T_{Listen} + \\ &\quad N \frac{T_W}{L} (T_M + T_{SwTxRx} + T_{Ack})]. \end{aligned}$$

6.4.2 Latency

We reuse the approach presented in section 3.6. The average delay of a $M/D/1/\infty$ system is given by:

$$E[Delay] = \frac{1}{\mu} + \frac{\rho}{2\mu(1-\rho)}, \quad (6.4)$$

where $\rho = \lambda/\mu$ is the traffic intensity, $\lambda = N\lambda_s = N/L$ is the packet rate and μ^{-1} is the service time. This model holds only for values of ρ smaller than 1.

For the ideal protocol, we have:

$$\mu_{Ideal}^{-1} = T_{SetupTx} + T_M$$

as the protocol can immediately transmit its packet and does not use acknowledgments.

For ALOHA, we obtain:

$$\mu_{ALOHA}^{-1} = \overline{T_{Backoff}} + T_M + T_{SwTxRx} + T_{Ack},$$

where $\overline{T_{Backoff}}$ is defined as in subsection 3.6.1. This assumes an idle channel.

For WideMac and at low data rates:

$$\mu_{WideMac}^{-1} = \frac{T_W}{2} + T_{SwRxTx} + T_M + T_{SwTxRx} + T_{Ack}$$

as the packet can arrive at the MAC layer at any time, and thus the node must wait on average $T_W/2$ for the destination node to wake up and send its beacon.

6.5 Evaluation

This section compares analytical power consumption and latency results for the ideal MAC, WideMac and Aloha in the context of a small medical body area network. This enables to quantify the performance improvements offered by WideMac-HA. While WideMac-HA consists of two key features, dual-mode operation and the multi-band channel access, only the first is considered here. The second one is discussed in the next section.

This section is structured as follows. Subsection 6.5.1 gives the values of all characteristic radio parameters. Subsection 6.5.2 and 6.5.3 study respectively the power consumption and the latency results.

6.5.1 Parameters

All results assume an IEEE 802.15.4A mandatory mode compliant UWB-IR radio transceiver with the following characteristics: $P_{Rx} = 30$ mW, $P_{Tx} = 1$ mW, $P_{Sync} = 45$ mW, $P_{Sleep} = 60$ μ W, bit rate = 0.85 Mbps, $T_{SetupRx} = T_{SetupTx} = 0.2$ ms, turnaround time = 20 μ s. Data packet size is set to 100 bytes, WideMac beacons to 30 bytes and acknowledgments are 4 bytes long.

6.5.2 Power Consumption

Figures 6.9 and 6.10 represent respectively the power consumption as a function of the time between two packets, of respectively the sink and the sensor, with only one sink and one sensor. The green solid line represents the lower bound of the ideal protocol, the red dashed line represents WideMac and the dotted blue line Aloha. Three values of the WideMac wake-up interval T_W are considered: 0.25, 0.5 and 1 s.

For the sink power consumption on figure 6.9, WideMac results are surprisingly close to the ideal. The spread diminishes for larger values of T_W . Two

reasons explain this small difference with ideality. First, the low beacon emission cost: $E_{Beacon} = E_{SetupTx} + T_{Beacon}P_{Tx} + E_{SwTxRx} + T_{Listen}P_{Rx} = 1.36 \mu\text{J}$, or respectively 5.44, 2.72 and 1.36 μW for the three wake-up interval values. Therefore the beacon emission cost is almost negligible. The second reason is that when a reception occurs, WideMac overhead is equal to the acknowledgment transmission (combined to the energy required to switch the radio state from reception to transmission). Since $P_{Rx} = 30 \text{ mW} = 30P_{Tx}$, the acknowledgment transmission cost is also barely noticeable.

Even with 10 packets received per second, the sink power consumption in Low Power WideMac mode is below 500 μW , or 100 times lower than Aloha's power consumption.

The results are significantly different for the sensor power consumption. As in the previous case, the wake-up interval effects the minimum power consumption (on the right side of the graphs). This time, however, the deviation from ideality is significantly higher and increases with the traffic intensity. For high traffic rates, the radio reception mode setup cost dominates (95% of the transmission cost), while for lower rates it goes down to 10% of the transmission costs, the beacon detection time quickly becoming the dominant factor (even if its duration is here never longer than 8 ms) because of the high reception mode power consumption. The slight decrease of per transmission energy cost is offset by the increased number of transmissions at high traffic rates.

But even with these significant deviations from ideality, the worst observed power consumption results of WideMac remain 50 times lower than Aloha. Figures 6.11 and 6.12 represent the same information, this time for a network of four sensors and one sink (similar to the one studied in chapter 4). The results do not change for the sensor power consumption, as the analytical expressions do not capture the effect of collisions and retransmissions that this situation would incur, especially for the higher data rates. The sink power consumption results show an increased power consumption, in the same proportion for both the ideal protocol and for WideMac.




Figure 6.9: Sink power consumption with 1 sensor.




Figure 6.10: Sensor power consumption with 1 sensor.




Figure 6.11: Sink power consumption with 4 sensors.




Figure 6.12: Sensor power consumption with 4 sensors.

6.5.3 Latency




Figure 6.13: Average latency with one sensor.

Figures 6.13 and 6.14 represent the average latency as a function of the time between two packets, respectively with one and four sensors, and for three MAC protocols: the ideal low power MAC protocol, WideMac and Aloha, plotted

respectively in solid green, dashed red and dotted blue. The ideal low power MAC latency is dominated by the radio reception mode setup time and is not affected by the number of sensors. For Aloha, the main factor in the latency is the mean backoff time $\overline{T_{Backoff}}$, and for WideMac it is the wake-up interval T_W . For all three protocols, there is a point after which the latency begins to increase with the traffic rate. This point is reached sooner with WideMac than for the other protocols, because of its periodic duty-cycling. For WideMac, this point also clearly depends on the wake-up interval, with longer wake-up intervals leading to lower access capacity.

These results illustrate the lower capacity of WideMac, and its higher latency. Switching from WideMac to Aloha can reduce latency from 500 ms to 5 ms (with $T_W = 1$ s), or two orders of magnitude (in the worst case). With a smaller T_W value of 250 ms, the latency is still reduced from 125 ms to 5 ms, a non negligible improvement.




Figure 6.14: Average latency with four sensors.

6.6 Observations

By allowing each network node to independently switch between the ultra low power mode WideMac and the low latency mode Aloha, WideMac-HA offers an ultra low power consumption and associated long operating life to body area networks without compromising on performance.

It does so without imposing any constraints on the network topology: while this work focuses on star topology networks and convergecast traffic, the protocol can offer the same level of performance for fully distributed mesh networks.

Section 6.5 quantified the performance improvements that can be obtained in a star topology MBAN network. By configuring the sink node to remain in Aloha mode (and continue sending WideMac beacons), latency is divided by a factor of at least 20. This comes at the price of an increased power consumption at the sink, by a factor 100. The exact value depends on the SYNC mode power consumption of future UWB-IR transceivers. Therefore, the application running at the sink should intelligently switch between the two MAC access modes depending on the evolving application requirements: for instance, it could switch between a normal monitoring mode of operation and a low latency alert mode. An additional advantage of the Aloha mode for sink access occurs when multiple BAN operate on the same channel. In that case, the greater availability provided by Aloha leads to much higher robustness, as can be deduced from figures 4.28, 4.29, 4.30, 4.31 and 4.32 in chapter 4 (when considering CSMA on a CC2420 transceiver).

The robustness to interference offered by the combined use of a Detect-and-Avoid mechanism and a node rediscovery procedure can greatly enhance the communications reliability. However, this has not been studied here because of its potential complexity, its dependence on application requirements, radio transceiver capabilities and electromagnetic interference from other wireless systems. The same goes for the mode switching strategy, which is highly dependent on the application.

Considering the MBAN application, one possible switching strategy is to keep all sensor devices in WideMac mode to preserve energy, and keep the sink device always in Aloha mode. An improvement of this strategy is, if the application allows it, to make the sink device switch between Aloha and WideMac depending on the analysis of the received data.

The detect-and-avoid procedure could decide to switch to another channel after five erroneous reception tentatives during the last 10 channel pollings. The exact values to use depend on expected channel usage and radio performance. The rediscovery procedure should be triggered after the maximum number of retransmissions has been reached.

The choice of efficient and robust strategies is the subject of further work and implementation in a network simulator will help refine the evaluation and specification of WideMac-HA. The adaptation of WideMac-HA to a narrowband radio transceiver is also under evaluation, under the name WiseMAC-HA. Whereas WideMac-HA combines WideMac and Aloha on an UWB-IR transceiver, WiseMAC-HA combines WiseMAC and CSMA on a narrowband radio.

Conclusion

This thesis has evaluated the potential of UWB-IR for wireless sensor networks. While UWB-IR offers a multitude of modulation types and parametrizations, we have restricted our work to IEEE 802.15.4A UWB PHY mandatory mode compatible transceivers, as this standard is specifically designed for wireless sensor networks, and is as of now the UWB-IR transceiver most likely to be one day widely available. The current lack of hardware platforms for experimentation has constrained us to rely extensively on mathematical and network simulation models. The combination of these two modeling approaches has helped us to better understand the observed results and has increased our confidence in them.

In chapter 2, the limitations of existing analytical approaches to the modeling of UWB-IR Multiple Access Interference in terms of receiver architectures, channel models, modulation type, and consideration of the synchronization preamble has motivated us to adopt a novel symbol-level approach named Maximum Pulse Amplitude Estimation. Results obtained with this technique are close to MATLAB simulation results for a sophisticated energy detection receiver. The simulation tool has enabled us to evaluate the packet error rate as a function of distance for various packet sizes and channel models, and the system robustness to multiple access interference in more configurations than what was considered for the model validation. These results have enhanced our understanding of the UWB-IR technology and its performance characteristics.

In chapter 3, the lack of clear channel assessment capabilities has been identified as a key difficulty to operate an ultra low power MAC protocol on UWB-IR. This mechanism is widely used in the literature to avoid collisions and enables nodes to detect incoming transmissions when leaving sleep mode, thereby addressing two types of energy waste, collisions and idle listening. A novel approach to ULP MAC, named WideMac, has been presented in the chapter. It is based on periodic and asynchronous beacon emissions, between which nodes independently enter sleep mode to save energy. A node with data to send simply listens to the channel until it receives the destination beacon, after which it can send its packets. As CCA is impossible, or at least difficult, the usual backoff mechanism based on the channel state at the sender side could not be chosen. Instead, a receiver parametrized backoff exponent value is used. This value is increased by a node whenever it misses a frame reception and decreased when a frame is successfully received. Mathematical models of

power consumption and latency showed WideMac to be a viable alternative to narrowband ULP solutions.

Chapter 4 has further evaluated WideMac and UWB-IR, through numerical network simulations and considering a realistic medical body area network. Systems level and application related metrics have been defined, and the technology was compared with a IEEE 802.15.4 non beacon enabled mode (CSMA) MAC protocol running on a IEEE 802.15.4 Texas Instruments CC 2420 radio transceiver and with another narrowband radio (Texas Instruments CC 1100) combined to the ultra low power WiseMAC protocol. The three communication protocol stacks were studied with numerous parameters, and performance limits and optimal operating points were identified. It was found that for the considered application, data packets should be aggregated so that no sensor generates more than 1 packet every 5 seconds, to maximize packet success rate and minimize power consumption.

Further, the same systems were reevaluated, considering this time network coexistence. It was found that the systems optimal operating points were modified, and that all considered technologies could tolerate some amount of coexistence, provided that they were adequately configured. This suggests that to reach the best performance, the system may dynamically reconfigure itself at run-time depending on its environment. This poses however stability and global state consistency problems, as all nodes must use the same parameter values.

In chapter 5, the performance of WideMac in a static multihop wireless sensor network was evaluated and compared to WideMac on a narrowband radio and to IEEE 802.15.4. All systems used the same routing protocol. Simulations have considered various network sizes and two topologies: a line and a grid. It was found that the lack of clear channel assessment of UWB-IR, which could have been particularly problematic in large networks, was adequately addressed by WideMac, as this solution reached performance results on par with a mature wireless sensor network solution (WiseMAC on TI CC 1100). Even better performance results may be obtained by optimizing the routing layer implementation so that it makes use of WideMac periodic beacon emissions. The long computing time for these simulations has limited the range of the parameter study. Further work would benefit from the implementation of the simulator optimization techniques identified in chapter 2.

Finally, chapter 6 has studied how to extend the performance of ULP MAC protocols without compromising on power consumption. WideMac - High Availability, a novel, dual mode and multichannel MAC protocol was presented. It allows each node to switch between two access modes: the ULP WideMac or the low latency Aloha, depending on current application requirements and available energy. A mechanism to extend the operation of both WideMac and Aloha over multiple channels was also described. This allows nodes to detect and avoid interference and to minimize the effect of coexistence. An analytical study has quantified the latency improvement that offers this mode switch (a factor 20 to 100) and the associated energy cost (a factor 100). It was observed that the projected performance improvements are significant and that further studies should consider the mode switch algorithm. Also, the dual mode mechanism was deemed worth evaluating for a narrowband system, combining WiseMAC and CSMA under the name WiseMAC-HA.

Overall, this thesis has confirmed the potential of UWB for wireless sen-

sensor networks as a replacement for existing narrowband technology. One of its more attracting features, however, may lie in its high precision ranging capabilities, which, if realized, would bring significant added value to the application. While the simulation model presented in chapter 2 did not explicitly consider it, UWB ranging algorithms could be modeled as well by exploiting the high resolution signal representation. Further work may consider using WideMac periodic beacon emissions to perform ranging “for free”, and the evaluation of network-level localization algorithms that improve the resolution of raw ranging measurements. These algorithms are particularly interesting in case of non line of sight situations. WideMac offers other interesting features: one of them, is the fast neighbour discovery that can be performed simply by listening for the other nodes’ periodic beacons. Such a mechanism is currently being implemented on a narrowband transceiver to enable coordinated flying of small autonomous helicopters. Further, WideMac-HA and its sibling WiseMAC-HA open novel application domains for wireless sensor networks, by dynamically reconfiguring the nodes at run-time to reach the optimal operating points and increase the reliability of wireless communications. While this poses stability issues, the potential of improvement is significant and justifies further work in this field.

Bibliography

- [1] Google Health. <http://www.google.com/health>, retrieved on Sept. 30, 2009.
- [2] NETC@RDS, a step towards the electronic EHIC. <http://netcards-project.eu>, retrieved on Sept. 30, 2009.
- [3] Standard Performance Evaluation Corporation. <http://www.spec.org>.
- [4] J. Salehi A. Forouzan, M. Nasiri-Kenari. Performance Analysis of Time-Hopping Spread-Spectrum Multiple-Access Systems: Uncoded and Coded Schemes. *IEEE Transactions on Wireless Communications*, 1(4):671–681, October 2002.
- [5] Norman Abramson. The ALOHA System—Another Alternative for Computer Communication. In *Proceedings of the AFIPS Fall Joint Computer Conference*, volume 37, pages 281–285, 1970.
- [6] Akyildiz. Wireless sensor networks: a survey. *Computer Networks*, 38(4):393, 2002.
- [7] Jamal N. Al-Karaki and Ahmed E. Kamal. Routing techniques in wireless sensor networks: a survey. *IEEE Wireless Communications*, 11(6):6, 2004.
- [8] B. Alawieh, Yongning Zhang, C. Assi, and H. Moutah. Improving Spatial Reuse in Multihop Wireless Networks - A Survey. *Communications Surveys & Tutorials, IEEE*, 11(3):71–91, Quarter 2009.
- [9] J. B. Andersen, T. S. Rappaport, and S. Yoshida. Propagation measurements and models for wireless communications channels. *Communications Magazine, IEEE*, 33(1):42–49, jan 1995.
- [10] Hicham Anouar, Younes Souilmi, and Christian Bonnet. Self-Balanced Receiver-Oriented MAC for Ultra-Wideband Mobile Ad Hoc Networks. In *Proceedings of the IEEE International Workshop on Ultra Wideband Systems*, 2003.
- [11] Nathaniel J. August, Hyung-Jin Lee, and Dong Sam Ha. Pulse Sense: a Method to Detect a Busy Medium in Pulse-Based Ultra Wideband (UWB) Networks. In *Proceedings of the IEEE Conference on Ultra Wideband Systems and Technologies*, pages 366–370, 2004.

-
- [12] Nathaniel J. August, Hyung-Jin Lee, and Dong Sam Ha. Enabling Distributed Medium Access Control for Impulse-Based Ultrawideband Radios. *IEEE Transactions on Vehicular Technology*, 56(3):1064–1075, 2007.
- [13] Elyes Ben Hamida, Guillaume Chelius, and Jean Marie Gorce. Impact of the Physical Layer Modeling on the Accuracy and Scalability of Wireless Network Simulation. *SCS SIMULATION 2009: Transactions of the Society for Modeling and Simulation International*, 85(9):574–588, 2009.
- [14] R. Bernhardt. Macroscopic Diversity in Frequency Reuse Radio Systems. *Selected Areas in Communications, IEEE Journal on*, 5(5):862–870, jun 1987.
- [15] Dimitri Bersekas and Robert Gallager. *Data networks*. Prentice-Hall, 2nd edition, 1992.
- [16] D. Bertsekas and R. Gallager. *Data Networks*. Prentice Hall, 1992.
- [17] Vaduvur Bharghavan, Alan Demers, Scott Shenker, and Lixia Zhang. MACAW: a media access protocol for wireless LAN's. In *SIGCOMM '94: Proceedings of the conference on Communications architectures, protocols and applications*, pages 212–225, New York, NY, USA, 1994. ACM.
- [18] Ioannis Broustis, Mart Molle, Srikanth Krishnamurthy, Michalis Faloutsos, and Jeffrey Foerster. A new binary conflict resolution-based MAC protocol for impulse-based UWB ad hoc networks. *Wireless Communications and Mobile Computing*, 6(7):933, 2006.
- [19] Ioannis Broustis, Angelos Vlavianos, Prashant Krishnamurthy, and Srikanth Krishnamurthy. CTU: Capturing Throughput Dependencies in UWB Networks. In *Proceedings of the IEEE Conference on Computer Communications (INFOCOM 2008)*, pages 412–420, 2008.
- [20] Michael Buettner, Gary Yee, Eric Anderson, and Richard Han. X-MAC: a short preamble MAC protocol for duty-cycled wireless sensor networks. In *Proceedings of the 4th international conference on Embedded networked sensor systems*, pages 307–320, 2006.
- [21] D. Cassioli, M. Z. Win, F. Vatalaro, and A. F. Molisch. Low Complexity Rake Receivers in Ultra-Wideband Channels. *IEEE Transactions on Wireless Communications*, 6(4):1265–1275, April 2007.
- [22] David Cavin, Yoav Sasson, and André Schiper. On the accuracy of MANET simulators. In *POMC '02: Proceedings of the second ACM international workshop on Principles of mobile computing*, pages 38–43, New York, NY, USA, 2002. ACM.
- [23] S. Chatterjea, L.F.W. van Hoesel, and P.J.M. Havinga. AI-LMAC: An Adaptive, Information-centric and Lightweight MAC Protocol for Wireless Sensor Networks. In *Proceedings of the 2nd International Conference on Intelligent Sensors, Sensor Networks and Information Processing*, pages 381–388, 2004.

- [24] O. Chetelat, J. Sola i Caros, J. Krauss, S. Dasen, S. Droz, R. Gentsch, J.-M. Koller, J. Luprano, A. O'Hare, P. Pilloud, and P. Theurillat. Continuous multi-parameter health monitoring system. In Springer, editor, *World Congress on Medical Physics and Biomedical Engineering 2006*, volume 9, pages 684–687, Sept. 2007.
- [25] Chia-Chin Chong and Su Khiong Yong. UWB Direct Chaotic Communication Technology for Low-Rate WPAN Applications. *IEEE Transactions on Vehicular Technology*, 57(3):1527, 2008.
- [26] Lars Christensen and Gitte Hansen. The Optimized Link State Routing Protocol - Performance Analysis through Scenario-based Simulations. Master's thesis, Aalborg University, Sweden, September 2001.
- [27] Yuechun Chu and Aura Ganz. MAC Protocols for Multimedia Support in UWB-based Wireless Networks. In *Proc. of BROADNETS 2004*, 2004.
- [28] Sengul Cigdem, Indranil Gupta, and Matthew J. Miller. Adaptive probability-based broadcast forwarding in energy-saving sensor networks. *ACM Transactions on Sensor Networks (TOSN)*, 4(2):1–32, 2008.
- [29] Ugo Maria Colesanti, Carlo Crociani, and Andrea Vitaletti. On the accuracy of omnet++ in the wireless sensor networks domain: simulation vs. testbed. In *PE-WASUN '07: Proceedings of the 4th ACM workshop on Performance evaluation of wireless ad hoc, sensor, and ubiquitous networks*, pages 25–31, New York, NY, USA, 2007. ACM.
- [30] CEPT 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, July 2007.
- [31] D. C. Cox, R. R. Murray, and A. W. Norris. Antenna height dependence of 800 MHz attenuation measured in houses. *Vehicular Technology, IEEE Transactions on*, 34(2):108–115, may 1985.
- [32] Cuomo. Radio resource sharing for ad hoc networking with UWB. *IEEE Journal on Selected Areas in Communications*, 20(9):1722, 2002.
- [33] F. Cuomo and C. Martello. MAC principles for an ultra wide band wireless access. In *Proceedings of the IEEE Global Telecommunications Conference (GLOBECOM 2001)*, volume 6, pages 3548–3552, 2001.
- [34] Douglas S. J. De Couto. *High-Throughput Routing for Multi-Hop Wireless Networks*. PhD thesis, Massachusetts Institute of Technology, 2004.
- [35] L. De Nardis and M. G. Di Benedetto. Cognition in routing for low rate UWB networks. volume 3, pages 97–100, sept. 2008.
- [36] Maria-Gabriella Di Benedetto, Luca De Nardis, Guerino Giancola, and Daniele Domenicali. The Aloha access (UWB)² protocol revisited for IEEE 802.15.4a. *ST Journal of Research*, 2006.

- [37] Maria-Gabriella Di Benedetto, Luca De Nardis, Matthias Junk, and Guerino Giancola. (UWB)2: Uncoordinated, Wireless, Baseborn Medium Access for UWB Communication Networks. *Mobile Networks and Applications*, 10(5):663, 2005.
- [38] Jin Ding, Li Zhao, Sirisha R. Medidi, and Krishna M. Sivalingam. MAC Protocols for Ultra-Wide-Band (UWB) Wireless Networks: Impact of Channel Acquisition Time. In *Proceedings of SPIE ITCOM'2002*, pages 1953–54, 2002.
- [39] Guido Dolmans, Olivier Rousseaux, Li Huang, Ting Fu, Bert Gyselinkx, Stefano d'Amico, Andrea Baschiroto, Julien Ryckaert, and Bart van Poucke. UWB Radio Transceivers For Ultra Low Power and Low Data Rate Communications. In *Proc. of the 2007 IEEE International Conference on Ultra-Wideband*, pages 152–157, 2007.
- [40] G. Durisi and G. Romano. Performance Evaluation and Comparison of Different Modulation Schemes for UWB Multiaccess Systems. In *Proceedings of the IEEE International Conference on Communications ICC'03*, pages 2187–2191, 2003.
- [41] Giuseppe Durisi and Giovanni 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 (ICUWST)*, pages 157–161, 2002.
- [42] G.P. Efthymoglou, V.A. Aalo, and H. Helmken. Performance analysis of coherent DS-CDMA systems in a Nakagami fading channel with arbitrary parameters. *IEEE Transactions on Vehicular Technology*, 2(46):289–297, 1997.
- [43] Amre El-Hoiydi. Spatial TDMA and CSMA with Preamble Sampling for Low Power Ad Hoc Wireless Sensor Networks. In *Proceedings of the IEEE International Conference on Computers and Communications (ISCC), Taormina, Italy*, pages 685–692, 2002.
- [44] Amre El-Hoiydi. *Energy Efficient Medium Access Control for Wireless Sensor Networks*. PhD thesis, Swiss Federal Institute of Technology (EPFL), 2005.
- [45] Amre El-Hoiydi and Jean-Dominique Decotignie. WiseMAC: An Ultra Low Power MAC Protocol for Multi-hop Wireless Sensor Networks. In *1st International Workshop on Algorithmic Aspects of Wireless Sensor Networks*, 2004.
- [46] Christian Enz, Amre El-Hoiydi, Jean-Dominique Decotignie, and Vincent Peiris. WiseNET: an ultralow-power wireless sensor network solution. *IEEE Computer*, 37(8):62–70, 2004.
- [47] S. C. Ergen and P. Varaiya. PEDAMACS: power efficient and delay aware medium access protocol for sensor networks. *IEEE Transactions on Mobile Computing*, 5(7):920–930, July 2006.

- [48] Deborah Estrin, W. Ye, and J. Heidemann. Medium Access Control With Coordinated Adaptive Sleeping for Wireless Sensor Networks. *IEEE/ACM Transactions on Networking*, 12(3):493–506, June 2004.
- [49] K. Fall and K. Varadhan. *The NS Manual*, 2002.
- [50] John Farserotu, John Gerrits, Jérôme Rousselot, Gerrit van Veenendaal, Manuel Lobeira, and John Long. CSEM FM-UWB Proposal in response to the TG6 Call for Proposals. Technical Report DCN 276, IEEE P802.15 Working Group for Wireless Personal Area Networks, May 2009.
- [51] Federal Communications Commission (FCC). Revision of Part 15 of the Commission’s Rules Regarding Ultra-Wideband Transmission Systems. Technical report, 2002.
- [52] Laura Marie Feeney and Daniel Willkomm. Energy Framework: An extensible framework for simulating battery consumption in wireless networks. In *Proceedings of the 3rd International Omnet++ Workshop (OM-NeT’10), part of SimuTools 2010*, 2010.
- [53] Manuel Flury, Ruben Merz, and Jean-Yves Le Boudec. Clock-Offset Tracking Software Algorithms For IR-UWB Energy-Detection Receivers. *IEEE International Conference on Ultra-Wideband (ICUWB 2009)*, 2009.
- [54] Manuel Flury, Ruben Merz, Jean-Yves Le Boudec, and Julien Zory. Performance Evaluation of an IEEE 802.15.4a Physical Layer with Energy Detection and Multi-User Interference. *Proceedings of the IEEE International Conference on Ultra Wideband (ICUWB 2007)*, pages 663–668, 2007.
- [55] S. Krishnamurthy J. Stankovic G. Zhou, T. He. Models and Solutions for Radio Irregularity in Wireless Sensor Networks. *ACM Transactions on Sensor Networks*, 2(2):221–262, May 2006.
- [56] Pierre Gandolfo, Dusan Radovic, Milan Savic, and Djordje Simic. IEEE 802.15.4A UWB-IR Radio System for Telemedicine. *Proc. of the 2008 IEEE International Conference on Ultra-Wideband (ICUWB2008)*, 3:11–14, 2008.
- [57] J.J. Garcia-Luna-Aceves and Jyoti Raju. Distributed Assignment of Codes for Multihop Packet-Radio Networks. In *Proceedings of MILCOM 97*, pages 450–454, 1997.
- [58] John F. M. Gerrits, Michiel H. L. Kouwenhoven, Paul R. van der Meer, John R. Farserotu, and John R. Long. Principles and Limitations of Ultra-Wideband FM Communications Systems. *EURASIP Journal on Applied Signal Processing*, 3:382–396, 2005.
- [59] S. Gezici, Zhi Tian, G. B. Giannakis, H. Kobayashi, A. F. Molisch, H. V. Poor, and Z. Sahinoglu. Localization via ultra-wideband radios: a look at positioning aspects for future sensor networks. *Signal Processing Magazine, IEEE*, 22(4):70–84, July 2005.

-
- [60] S.S. Ghassemzadeh. Measurement and Modeling of an Ultra-Wide Bandwidth Indoor Channel. *IEEE Transactions on Communications*, 52(10):1786, 2004.
- [61] S.A. Ghorashi, B. Allen, M. Ghavami, and A.H. Aghvami. An Overview of MB-UWB OFDM. In *Proceedings of the IEE Seminar on Ultra Wideband Communication Technologies and System Design*, pages 107–110, 2004.
- [62] Guerino Giancola and Maria-Gabriella Di Benedetto. A novel approach for estimating multi-user interference in impulse radio UWB networks: The pulse collision model. *Signal Processing*, 86(9):2185–2197, 2006.
- [63] Gianni Giorgetti, Sandeep K. S. Gupta, and Gianfranco Manes. Localization using signal strength: to range or not to range? In *MELT '08: Proceedings of the first ACM international workshop on Mobile entity localization and tracking in GPS-less environments*, pages 91–96, New York, NY, USA, 2008. ACM.
- [64] Chunlong Guo, Lizhi Charlie Zhong, and J. M. Rabaey. Low power distributed MAC for ad hoc sensor radio networks. In *Global Telecommunications Conference, 2001. GLOBECOM '01. IEEE*, volume 5, pages 2944–2948, 2001.
- [65] Ashima Gupta and Prasant Mohapatra. A survey on ultra wide band medium access control schemes. *Computer Networks*, 51(11):2976–2993, 2007.
- [66] G. P. Halkes and K. G. Langendoen. Crankshaft: An Energy-Efficient MAC-Protocol for Dense Wireless Sensor Networks. In *4th European Conference on Wireless Sensor Networks (EWSN 2007)*, 2007.
- [67] G. P. Halkes and K. G. Langendoen. Experimental Evaluation of Simulation Abstractions for Wireless Sensor Network MAC Protocols. In *Proceedings of the 14th IEEE International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD 2009), Pisa, Italy*, 2009.
- [68] Guoyou He. Destination-Sequenced Distance Vector (DSDV) Protocol. *unknown*, 2003.
- [69] W. Heinzelman, J. Kulik, and H. Balakrishnan. Adaptive Protocols for Information Dissemination in Wireless Sensor Networks. In *Proceedings of the 5th ACM/IEEE Mobicom*, pages 174–185, 1999.
- [70] Jason L. Hill and David E. Culler. Mica: A Wireless Platform for Deeply Embedded Networks. *IEEE Micro*, 22(6):12–24, Nov./Dec. 2002.
- [71] Ralph Hocht and Harold Tomlinson. Delay-hopped transmitted-reference RF communications. In *Proc. of the IEEE International Conference on Ultra Wideband Systems and Technologies*, pages 265–269, 2002.

- [72] Chalermek Intanagonwiwat, Ramesh Govindan, and Deborah Estrin. Directed diffusion: a scalable and robust communication paradigm for sensor networks. In *MobiCom '00: Proceedings of the 6th annual international conference on Mobile computing and networking*, pages 56–67, New York, NY, USA, 2000. ACM.
- [73] Svilen Ivanov, André Herms, and Georg Lukas. Experimental validation of the ns-2 wireless model using simulation, emulation, and real network. In *In 4th Workshop on Mobile Ad-Hoc Networks (WMAN'07)*, pages 433–444, 2007.
- [74] Kyle Jamieson, Hari Balakrishnan, and Y. C. Tay. Sift: A MAC Protocol for Event-Driven Wireless Sensor Networks. In *Proceedings of the 3rd European Workshop on Wireless Sensor Networks (EWSN)*, pages 260–275, 2006.
- [75] David B. Johnson and David A. Maltz. *Mobile Computing*, volume 353 of *The International Series in Engineering and Computer Science*, chapter Dynamic Source Routing in Ad Hoc Wireless Networks, pages 153–181. Kluwer Academic Publishers, 1996.
- [76] Jurdak. U-MAC: a proactive and adaptive UWB medium access control protocol. *Wireless Communications and Mobile Computing*, 5(5):551, 2005.
- [77] P. Karn. MACA - A New Channel Access Method for Packet Radio. In *ARRL/CRRL Amateur Radio 9th Computer Networking Conference, September 22, 1990*.
- [78] Brad Karp and H. T. Kung. GPSR: greedy perimeter stateless routing for wireless networks. In *Proc. of the ACM International Conference on Mobile Computing and Networking (Mobicom '00)*, pages 243–254, New York, NY, USA, 2000. ACM.
- [79] Muhammad Gufran Khan, Jörgen Nordberg, Abbas Mohammed, and Ingvar Claesson. Performance Evaluation of RAKE Receiver for UWB Systems using Measured Channels in Industrial Environments. In *Proc. of the International Conference on Wireless Broadband and Ultra Wideband Communication (Auswireless 2006)*, 2006.
- [80] Y. B. Ko and N. H. Vaidya. Location Aided Routing in Mobile Ad Hoc Networks. In *Proceedings of the fourth annual ACM/IEEE international Conference on Mobile Computing and Networking (MobiCom '98)*, pages 66–75, 1998.
- [81] A. Köpke, M. Swigulski, K. Wessel, D. Willkomm, P.T. Klein Haneveld, T.E.V. Parker, O.W. Visser, H.S. Lichte, and S. Valentin. Simulating Wireless and Mobile Networks in OMNeT++: The MiXiM Vision. In *Simutools '08: Proceedings of the 1st international conference on Simulation tools and techniques for communications, networks and systems & workshops*, pages 1–8, ICST, Brussels, Belgium, 2008. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering).

- [82] Juha Korhonen. *Introduction to 3G mobile communications*, chapter 14: 3G Applications, page 423. Mobile Communications Series. Artech House, second edition edition, 2003.
- [83] Kurtis Kredon and Prasant Mohapatra. Medium access control in wireless sensor networks. *Computer Networks*, 51(4):961, 2007.
- [84] J. Kulik, W. R. Heinzelman, and H. Balakrishnan. Negotiation-Based Protocols for Disseminating Information in Wireless Sensor Networks. *Wireless Networks*, 8:169–85, 2002.
- [85] David Lachartre, Benoit Denis, Dominique Morche, Laurent Ouvry, Manuel Pezzin, Bernard Piaget, Jerome Prouvee, and Pierre Vincent. A 1.1nJ/b 802.15.4a-compliant fully integrated UWB transceiver in 0.13 μ m CMOS. *IEEE International Solid State Circuits Conference (ISSCC 2009)*, pages 312–313,313a, 2009.
- [86] Jean-Yves Le Boudec, Ruben Merz, Bozidar Radunovic, and Jorg Widmer. DCC-MAC: A Decentralized MAC Protocol for 802.15.4a-like UWB Mobile Ad-Hoc Networks Based on Dynamic Channel Coding. In *Proceedings of the First International Conference on Broadband Networks (BroadNets 2004)*, 2004.
- [87] F. Legrand, I. Bucaille, S. Hethuin, L. De Nardis, G. Giancola, M. di Benedetto, L. Blazevic, and P. Rouzet. U.C.A.N.'s Ultra Wide Band System: MAC and Routing protocols. In *Proceedings of International Workshop on Ultra Wideband Systems*, 2003.
- [88] E.-Y.A. Lin, J.M. Rabaey, and A. Wolisz. Power-efficient rendez-vous schemes for dense wireless sensor networks. In *IEEE International Conference on Communications*, volume 7, pages 3769–3776, 2004.
- [89] Huaping Liu. Multicode ultra-wideband scheme using chirp waveforms. *IEEE Journal on Selected Areas in Communications*, 24(4):885–891, 2006.
- [90] Shou-Chih Lo. Design of multichannel MAC protocols for wireless ad hoc networks. *International Journal of Network Management*, 19(5):399, 2009.
- [91] T. Luo, M. Motani, and V. Srinivasan. Cooperative Asynchronous Multichannel MAC: Design, Analysis, and Implementation. *Mobile Computing, IEEE Transactions on*, 8(3):338–352, March 2009.
- [92] H. Arslan M. E. Sahin, I. Guvenc. Optimization of energy detector receivers for UWB systems. 2:1386–1390, June 2005.
- [93] J. Li M. Jiang and Y. Tay. Cluster Based Routing Protocol (CBRP) Functional Specification. IETF MANET Working Group Internet Draft, 1999, August 1998. Available online at <http://tools.ietf.org/html/draft-ietf-manet-cbrp-spec-01>.

- [94] Stefan Mahlknecht and Michael Bock. CSMA-MPS: a minimum preamble sampling MAC protocol for low power wireless sensor networks. In *Proceedings of the 2004 IEEE International Workshop on Factory Communication Systems*, pages 73–80, 2004.
- [95] A. Manjeshwar and D. P. Agrawal. TEEN: a routing protocol for enhanced efficiency in wireless sensor networks. pages 2009–2015, apr 2001.
- [96] Ruben Merz, Jean-Yves Le Boudec, and Joerg Widmer. An architecture for wireless simulation in NS-2 applied to impulse-radio ultra-wide band networks. In *Proc. of the 10th Communications and Networking Simulation Symposium (CNS07)*, pages 256–263, San Diego, CA, USA, 2007. Society for Computer Simulation International.
- [97] Ruben Merz and Jean-Yves Le Boudec. Conditional bit error rate for an impulse radio UWB channel with interfering users. In *Proceedings of the IEEE International Conference on Ultra-Wideband (ICUWB 2005)*, pages 130–135, 2005.
- [98] Jeonghoon Mo, Jeonghoon Mo So, and Jean Walrand. Comparison of Multichannel MAC Protocols. *IEEE Transactions on Mobile Computing*, 7(1):50–65, 2008.
- [99] A.F. Molisch. Ultrawideband Propagation Channels-Theory, Measurement, and Modeling. *IEEE Transactions on Vehicular Technology*, 54(5):1528, 2005.
- [100] Andreas F. Molisch, Kannan Balakrishnan, Dajana Cassioli, Chia-Chin Chong, Shahriar Emami, Andrew Fort, Johan Karedal, Juergen Kunisch, Hans Schantz, Ulrich Schuster, and Kai Siwiak. IEEE 802.15.4a channel model - final report. Technical report, IEEE 802.15.4 Working Group, 2005.
- [101] S. Niranjayan, A. Nallanathan, and B. Kannan. Modeling of Multiple Access Interference and BER Derivation for TH and DS UWB Multiple Access Systems. *IEEE Transactions on Wireless Communications*, 5(10):2794–2804, 2006.
- [102] Bob O’Hara and Al Petrick. *The IEEE 802.11 Handbook: A Designer’s Companion*. IEEE Press, 1999.
- [103] Ian Oppermann, Lucian Stoica, Alberto Zabbachin, Zack Shelby, and Jussi Haapola. UWB wireless sensor networks: UWEN-a practical example. *IEEE communications magazine*, 42(12):S27–S28, 2004.
- [104] Laurent Ouvry. UWB bas débit et standardisation à l’IEEE 802.15.4a. Technical report, CEA LETI, <http://gdr-ondes.lss.supelec.fr/actu/231006Ouvry.pdf>, 2006.
- [105] V. D. Park and M. S. Corson. . In *Proc. of the IEEE International Conference on Computer Communications (Infocom 97)*, pages 1405–1413, April 1997.

- [106] N. Patwari, III Hero, A.O., M. Perkins, N. S. Correal, and R. J. O'Dea. Relative location estimation in wireless sensor networks. *IEEE Transactions on Signal Processing*, 51(8):2137–2148, Aug. 2003.
- [107] Vincent Peiris, Claude Amre, and Stéphane Bories. A 1V 433/868MHz 25kb/s-FSK 2kb/s-OOK RF Transceiver SoC in Standard Digital 0.18 μ m CMOS. *Proceedings of the International Solid State Circuits Conference (ISSCC 2005)*, pages 258–259, 2005.
- [108] Jun Peng, Liang Cheng, and Biplab Sikdar. A Wireless MAC Protocol with Collision Detection. *IEEE Transactions on Mobile Computing*, 6(12):1357–1369, 2007.
- [109] Charles E. Perkins and Pravin Bhagwat. Highly dynamic Destination-Sequenced Distance-Vector routing (DSDV) for mobile computers. *SIGCOMM Comput. Commun. Rev.*, 24(4):234–244, 1994.
- [110] Charles E. Perkins and Elizabeth M. Royer. Ad-hoc On-Demand Distance Vector Routing. In *WMCSA '99: Proceedings of the Second IEEE Workshop on Mobile Computer Systems and Applications*, page 90. IEEE Computer Society, 1999.
- [111] M. Pezzin and D. Lachartre. A Fully Integrated LDR IR-UWB CMOS Transceiver Based on "1.5-bit" Direct Sampling. *IEEE International Conference on Ultra-Wideband (ICUWB 2007)*, pages 642–647, 2007.
- [112] Joseph Polastre, Jason Hill, and David Culler. Versatile low power media access for wireless sensor networks. In *Proceedings of the 2nd international conference on Embedded networked sensor systems*, volume 10, pages 95–107, New York, NY, USA, 2004. ACM.
- [113] Joseph Polastre, Jason Hill, and David Culler. Versatile low power media access for wireless sensor networks. In *Proceedings of the 2nd international conference on Embedded networked sensor systems*, pages 95–107, New York, NY, USA, 2004. ACM.
- [114] Petar Popovski, Hiroyuki Yomo, and Ramjee Prasad. Dynamic Adaptive Frequency Hopping for Mutually Interfering Wireless Personal Area Networks. *IEEE Transactions on Mobile Computing*, 5(8):991–1003, 2006.
- [115] John G. Proakis. *Digital Communications*. McGraw-Hill, 1983.
- [116] M. Pursley. Performance Evaluation for Phase-Coded Spread-Spectrum Multiple-Access Communication - Part I: System Analysis. *IEEE Transactions on Communications*, 30(5):795–799, August 1977.
- [117] Gabriela Quintero Díaz de León and A. K. Skrivervik. Analysis of UWB Antennas for Carrier-Based UWB Impulse Radio. In *Proc. of the Second European Conference on Antennas and Propagation (EuCAP 2007)*, pages 1–5, 2007.
- [118] B. Radunovic and J.-Y. LeBoudec. Optimal Power Control, Scheduling, and Routing in UWB Networks. *IEEE Journal on Selected Areas in Communications*, 22(7):1252, 2004.

-
- [119] V. Rajendran, J. J. Garcia-Luna-Aceves, and K. Obraczka. Energy-efficient, application-aware medium access for sensor networks. In *Proceedings of the 2nd IEEE International Conference on Mobile Adhoc And Sensor Systems (MASS 2005)*, page 630, Nov. 2005.
- [120] Venkatesh Rajendran, Katia Obraczka, and J. J. Garcia-Luna-Aceves. Energy-Efficient, Collision-Free Medium Access Control for Wireless Sensor Networks. *Wireless Networks*, 12(1):63–78, 2006.
- [121] T. S. Rappaport, S. Y. Seidel, and K. Takamizawa. Statistical channel impulse response models for factory and open plan building radio communication system design. *IEEE Transactions on Communications*, 39:794–806, 1991.
- [122] Theodore S. Rappaport. *Wireless Communications: Principles & Practice*, chapter 8 Multiple Access Techniques for Wireless Communications, pages 410–417. Prentice Hall PTR, 1996.
- [123] Bixio Rimoldi. *Principles of Digital Communications*. 2007.
- [124] Roberts. ALOHA packet system with and without slots and capture. *ACM SIGCOMM Computer Communication Review*, 5(2):28, 1975.
- [125] Jérôme Rousselot and Jean-Dominique Decotignie. A High-Precision Ultra Wideband Impulse Radio Physical Layer Model for Network Simulation. In *Proc. of the Second Int. Omnet++ Workshop, Simu'TOOLS (Omnet2009)*, pages 1–8, 2009.
- [126] Jérôme Rousselot and Jean-Dominique Decotignie. On the best way to cut a body area network's wires. In *Proc. of the IEEE International Conference on Communications (ICC 2010)*, 2010.
- [127] Jérôme Rousselot and Jean-Dominique Decotignie. When High Performance meet Ultra Low Power: the WideMac High Availability Protocol. In *Proc. of the IEEE International Conference on Ultra Wideband (ICUWB 2010)*, 2010.
- [128] Jérôme Rousselot, Jean-Dominique Decotignie, Marc Aoun, Peter van der Stok, Ramon Serna Olivera, and Gerhard Fohler. Accurate Timeliness Simulations for Real-Time Wireless Sensor Networks. In *Proc. of the 3rd UKSim European Symposium on Computer Modeling and Simulation (EMS 2009)*, pages 476–481. IEEE Press, 2009.
- [129] Jérôme Rousselot, Amre El-Hoiydi, and Jean-Dominique Decotignie. Low Power Medium Access Control Protocols for Wireless Sensor Networks. In *European Wireless Conference 2008 (EW 2008)*, pages 1–5, 2008.
- [130] Jérôme Rousselot, Amre El-Hoiydi, and Jean-Dominique Decotignie. WideMac: a low power and routing friendly MAC protocol for Ultra Wideband sensor networks. In *Proc. of the International Conference on Ultra-Wideband*, pages 105–108, 2008.

- [131] Julien Ryckaert, M. Badaroglu, C. Desset, V. De Heyn, G. van der Plas, Piet Wambacq, Bart van Poucke, and Stéphane Donnay. Carrier-based UWB impulse radio: simplicity, flexibility, and pulser implementation in 0.18-micron CMOS. In *IEEE International Conference on Ultra-Wideband (ICUWB 2005)*, pages 432–437, 2009.
- [132] Julien Ryckaert, Mustafa Badaroglu, Vincent De Heyn, Geert Van der Plas, Pierluigi Nuzzo, Andrea Baschirotto, Stefano D’Amico, Claude Desset, Hans Suys, Michael Libois, Bart Van Poucke, Piet Wambacq, and Bert Gyselinckx. A 16 mA UWB 3-to-5GHz 20 Mpulses/s Quadrature Analog Correlation Receiver in 0.18micrometer CMOS. In *International Solid-State Circuits Conference*, pages 368–377, 2006.
- [133] Julien Ryckaert, Geert Van der Plas, Vincent De Heyn, Claude Desset, Geert Vanwijnsberghe, Bart Van Poucke, and Jan Craninckx. A 0.65-to-1.4nJ/burst 3-to-10GHz UWB Digital TX in 90nm CMOS for IEEE 802.15.4a. In *International Solid State Circuits Conference*, pages 120–591, 2007.
- [134] A. Saleh and R. Valenzuela. A Statistical Model for Indoor Multipath Propagation. *IEEE Journal on Selected Areas in Communications*, 5(2):128–137, Feb 1987.
- [135] Y. Sasson, D. Cavin, and A. Schiper. Probabilistic broadcast for flooding in wireless mobile ad hoc networks. In *Proc. of the IEEE Wireless Communications and Networking Conference (WCNC 2003)*, volume 2, pages 1124–1130, march 2003.
- [136] R. A. Scholtz. Multiple Access with Time-Hopping Impulse Modulation. In *Proc. of the Military Communications Conference (MILCOM’93)*, pages 447–450, 1993.
- [137] R. A. Scholtz, D. M. Pozar, and W. Namgoong. Ultra-Wideband Radio. *EURASIP Journal on Applied Signal Processing*, pages 252–272, 2005.
- [138] S. Y. Seidel and T. S. Rappaport. 914 MHz path loss prediction models for indoor wireless communications in multifloored buildings. *Antennas and Propagation, IEEE Transactions on*, 40(2):207–217, feb 1992.
- [139] Claude E. Shannon. Communication in the Presence of Noise. *Proceedings of the IEEE*, 86(2):447–457, 1998.
- [140] Xuemin S. Shen, Weihua Zhuang, Hai Jiang, and Jun Cai. Medium Access Control in Ultra-Wideband Wireless Networks. *IEEE Transactions on Vehicular Technology*, 54(5):1663, 2005.
- [141] Hongsan Sheng, P. Orlik, A. M. Haimovich, . Jr. Cimini, L. J., and Jinyun Zhang. On the spectral and power requirements for ultra-wideband transmission. *IEEE International Conference on Communications (ICC’03)*, 1:738–742, May 2003.
- [142] Nan Shi, Liang Xia, and Ignas G. Niemegeers. A Novel Approach for the Link Layer in Impulse-based UWB Ad Hoc Networks. *Wireless Personal Communications*, 42(2):143–159, 2007.

- [143] Nan Shi, Yimeng Yang, and Ignas G. Niemegeers. MAC Protocol Design for Impulse Radio UWB Based WPANs. *Local Computer Networks, Annual IEEE Conference on*, 0:939–948, 2007.
- [144] Xiaolei Shi and Guido Stromberg. SyncWUF: An Ultra Low-Power MAC Protocol for Wireless Sensor Networks. *IEEE Transactions on Mobile Computing*, 6(1):115–125, 2007.
- [145] G. N. Shirazi. Joint cooperative MAC and routing in IR-UWB networks. pages 575–579, sept. 2009.
- [146] William Silvert. Modeling as a Discipline. *International Journal of General Systems*, 30(3):261–282, 2001.
- [147] SkyCross. SMT-3TO10M 3.1-10 GHz Ultra-Wideband Antenna. Technical report, 2009.
- [148] Jungmin So and Nitin H. Vaidya. Multi-channel mac for ad hoc networks: handling multi-channel hidden terminals using a single transceiver. In *MobiHoc '04: Proceedings of the 5th ACM international symposium on Mobile ad hoc networking and computing*, pages 222–233, New York, NY, USA, 2004. ACM.
- [149] IEEE Computer Society. IEEE 802.15 WPAN™ Task Group 6 Body Area Networks (BAN). <http://www.ieee802.org/15/pub/TG6.html>.
- [150] IEEE Computer Society. *IEEE 802.15.4 WPAN Low Data Rate Alternative PHY Layer*. <http://www.ieee802.org/15/pub/TG4a.html>.
- [151] IEEE Computer Society. *IEEE Std 802.15.4-2006, IEEE Standard for Information technology-Telecommunications and information exchange between systems-Local and metropolitan area networks-Specific requirements- Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (LR-WPANs)*. IEEE Press.
- [152] IEEE Computer Society. *IEEE Std 802.11-1999, IEEE Standard for Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements - Part 11*. IEEE Press, 1999.
- [153] IEEE Computer Society. *IEEE Std 802.15.1-2005, IEEE Standard for Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements - Part 15.1*. IEEE Press, 2005.
- [154] IEEE Computer Society. *IEEE Std 802.11-2007, IEEE Standard for Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements - Part 11*. IEEE Press, 2007.
- [155] IEEE Computer Society. IEEE P802.15 Wireless Personal Area Networks - MedWiN MAC and Security Proposal – Documentation. Technical report, September 2009. <http://mentor.ieee.org/802.15/dcn/09/15-09-0327-01-0006-medwin-mac-and-security-proposal-documentation.pdf>.

- [156] David H. Staelin, Ann W. Morgenthaler, and Jin Au Kong. *Electromagnetic Waves*. Prentice-Hall International, 1994.
- [157] Yanjun Sun, Omer Gurewitz, and David B. Johnson. RI-MAC: a receiver-initiated asynchronous duty cycle MAC protocol for dynamic traffic loads in wireless sensor networks. In *SenSys '08: Proceedings of the 6th ACM conference on Embedded network sensor systems*, pages 1–14, New York, NY, USA, 2008. ACM.
- [158] Cisco Systems. Cisco Aironet Access Point FAQ. Technical report, December 2009.
- [159] Texas Instruments. *2.4 GHz IEEE 802.15.4 / ZigBee-ready RF Transceiver*, 2007. <http://www.ti.com>.
- [160] Texas Instruments. *CC1100: Low-Cost Low-Power Sub-1 GHz RF Transceiver*, 2008. <http://www.ti.com>.
- [161] Z. Tian and B. M. Sadler. Weighted energy detection of ultra-wideband signals. In *Proc. of the IEEE 6th Workshop on Signal Processing Advances in Wireless Communications*, pages 1068–1072, June 2005.
- [162] F. Tufvesson, S. Gezici, and A. F. Molisch. Ultra-Wideband Communications using Hybrid Matched Filter Correlation Receivers. *Wireless Communications, IEEE Transactions on*, 5(11):3119–3129, November 2006.
- [163] Tijs van Dam and Koen Langendoen. An Adaptive Energy-Efficient MAC Protocol for Wireless Sensor Networks. In *Proceedings of the 1st international conference on Embedded networked sensor systems*, pages 171–180, 2003.
- [164] L.F.W. van Hoesel and P.J.M. Havinga. A Lightweight Medium Access Protocol (LMAC) for Wireless Sensor Networks. In *Proceedings of the 1st International Workshop on Networked Sensing Systems (INSS 2004)*, pages 946–953, 2004.
- [165] András Varga. The OMNeT++ Discrete Event Simulation System. In *Proceedings of the European Simulation Multiconference (ESM'2001)*, pages 319–324, 2001.
- [166] Vladimir M. Vishnevsky, Andrey I. Lyakhov, Alexander A. Safonov, Shaomin S. Mo, and Alexander D. Gelman. Study of Beaconing in Multihop Wireless PAN with Distributed Control. *IEEE Transactions on Mobile Computing*, 7(1):113–126, 2008.
- [167] Jinghao Xu ; Bojan Peric ; Branimir Vojcic. Performance of energy-aware and link-adaptive routing metrics for ultra wideband sensor networks. *Mobile Networks and Applications*, 11(4):509–519, August 2006.
- [168] M. Weisenhorn and W. Hirt. Robust noncoherent receiver exploiting UWB channel properties. In *Proc. of the International Workshop on Ultra Wideband Systems, joint with Conference on Ultrawideband Systems and Technologies.*, pages 156–160, May 2004.

- [169] M. L. Welborn. System considerations for ultra-wideband wireless networks. In *Radio and Wireless Conference, 2001. RAWCON 2001. IEEE*, pages 5–8, 2001.
- [170] Moe Z. Win and Robert A. Scholtz. Ultra-Wide Bandwidth Time-Hopping Spread-Spectrum Impulse Radio for Wireless Multiple-Access Communications. *IEEE Transactions on Communications*, 48(4):679–691, 2000.
- [171] Klaus Witrisal. Statistical Analysis of the IEEE 802.15.4a UWB PHY over Multipath Channels. In *Proceedings of the IEEE Wireless Communications & Networking Conference (WCNC'08)*, pages 130–135, 2008.
- [172] Alec Woo. *A Holistic Approach to Multihop Routing in Sensor Networks*. PhD thesis, University of California, Berkeley, 2004.
- [173] Alec Woo, Terence Tong, and David Culler. Taming the underlying challenges of reliable multihop routing in sensor networks. In *SenSys '03: Proceedings of the 1st international conference on Embedded networked sensor systems*, pages 14–27, New York, NY, USA, 2003. ACM.
- [174] Yafeng Wu, J. A. Stankovic, Tian He, and Shan Lin. Realistic and Efficient Multi-Channel Communications in Wireless Sensor Networks. In *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, pages 1193–1201, April 2008.
- [175] Kaixin Xu, M. Gerla, and Sang Bae. How effective is the IEEE 802.11 RTS/CTS handshake in ad hoc networks. In *GLOBECOM '02: Proceedings of the IEEE Global Telecommunications Conference*, volume 1, pages 72–76, Nov. 2002.
- [176] Wei Ye, John Heidemann, and Deborah Estrin. An energy-efficient MAC protocol for wireless sensor networks. In *21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, volume 3, pages 1567–1576, 2002.
- [177] Wei Ye, Fabio Silva, and John Heidemann. Ultra-Low Duty Cycle MAC with Scheduled Channel Polling. In *4th international conference on Embedded networked sensor systems*, pages 321–334, 2006.
- [178] Fan Yu, Tao Wu, and S. Biswas. Toward In-Band Self-Organization in Energy-Efficient MAC Protocols for Sensor Networks. *Mobile Computing, IEEE Transactions on*, 7(2):156–170, Feb. 2008.
- [179] Fan Yu, Tao Wu, and S. Biswas. Toward In-Band Self-Organization in Energy-Efficient MAC Protocols for Sensor Networks. *Mobile Computing, IEEE Transactions on*, 7(2):156–170, Feb. 2008.
- [180] Y. Yu, R. Govindan, and D. Estrin. Geographical and Energy Aware Routing: A Recursive Data Dissemination Protocol for Wireless Sensor Networks. Technical report, 2001.
- [181] T. Yucek and H. Arslan. A survey of spectrum sensing algorithms for cognitive radio applications. *Communications Surveys & Tutorials, IEEE*, 11(1):116–130, Quarter 2009.

-
- [182] J. Zhang, P. Orlik, Z. Sahinoglu, A. F. Molisch, and P. Kinney. UWB Systems for Wireless Sensor Networks. *Proceedings of the IEEE*, 97(2):313–331, 2009.
 - [183] Jingbin Zhang, Gang Zhou, Chengdu Huang, S. H. Son, and J. A. Stankovic. TMMAC: An Energy Efficient Multi-Channel MAC Protocol for Ad Hoc Networks. In *Communications, 2007. ICC '07. IEEE International Conference on*, pages 3554–3561, June 2007.
 - [184] G. Zhou, C. Huang, T. Yan, T. He, J. A. Stankovic, and T. F. Abdelzaher. MMSN: Multi-Frequency Media Access Control for Wireless Sensor Networks. In *INFOCOM 2006. 25th IEEE International Conference on Computer Communications. Proceedings*, pages 1–13, April 2006.
 - [185] Jiang Zhu and Abraham O. Fapojuwo. A complementary code-CDMA-based MAC protocol for UWB WPAN system. *EURASIP journal on wireless communications and networking*, 2:249–259, 2005.
 - [186] Jihui Zhang Qian Zhang Bo Li Xiaonan Luo Wenwu Zhu. Energy-efficient routing in mobile ad hoc networks: mobility-assisted case. *IEEE Transactions on Vehicular Technology*, 55(1):369–379, January 2006.

Publications

Journal papers

- J. Rousselot, J.-D. Decotignie, An Ultra Wideband Impulse Radio PHY Layer Model for Network Simulation. In *Simulation: Transactions of the Society for Modeling and Simulation International, Special Issue on Software Tools, Techniques and Architectures for Computer Simulation*, accepted, 2010.
- J. Farserotu, J. Gerrits, J. Rousselot, Low power and robust PHY-MAC solution for Medical BAN. In *IEICE Transactions on Communications, Special Issue on Body Area Networks, Invited Paper*, vol. E93-B, nb. 4, April 2010.

Conference papers

- J. Rousselot, J.-D. Decotignie, When High Performance meets Ultra Low Power: the WideMac High Availability Protocol. In *Proc. of the IEEE International Conference on Ultra Wideband (ICUWB 2010)*, accepted, Nanjing, China, 20-23 September 2010.
- J. Rousselot, J.-D. Decotignie, On the best way to cut a body area network's wires. In *Proc. of the IEEE International Conference on Communications (ICC 2010)*, Cape Town, South Africa, 23-27 May 2010.
- J. Rousselot, Ph. Dallemagne, J.-D. Decotignie, Deployments of Wireless Sensor Networks performed by CSEM. In *Proc. of the International Conference on Cognitive Systems with Interactive Sensors (COGIS'09)*, Paris, France, 16-18 November 2009.
- J. Rousselot, J.-D. Decotignie, M. Aoun, P. van der Stok, R. Serna Oliver, G. Fohler, Accurate Timeliness Simulations for Real-Time Wireless Sensor Networks. In *Proc. of the 3rd European Symposium on Computer Modeling and Simulation (EMS 2009)*, pp. 476-481, Athens, Greece, 2009.
- J. Rousselot, J.-D. Decotignie, Wireless Communication Systems for Continuous Multiparameter Health Monitoring. In *Proc. of the IEEE International Conference on Ultra-Wideband (ICUWB 2009)*, invited paper, pp. 480-484, Vancouver, Canada, September 2009.

- J. Rousselot, J.-D. Decotignie, A High-Precision Ultra Wideband Impulse Radio Physical Layer Model for Network Simulation. In Proc. of the Second International Omnet++ Workshop, part of the 3rd ICST International Conference on Simulation Tools and Techniques (Simu'TOOLS 2009), pp. 1-8, Rome, 6 March 2009.
- J. Rousselot, A. El-Hoiydi, J.-D. Decotignie, WideMac: a Low Power and Routing Friendly MAC Protocol for Ultra Wideband Sensor Networks. In Proc. of the IEEE International Conference on Ultra-Wideband (ICUWB 2008), pp. 105-108, Hannover, Germany, 10-12 September 2008.
- J. Rousselot, J. Farserotu and J.-D. Decotignie, Reconciling Ultra Low Power Consumption, Low Latency and High Throughput Communications : The WiseMAC-HA Protocol. In Proc. of International Workshop on Future Wellness and Medical ICT Systems (FEELIT 2008), part of the 11th International Symposium on Wireless Personal and Multimedia Communications (WPMC 2008), Invited Paper, Saariselkä, Finland, 9 September 2008.
- J. Rousselot, A. El-Hoiydi, J.-D. Decotignie, Low Power Medium Access Control Protocols for Wireless Sensor Networks. In Proc. of the 14th European Wireless Conference 2008 (EW 2008), pp. 1-5, Prague, Czech Republic, 22-25 June 2008.
- J. Rousselot, A. El-Hoiydi, J.-D. Decotignie, Trade-off Analysis of Communication Protocols for Wireless Sensor Networks. Proc. of the 1st Workshop on Adaptive and Reconfigurable Embedded Systems (APRES 2008), part of the 14th IEEE Real-Time and Embedded Technology and Applications Symposium, pp. 36-39, St-Louis, USA, 21 April 2008.
- J. Rousselot, A. El-Hoiydi, J.-D. Decotignie, Performance evaluation of the IEEE 802.15.4A UWB physical layer for Body Area Networks. Proc. of the 12th IEEE International Symposium on Computers and Communications (ISCC 2007), pp. 969-974, Aveiro, Portugal, July 1-4 2007.
- J. Rousselot, A. El-Hoiydi, J.-D. Decotignie, On the Problem of Near-Far Interference with Impulse Ultra Wide Band Radios. European Ultra Wide Band Radio Technology Workshop (UWB 2007), Grenoble, France, May 10-11 2007.

Standard Proposal & Technical Reports

- J. Farserotu, J. Gerrits, J. Rousselot, CSEM FM-UWB Proposal in response to the TG6 Call for Proposals, May 2009.
- Ph. Stanley-Marbell, T. Basten, J. Rousselot, R. Serna Oliver, K. Holger, M. Geilen, R. Hoes, J.-D. Decotignie, G. Fohler. System Models in Wireless Sensor Networks, Electronic Systems Reports, Eindhoven University of Technology, Department of Electrical Engineering, ISSN 1574-9517, May 2008.

Curriculum Vitae

Jérôme Rousselot was born in Metz, France in 1981. He studied at Université Libre de Bruxelles, Ecole Polytechnique and earned a Master of Science in Computer Engineering in September 2005. For his master's thesis ("Déploiement d'un réseau sans fil pour des instituts de microfinance dans le cadre de la coopération au développement"), he studied the problem of enabling digital communications between isolated agencies of the microfinance institute Coopedu (Kigali, Rwanda), in cooperation with the international non governmental organization Aquadev. He competitively obtained public funding for the travel expenses from the Commission Universitaire au Développement and private sponsorship from Cisco Systems Belgium in the form of hardware. After this experience in Kigali, he was an exchange student at Ecole Polytechnique Fédérale de Lausanne from March 2005 to July 2005.

He started working as a PhD student at Centre Suisse d'Electronique et de Microtechnique (CSEM), Neuchâtel, in October 2005, and was accepted at the EPFL Doctoral School of Computer and Communication Sciences in May 2006. He was involved in the National Competence Center for Research on Mobile Information and Communication Systems (NCCR-MICS), and in the EU IST Wirelessly Accessible Sensor Populations (WASP) and PULSERS-II projects. During his thesis, he participated to the inclusion of CSEM FM-UWB technology in the upcoming IEEE 802.15.6 standard on medical body area networks.

Jérôme Rousselot is or has been a reviewer for Elsevier Computer Networks, Springer Wireless Networks, Wiley Wireless Communications and Mobile Computing, Hindawi Research Letters, and ETRI Journal. He is a member of the International Omnet++ Workshop organizing committee.