



Universidade do Minho
Escola de Engenharia

Helena Fernández López


Remote Vital Signs Monitoring Based on Wireless Sensor Networks

Helena Fernández López **Remote Vital Signs Monitoring Based on Wireless Sensor Networks**

UMinho | 2011

Outubro de 2011





Universidade do Minho
Escola de Engenharia

Helena Fernández López

Remote Vital Signs Monitoring Based on Wireless Sensor Networks

Programa Doutoral em Líderes para as Indústrias Tecnológicas

Trabalho efectuado sob a orientação de:
Professor Doutor Ricardo João Ferreira Simões
Professor Doutor José Higinio G. Correia
Professor Doutor José Augusto Afonso

Outubro de 2011

DECLARAÇÃO

Nome: Helena Fernández López

Endereço electrónico: hlopez@dei.uminho.pt

Telefone: 911527627

Número do Passaporte: AE474577R

Título dissertação /tese

Remote Vital Signs Monitoring Based on Wireless Sensor Networks

Orientador(es):

Professor Doutor Ricardo João Ferreira Simões

Professor Doutor José Higinio G. Correia

Professor Doutor José Augusto Afonso

Ano de conclusão: 2011

Designação do Doutoramento:

Programa Doutoral em Líderes para as Indústrias Tecnológicas

É AUTORIZADA A REPRODUÇÃO INTEGRAL DESTA TESE/TRABALHO APENAS PARA EFEITOS DE INVESTIGAÇÃO, MEDIANTE DECLARAÇÃO ESCRITA DO INTERESSADO, QUE A TAL SE COMPROMETE.

Universidade do Minho, de Outubro de 2011.

Assinatura: _____

Acknowledgements

I would like to thank my supervisors, Professors Ricardo Simões, José Higinio Correia, and José Afonso, for their continuous support and guidance during this work. I also would like to thank Professor Christopher Magee, who kindly received me at Massachusetts Institute of Technology.

I am grateful to Fundação para a Ciência e a Tecnologia, who supported this work under grant SFRH/BD/39408/2007, 3^o Quadro Comunitário de Apoio, and project “AAL4ALL”, n^o 13852 – QREN SI&IDT, co-financed by FEDER, Programa Operacional Factores de Competitividade. Clinical and partial financial support for the case-study has been provided by Grupo AMI - Assistência Médica Integral (Casa de Saúde de Guimarães, SA). In addition, I would like to thank the health care providers from Hospital Privado de Guimarães for their continuous support and to the questionnaires respondents for their invaluable contribution.

I would like to express my deepest appreciation to Gregory Stern, Ricardo Cunha-Vaz and André Lemos for their guidance and support during the internship.

I am in debt to my dear friend Ana Carolina Matos, who decisively contributed to the developed system. I am also grateful to Duarte Pereira and Pedro Macedo, two great programmers, who developed the Web-based applications and simulations. Additionally, I am grateful to David Araújo and Bruno Fernandes for their contributions and enthusiasm, and to Carlos Torres, Joel Almeida e Ângela Macedo, who provided invaluable support in several technical tasks.

I would like to acknowledge Professors Paulo Mendes and Carlos Lima, who helped me and Ana Carolina during the development of the electrocardiogram sensor, and to Professor Adriano Tavares who allowed us to work at his laboratory.

Finally, I am grateful to my in-laws, Cândida and Francisco; my parents, Maria del Carmen and Manuel; my children, Mariana and Carlos Eduardo; and my husband, Carlos, for their love and support. I have the most caring parents and in-laws, the most incredible children and the most adorable husband.

Remote Vital Signs Monitoring Based on Wireless Sensor Networks

Abstract

Governmental and private institutions face a major challenge to provide quality health care to a population consisting of a growing number of elderly and chronically ill patients. According to the World Health Organization, in 2006, the total global health expenditures exceeded US\$ 4 trillion and are rising in the majority of countries including Portugal which, during 2006, expended 9.9% of its gross domestic product in health care.

The use of remote vital signs monitoring systems increases the probability of early detection of risky situations, allows frequent monitoring of in-patients, elderly and chronically ill patients, and streamlines the work of health professionals. However, at present, these systems are expensive, complex and employ obtrusive sensors, which limit their application to intensive care units and cardiac intermediate care units.

This work is part of a project that aims to design, prototype and evaluate a remote vital signs monitoring system based on the IEEE 802.15.4 and ZigBee protocols, which allow the development of small low-power sensors. The prototype system comprises electrocardiogram/heart rate and axillary thermometer sensors, networking devices and three informatics applications that collect, process, and exhibit medical data.

The wireless sensors, the networking devices and one of the applications were developed under this work. Additionally, the wireless sensor network was evaluated through simulations at the MAC level and experimental and field tests. Field tests were performed at an in-patient floor of *Hospital Privado de Guimarães*, a Portuguese hospital. Finally, questionnaires were used to measure the satisfaction of users and catalog their critics and suggestions for improvement.

Simulations considered different topologies, operation modes and a crescent number of sensors and hops. Experimental and field tests confirmed most of the results obtained by simulations, but revealed that networks which did not assign transmission time slots to electrocardiogram sensors were unable to maintain a high delivery ratio. Contention between devices, aggravated by the inability of routers in receiving incoming packets during backoff, and collisions between packets generated by hidden-nodes were responsible for most message losses. On the other hand, beacon-enabled star IEEE 802.15.4 networks that assigned a guaranteed time slot to sensors were able to maintain a very high delivery ratio. In contrast, these networks are restricted in terms of the coverage area and the number of sensors. Also, field tests showed that under low traffic scenarios ZigBee nonbeacon-enabled networks can achieve a high delivery ratio even in presence of a high percentage of hidden-nodes.

Monitorização Remota de Sinais Vitais Baseada em Redes de Sensores sem Fios

Resumo

Instituições governamentais e privadas enfrentam um grande desafio para prestar cuidados de saúde de qualidade a uma população constituída por um número crescente de idosos e doentes crónicos. Segundo a Organização Mundial de Saúde, em 2006, a despesa mundial em saúde ultrapassou a quantia de 4 bilhões de dólares americanos e cresce anualmente na maioria dos países, incluindo Portugal, o qual, em 2006, gastou 9,9% do seu produto interno bruto em cuidados de saúde.

O uso de sistemas de monitorização remota de sinais vitais aumenta a probabilidade de deteção precoce de situações de risco, permite que doentes internados, idosos ou doentes crónicos sejam frequentemente monitorizados e agiliza o trabalho dos profissionais de saúde. No entanto, atualmente, estes sistemas são caros e complexos, o que limita a sua aplicação a alguns setores dos hospitais, tais como as unidades de cuidados intensivos e as unidades de cuidados intermédios na área da cardiologia.

O projeto no qual insere-se este trabalho visa a conceção, a prototipagem e a avaliação de um sistema de monitorização remota de sinais vitais com base nos protocolos IEEE 802.15.4 e ZigBee, os quais oferecem a possibilidade de construção de sensores com consumos energéticos muito baixos e reduzidas dimensões. O sistema consiste em sensores de eletrocardiograma/frequência cardíaca e temperatura axilar, dispositivos de rede e três aplicações que coletam, processam e apresentam o eletrocardiograma e os sinais vitais.

No âmbito deste trabalho foram desenvolvidos os sensores sem fios, os dispositivos de rede e uma das aplicações informáticas. Além disso, foi feita a avaliação do desempenho da rede de sensores sem fios através da análise de simulações a nível da camada de acesso ao meio (MAC) e de testes de laboratório e de campo. Os testes de

campo da rede de sensores sem fios foram executados em um dos pisos de internamento do Hospital Privado de Guimarães. Finalmente, foram usados questionários para medir a satisfação dos utilizadores e recolher críticas e sugestões de melhoria.

As simulações consideraram diferentes topologias e modos de operação, além de um número crescente de sensores e saltos. Testes experimentais e de campo confirmaram grande parte dos resultados obtidos por simulação mas, adicionalmente, revelaram que as redes constituídas por vários sensores de eletrocardiograma e que não reservaram um intervalo de tempo de transmissão aos sensores não foram capazes de manter uma elevada taxa de entrega de mensagens. Perdas de mensagens ocorreram devido a disputas entre sensores pelo acesso ao canal sem fios e devido a ocorrência de colisões de pacotes transmitidos por nós escondidos. Por outro lado, as redes baseadas no protocolo IEEE 802.15.4 que atribuíram um intervalo de tempo de transmissão a cada sensor conseguiram manter uma elevada taxa de entrega. Entretanto, essas redes são limitadas em termos da área de cobertura e do número de sensores. Adicionalmente, durante os testes de campo em cenários de tráfego reduzido, as redes ZigBee que não empregaram *beacons* atingiram uma elevada taxa de entrega mesmo na presença de uma grande percentagem de nós escondidos.

Contents

Acknowledgements	iii
Abstract	v
Resumo	vii
Contents.....	ix
Figures	xiv
Tables.....	xxi
Acronyms and abbreviations.....	xxv
Chapter 1.....	1
Introduction	1
1.1 Wireless communications.....	1
1.2 Vital signs.....	4
1.3 Acquisition of vital signs.....	8
1.3.1 ECG acquisition	8
1.3.2 Body temperature measurement.....	13
1.4 Vital signs monitors.....	15
1.4.1 Bedside monitors.....	15
1.4.2 Biotelemetry systems	16
1.5 Motivation	17
1.6 Approach and contributions.....	19
1.7 Thesis organization.....	21
Chapter 2.....	23

Wireless sensor networks and protocols.....	23
2.1 Wireless sensor networks: definition and applications.....	23
2.2 The IEEE 802.15.4 protocol.....	24
2.2.1 IEEE 802.15.4 protocol overview	24
2.2.2 Physical layer	27
2.2.3 Medium access control layer.....	31
2.3 The ZigBee protocol.....	35
2.3.1 ZigBee protocol overview	36
2.3.2 The Network layer.....	37
2.3.3 The Application layer.....	39
2.4 ZigBee versions comparison	45
2.5 Jennic’s programming environment	48
2.6 Other standard-based wireless communication protocols of interest	50
2.7 Summary.....	53
Chapter 3.....	57
Health monitoring systems based on wireless technologies	57
3.1 Patient monitoring general requirements.....	58
3.2 Regulation, standards and profiles.....	59
3.3 Wireless medical systems review	64
3.3.1 Non-wearable personal monitoring systems	64
3.3.2 Wearable personal monitoring systems.....	65
3.3.3 Systems based on wireless technologies for elderly care.....	68
3.3.4 Wireless hospital monitoring systems.....	72
3.4 Interviews	76
3.5 Summary.....	79
Chapter 4.....	81

The developed monitoring system.....	81
4.1 HM4All overview	81
4.2 Developed HM4All components	85
4.2.1 ECG sensor.....	85
4.3 Temperature sensor.....	93
4.3.1 Coordinators and routers	101
4.3.2 ZigBee-to-IP gateway	103
4.3.3 User acceptance.....	108
4.4 Summary.....	113
Chapter 5.....	115
HM4All evaluation based on simulation and laboratory tests	115
5.1 Performance assessment at MAC level	116
5.1.1 Simulation configuration and assumptions	116
5.1.2 Results	118
5.1.3 Simulation results outline.....	127
5.2 Performance assessment with hidden-nodes	128
5.2.1 Simulation configuration and assumptions	128
5.2.2 Results	129
5.2.3 Simulation results outline.....	132
5.3 Laboratory experiments.....	132
5.3.1 Experimental tests configuration.....	132
5.3.2 Results	135
5.3.3 Laboratory tests outline and discussion	137
5.3.4 Clock drift measurement	139
5.4 Summary.....	141
Chapter 6.....	143

HM4All evaluation based on field tests	143
6.1 Radio survey	143
6.2 Physical inspection	146
6.3 Connectivity tests	147
6.4 Range and link quality tests	149
6.5 Experimental tests with no hidden-nodes	152
6.5.1 Tests configuration	152
6.5.2 Results	156
6.5.3 Router deadlock.....	164
6.5.4 Discussion	168
6.6 Experimental tests with hidden-nodes	168
6.6.1 Tests configuration	169
6.6.2 Results for ECG traffic.....	169
6.6.3 Results for heart rate traffic.....	172
6.6.4 Discussion	173
6.7 Comparing the communication performance of HM4ALL with a commercial vital signs monitoring system.....	174
6.8 Critical factors on the performance of ZigBee networks	175
6.8.1 Mobility issues	175
6.8.2 Use of end-to-end acknowledged messages.....	177
6.8.3 An option to IEEE 802.15.4: the eLPTR protocol	177
6.9 Summary.....	178
Chapter 7	181
Conclusions and future work.....	181
7.1 Conclusions	181
7.2 Future work.....	186

Appendix A	Schematic diagram of the electrocardiogram sensor board.....	A-1
Appendix B	Schematic diagram of the temperature sensor board.....	B-1
Appendix C	Temperature sensor battery lifetime calculations.....	C-1
Appendix D	Schematic diagram of the coordinators and routers.....	D-1
Appendix E	Questionnaires	E-1
Appendix F	Range tests	F-1

Figures

Figure 1 – Protocol stack model.....	2
Figure 2 – ECG waveform (adapted from [25]).....	9
Figure 3 – Electrical conduction system of the heart (adapted from [157]).....	10
Figure 4 – Simplified single-channel ECG.....	12
Figure 5 – The wearable ECG sensor developed: (a) printed circuit board front side; (b) printed circuit board rear side; and (c) user case top side.	13
Figure 6 – Simplified block diagram of a thermistor-based temperature sensor.....	14
Figure 7 – The temperature sensor prototype version.	15
Figure 8 – The 1500 Patient Monitor (adapted from [213]).....	16
Figure 9 – The Dinamap Carescape V100 vital signs monitor (adapted from [72])......	16
Figure 10 – Micropaq patient unit: a) A patient unit worn in the hospital environment (adapted from [210]); b) User interface snapshot and external appearance (reproduced courtesy of Welch Allyn).	17
Figure 11 – Individual patient monitoring window.....	21
Figure 12 – Star (left) and peer-to-peer (right) topologies [91, 92].	25
Figure 13 – The cluster-tree topology [91, 92].....	26
Figure 14 – Superframe structure [91, 92].	26
Figure 15 – Alignment between IEEE 802.11b and IEEE 802.15.4 PHY channels [91, 92].....	29
Figure 16 – The PHY layer model [91, 92].....	29
Figure 17 – The MAC model [91, 92].....	31
Figure 18 – Unslotted version of the CSMA-CA mechanism defined in the 2003 version of the protocol [91, 92].	33

Figure 19 – Slotted version of the CSMA-CA mechanism defined in the 2003 version of the protocol [91, 92].	34
Figure 20 – Hidden-node scenarios.	35
Figure 21 – Exposed-node scenario.	35
Figure 22 – ZigBee model [220, 224].	37
Figure 23 – Profile definition.	41
Figure 24 – Multicasting in ZigBee PRO.	47
Figure 25 – Flow diagram of a generic application developed for the JN5139 module [113].	49
Figure 26 – Star and cluster tree topologies in ANT.	52
Figure 27 – Stages of a wireless medical device [87].	60
Figure 28 – WellDoc remote application is asking a patient to confirm the blood glucose value sent by the blood glucose meter using a Bluetooth connection (adapted from [214]).	65
Figure 29 – Cricket, Telos, Mica2 and MicaZ motes (adapted from [204]).	66
Figure 30 – CardioNet system patient devices (adapted from [26]).	67
Figure 31 – Nuvant system, from Corventis (adapted from [33]).	68
Figure 32 – A patch-like ECG sensor developed by IMEC/Holst Center (adapted from [158]).	68
Figure 33 – Flowie user interface (adapted from [6]).	70
Figure 34 – Actibelt patient unit embedded in a belt buckle (used with permission, adapted from [1]).	72
Figure 35 – ApexPro patient unit (adapted from [70]).	73
Figure 36 – Aingeal electrodes and clip-on patient device (adapted from [98]).	74
Figure 37 – The patient unit designed by MIT researchers under the SMART project (adapted from [36]).	76
Figure 38 – HM4All high-level system architecture.	82
Figure 39 – HM4All conceptual diagram.	84

Figure 40 – Monitoring page presenting data from several sensors (all data presented in this page was previously recorded for testing purposes).....	84
Figure 41 – ECG sensor block diagram.....	86
Figure 42 – ECG sensor case and electronic PCB.....	86
Figure 43 – ECG sensor cases: a) the silicone mold used to produce the final prototype cases; and b) an ECG case just extracted from the mold.....	87
Figure 44 – Rear side of the ECG sensor.	87
Figure 45 – Pan-Tompkins pre-processor.	88
Figure 46 – Results of ECG pre-processing. From top to bottom: ECG waveform; output of the band-pass filtering; output of the squaring filter; and output of the integrating moving-window filter.....	89
Figure 47 – Results of the Pan-Tompkins pre-processing using different HP filters. Top: original HP filter. Bottom: HP FIR filter designed to further attenuate P and T waves.	90
Figure 48 – User interface of the QRS detector implemented in C# language.	91
Figure 49 – JZA_vAppEventHandler flowchart.	92
Figure 50 – Temperature sensor prototype: (a) assembled with arm band and probe and (b) being worn.	94
Figure 51 – Exacon 400-series probes R-T characteristic.....	95
Figure 52 – Simplified block diagram of the temperature sensor.	96
Figure 53 – Schematics diagram of the power supply of the temperature sensor.....	96
Figure 54 – Schematics diagram of the front-end of the temperature sensor.....	97
Figure 55 – State diagram of the temperature sensor application.	100
Figure 56 – Coordinator/router (power on and activity LEDs not shown).	101
Figure 57 – Coordinator/router electronic PCB: (a) top side and (b) rear side.	102
Figure 58 – Data framing structure.....	102
Figure 59 – Class diagram of the ZigBee-to-IP application.....	104

Figure 60 – State diagram of the procedure used to process bytes received by the serial port interface.....	105
Figure 61 – Code snippet used to construct the string used to send data along with the post request.	106
Figure 62 – ZigBee-to-IP gateway user interface.....	107
Figure 63 – Evaluation of the ECG sensor made by patients.....	109
Figure 64 – Evaluation of the ECG sensor made by health care professionals.....	109
Figure 65 – Evaluation of the temperature sensor made by patients.....	110
Figure 66 – Evaluation of the temperature sensor made by health care professionals.	111
Figure 67 – System and concept evaluation.	112
Figure 68 – Missing functions.	112
Figure 69 – Negative points and suggestions.	113
Figure 70 – Star and tree topologies.....	117
Figure 71 – Collided packet ratio curves for an increasing number of EDs that transmit (a) raw ECG data and (b) compressed ECG data.	119
Figure 72 – Failed transmission attempt ratio curves for an increasing number of EDs that transmit (a) raw ECG data and (b) compressed ECG data.	119
Figure 73 – Normalized throughput curves for an increasing number of EDs considering: (a) star topology, raw ECG data; (b) tree topology, raw ECG data; (c) star topology, compressed ECG data; and (d) tree topology, compressed ECG data.	121
Figure 74 – DR curves for networks that relay the traffic generated by EDs that transmit (a) raw ECG data and (b) compressed ECG data.	122
Figure 75 – Mean and maximum delay for networks that only contain EDs that transmit raw ECG data.....	125
Figure 76 – Mean and maximum delay for networks that only contain EDs that transmit compressed ECG data.....	125
Figure 77 – JN5139 module current consumption per relevant activity.	126

Figure 78 – Energy consumption per message for networks that contain EDs that transmit (a) raw ECG data and (b) compressed ECG data.	127
Figure 79 – Collided packet ratio curves for networks that contain different percentages of hidden nodes and a crescent number of EDs that transmit (a) raw ECG data and (b) compressed ECG data.	129
Figure 80 – Failed transmission attempt ratio curves for networks that contain different percentages of hidden nodes and a crescent number of EDs that transmit (a) raw ECG data and (b) compressed ECG data.	130
Figure 81 – Throughput curves for networks that contain contain different percentages of hidden nodes and a crescent number of EDs that transmit (a) raw ECG data and (b) compressed ECG data.	131
Figure 82 – DR curves for networks that contain different percentages of hidden nodes and a crescent number of EDs that transmit (a) raw ECG data and (b) compressed ECG data.	131
Figure 83 – Test boards: (a) top side (b) bottom side.	133
Figure 84 – Schematic diagram of the developed test boards.	133
Figure 85 – Star and tree topologies used to measure the DR and delay.	134
Figure 86 – Lowest DR values measured for (a) raw ECG data transmission and (b) compressed ECG data transmission.	135
Figure 87 – Mean delay measured for (a) raw ECG data transmission and (b) ECG compressed ECG data transmission.	136
Figure 88 – Highest value of the delay measured for (a) raw ECG data transmission and (b) compressed ECG data transmission.	137
Figure 89 – Relative transmission sequence change due to clock drift.	141
Figure 90 – 2.4 GHz spectrum activity.	144
Figure 91 – Floor plan of the in-patient area on the second floor.	146
Figure 92 – Link survey involving all rooms accessible to patients.	148
Figure 93 – Survey used to select adequate spots, as presented on (a) through (f), to position coordinators and routers.	150

Figure 94 – Test routine user interface.	153
Figure 95 – Settings used for tests using the (a) star topology with no hidden-nodes; and (b) 2-hop tree topology with no hidden-nodes.	154
Figure 96 – The test setting inside the consultation room (the 2400E, connected to the notebook, was placed outside the room, in the corridor).....	155
Figure 97 – An ED on the head of a patient bed in room R202. The position of the device is assigned by the arrow.	155
Figure 98 – Packets captured for the IEEE 802.15.4-based networks with GTSs allocated to all EDs on (a) channel 26 and (b) channel 22.	158
Figure 99 – LQI values and lost data packets count for an ED operating in the IEEE 802.15.4 star network on channel 22.	161
Figure 100 – CDFs of DR for networks that operated on channel 26 in the absence of hidden-nodes: (a) obtained from the continuous values of the average DR and (b) obtained from the average of the DR over a 10-second sliding window.	162
Figure 101 – CDFs of average DR over a 10-second sliding window for networks that operated on channel 22 in the absence of hidden-nodes.	163
Figure 102 – DR per 2-second calculated using a 20-message length window for the 2- hop tree network on channel 26 with no hidden-nodes.	164
Figure 103 – ED1 accesses the channel just after a message from ED2 to the router is acknowledged. Unable to receive this message, the router just backs off while ED1 makes all possible retries.	165
Figure 104 – Packet timeline captured during a test using a 2-hop ZigBee-based tree network.	165
Figure 105 – Packet list that includes the packets 1 and 2 shown in Figure 104.	166
Figure 106 – Packet list that includes the packets 3 and 4 shown in Figure 104.	166
Figure 107 – 2-hop tree network.	167
Figure 108 – Setting used to evaluate the performance of a ZigBee-based, nonbeacon- enabled star network with 50% of hidden-nodes.....	169

Figure 109 – CDFs of the windowed DR for the test using a ZigBee-based, nonbeacon-enabled star network, on channel 26, 50% of hidden-nodes. ECG traffic transmission (parts 1 and 2).....	170
Figure 110 – DR per 2-second calculated using a 20-message length window for a ZigBee, nonbeacon-enabled, star network on channel 26 with 50% hidden-nodes. ECG traffic transmission.	171
Figure 111 – CDFs of windowed DR for the test using a ZigBee-based, nonbeacon-enabled star network, on channel 26, 50% of hidden-nodes. Heart rate data transmission.	172
Figure 112 – DR per 2-second calculated using a 20-message length window for a ZigBee, nonbeacon-enabled, star network on channel 26 with 50% hidden-nodes. Heart rate traffic transmission.	173
Figure 113 – DR test 1: coordinator at the entrance hall.	F-2
Figure 114 – DR test 2: coordinator at the reception.	F-3
Figure 115 – DR test 3: coordinator at the reception.	F-4
Figure 116 – DR test 4: coordinator at the pharmacy.....	F-5
Figure 117 – DR test 5: coordinator at the consultation room.	F-6
Figure 118 – DR test 6: coordinator at the consultation room.	F-6
Figure 119 – DR test 7: coordinator at the consultation room.	F-7

Tables

Table 1 – ISM bands.....	2
Table 2 – Normal temperature (°C) as a function of the body site and age (from Welch Allyn, [212]).	5
Table 3 – Normal vital sign values for various ages (from A.G. Perry and P.A. Potter [169]).	8
Table 4 – Primary ECG components and their durations [121].	11
Table 5 – PHY options introduced by the IEEE 802.15.4-2006 revision [92]......	27
Table 6 – PHY PIB attributes [91, 92].	30
Table 7 – CSMA-CA parameters and constants [91, 92]......	32
Table 8 – Some device identifiers defined by the Home Automation profile [223].	41
Table 9 – Some ZDP commands [220].	44
Table 10 – ZigBee versions feature comparison.	46
Table 11 – Basic characteristics of Wi-Fi technologies available.....	51
Table 12 – Wireless protocols that can be used to implement patient remote monitoring applications.....	54
Table 13 – QoS requirements for different categories of medical device data [87]......	62
Table 14 – Amount of data generated by sensors.....	82
Table 15 – Designed HP FIR filter coefficients.	90
Table 16 – Temperature sensor technical characteristics.	94
Table 17 – Exacon 400-series probes Steinhart-Hart constants.	95
Table 18 – Data sent to the Monitoring station application. Data are converted to string format before being passed to the method that constructs the data string sent along with the HTTP post request.	106
Table 19 – Parameters common to all simulations.....	118

Table 20 – Modes considered in the first set of simulations.	118
Table 21 – Maximum access delay as a function of the number of transmission attempts done.	123
Table 22 – Events and minimum periods of time involved in the transmission of an ECG data message for the star topology.	123
Table 23 – Events and minimum periods of time involved in the transmission of an ECG data message for the 2-hop tree topology.	124
Table 24 – Maximum delay an ECG data message can experience in a star network.	124
Table 25 – Energy consumption calculation for an ED that transmits compressed ECG data.	127
Table 26 – Maximum number of ECG sensors to achieve established QoS parameters.	127
Table 27 – Hidden-node analysis results summary for nonbeacon-enabled ZigBee/IEEE 802.15.4-based star networks.	132
Table 28 – Experimental results summary.	138
Table 29 – Experimental and simulation results comparison.	138
Table 30 – Transmission time measurements.	140
Table 31 – Wi-Fi report.	145
Table 32 – ZigBee report.	145
Table 33 – Icons used to represent network devices.	151
Table 34 – Field test results. No hidden-nodes are present.	156
Table 35 – Detailed results for field tests using star networks in the absence of hidden- nodes.	157
Table 36– Detailed results for long tests using 2-hop networks in the absence of hidden- nodes.	160
Table 37 – DR results obtained from simulation.	168

Table 38 – Average DR values for the test using a ZigBee-based, nonbeacon-enabled star network, on channel 26, 50% of hidden-nodes, ECG waveform transmission.	169
Table 39 – Detailed results for the test using a ZigBee-based nonbeacon-enabled star network on channel 26 with 50% of hidden-nodes. ECG data transmission.	170
Table 40 – Detailed results for the test using a ZigBee-based nonbeacon-enabled star network on channel 26 with 50% of hidden-nodes. Heart rate data transmission.	172
Table 41 – FlexNet (Welch Allyn) communication performance [210].	174
Table 42 – DR test 1 results.....	F-3
Table 43 – DR test 2 results.....	F-3
Table 44 – DR test 3 results.....	F-4
Table 45 – DR test 4 results.....	F-5
Table 46 - DR test 6 results.	F-6
Table 47 – DR test 7 results.....	F-7

Acronyms and abbreviations

ADC	Analog-to-digital converter
AODV	Ad hoc on Demand Distance Vector
AP	Access point
API	Application programming interface
APP	Application layer
ASK	Amplitude shift keying
AV	Atrioventricular
BAN	Body area network
BP	Blood pressure
BPSK	Binary phase-shift keying
CAD	Computer aided design
CAP	Contention access period
CDF	Cumulative distribution function
CSMA-CA	Carrier sense multiple access with collision avoidance
DR	Delivery ratio
DSP	Digital signal processor
DSSS	Direct sequence spread spectrum
DTED	Delay test end device
ED	End device
EEC	European economic area
EMR	Electronic medical record

ETCO ₂	End-tidal carbon dioxide
FFD	Full-function device
GTS	Guaranteed time slot
GUI	Graphical User Interface
HIS	Hospital information system
HM4ALL	Health Monitoring for All
HP	High pass
HPF	High pass filter
HTTP	Hypertext Transfer Protocol
IA	Instrumentation amplifier
ICU	Intensive care unit
ISM	Industrial, scientific and medical
ITU	International telecommunications union
ITU-R	ITU Radiocommunication Sector
LED	Light emitter diode
LP	Low pass
LPF	Low pass filter
LQI	Link Quality Indicator
LR-WPAN	Low-rate wireless personal area network
MAC	Medium access control
MAP	Mean arterial pressure
MGH	Massachusetts General hospital
MIT	Massachusetts Institute of Technology
NTC	Negative temperature coefficient (of resistance)
NWK	Network layer
O-QPSK	Offset quadrature phase-shift keying

PACS	Picture archiving and communication system
PAN	Personal area network
PCB	Printed circuit board
PDA	Personal digital assistant
PER	Packet error rate
PHY	Physical layer
PIB	PAN information base
PPG	Photoplethysmogram
PSSS	Parallel sequence spread spectrum
PTC	Positive temperature coefficient (of resistance)
PTT	Pulse transit time
QoS	Quality of service
RF	Radiofrequency
RFD	Reduced-function device
SA	Sinuatrial
SAP	Service access point
TSMP	Time synchronized mesh protocol
UART	Universal asynchronous receiver / transmitter
WLAN	Wireless local area network
WMAN	Wireless metropolitan area networks
WMTS	Wireless medical telemetry systems
WPAN	Wireless personal area network
WSN	Wireless sensor network
WWAN	Wireless wide area network

Chapter 1

Introduction

This chapter starts by introducing background information, which consists of basic concepts of mobile communications and topics related to vital signs, vital signs acquisition using electronic devices and vital signs monitors. These initial sections are followed by the thesis motivation, the approach taken and the list of contributions. Finally, the thesis organization is presented.

1.1 Wireless communications

Telecommunication activities are regulated worldwide by the International Telecommunications Union (ITU). A section of ITU, the ITU Radiocommunication sector (ITU-R), handles standardization in the wireless sector and is responsible for frequency spectrum management. Additionally, national and regional agencies are responsible for further regulation. Whereas several frequency bands have been allocated or licensed to particular services, such as radionavigation and terrestrial mobile telecommunication, few specific bands were assigned to industrial, scientific and medical (ISM) applications in a license-free basis [4]. These bands, which are called ISM bands, are listed in Table 1. Some of these bands, like the 2450 MHz band, are allocated worldwide, whereas others can only be used regionally. The main advantage of the ISM bands is that they can be used without license, provided that the device operating in these bands respects specific regulations. On the other hand, communication is subject to interference. Among the ISM bands, the 2450 MHz band is of special interest because it is used by wireless local area networks (WLANs) based on the IEEE 802.11 protocol, Bluetooth devices, cordless phones and baby monitors among others.

Table 1 – ISM bands.

Frequency band	Center frequency
6765–6795 kHz	6780 kHz
13,553–13,567 kHz	13,560 kHz
26,957–27,283 kHz	27,120 kHz
40.66–40.70 MHz	40.68 MHz
433.05–434.79 MHz	433.92 MHz
902–928 MHz	915 MHz
2400–2500 MHz	2450 MHz
5725–5875 MHz	5800 MHz
24–24.25 GHz	24.125 GHz
61–61.5 GHz	61.25 GHz
122–123 GHz	122.5 GHz
244–246 GHz	245 GHz

In order to exchange messages, the devices must agree on a communication protocol, which is a formal description of the digital message formats and the rules for exchanging these messages. Due to the complexity involved in the development of communication protocols, their design is structured in layers. The number of layers and their designations depend on the network type. Each layer uses services provided by lower layers and offers a set of services to the layers above it [193]. The protocol stack model used in this work consists of four layers, as shown in Figure 1, where the dark vertical arrows represent the interfaces between layers, which are called service access points (SAPs), whereas the horizontal arrows symbolize the logical information exchange between the same layers of different devices.

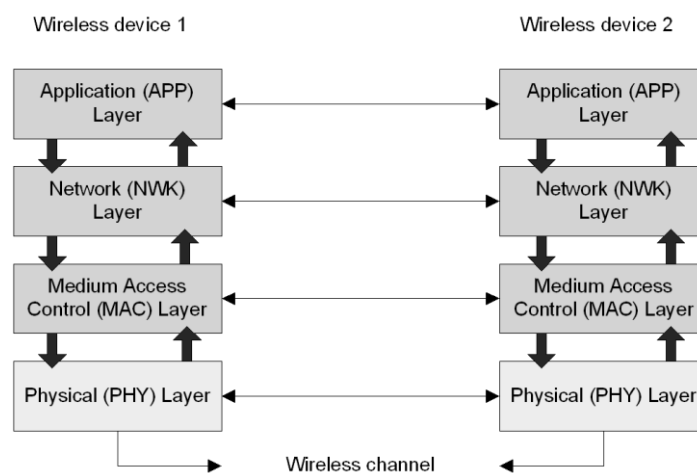


Figure 1 – Protocol stack model.

The physical layer (PHY) is responsible for data transmission and reception. Additionally, this layer is responsible for all activities directly involving the radio, such

as frequency selection, generation of the carrier frequency, signal detection and modulation. The medium access control layer (MAC) arbitrates the access to a shared wireless medium. The network layer (NWK) is responsible for routing frames to their intended destinations and provides functionalities such as network formation, address assignment, and mechanisms for devices to join and leave the network. The application layer (APP) is the highest hierarchy layer and provides an interface to applications objects [220].

Wireless networks are subject to issues intrinsic to the wireless channel, such as fading and interference. Unintentional noise consists of interference from other emitters that share the same frequency band and broadband interference from incidental radio frequency emitters (from instance, radiofrequency interference generated by electric power transmission lines [37]). Fading can occur due to multipath propagation, referred to as multipath induced fading, or due to shadowing from obstacles that affect the radio waves propagation. Multipath induced fading occurs in the presence of surfaces that reflect or cause scattering or diffraction of the radio waves. Multiple copies of the signal experience attenuation, delay and phase shift, which can result in constructive or destructive interference in the receiver. When severe destructive interference occurs, it is referred to as deep fade and results in loss of communication [182].

The quality of the service provided by wireless communication networks can be evaluated using quality of service (QoS) parameters such as packet delivery ratio, delay, packet delay variation, and throughput [58]. The packet delivery ratio represents the percentage of successfully delivered packets to the number of generated packets. Delay or latency refers to the time interval between the instant when a packet is generated and when it is delivered. The delay suffered by a packet can vary depending on several causes, such as the route taken and the amount of traffic in the network. This variation in delay is known as the packet delay variation or jitter. Throughput refers to the average rate of successful message delivery over a communication channel. If all emitters generate the same amount of traffic, ideally, the network throughput increases linearly with the number of emitters. However, in practice, when the network traffic exceeds a certain limit, the network performance degrades and, in the limit, the network may collapse.

Several applications require that certain QoS requirements are ensured. For instance, remote patient monitoring applications typically have low bandwidth requirements, but are highly sensitive to delay and data loss. On the other hand, file transfer applications can tolerate relatively large packet delays, but are extremely sensitive to data loss [188]. Whereas some wireless communication protocols can only offer a single best effort level of service, others can support different levels of QoS requirements. A network must be able to prioritize traffic and allocate specific resources to each traffic category to provide QoS. In this case, QoS parameters can be probabilistically bounded. Conversely, networks that only offer best effort services cannot support QoS. However, high quality communication can be provided by these networks by over-provisioning their capacity so that it is sufficient for the expected peak traffic load.

1.2 Vital signs

According to the Merriam-Webster's English Dictionary, vital signs are defined as signs of life. [144]. However, for health care providers vital signs are more than indicatives of life. Actually, they are used for diagnostic and evaluation purposes because they convey essential information about the body functioning. The primary vital signs are the body temperature, the heart rate, the blood pressure, and the respiration rate [169].

The body temperature is regulated by the hypothalamus, a portion of the brain [212]. Adults, to a higher degree than infants and children, are capable of maintaining their core temperature (the temperature of the deep tissues and organs within the cranial, thoracic and abdominal cavities), regardless of the external temperature. The skin temperature, however, is heavily dependent on the environmental temperature and the site where it is measured, not being directly correlated to the core temperature. It may be quite a few degrees below or above the core temperature without affecting the correct functioning of the body [169]. Nevertheless, a few accessible sites maintain a temperature that can be correlated to the core temperature. It is the case of the esophagus, the rectum, the mouth, the tympanic membrane and the axilla. Whereas measurements done using esophageal (mainly used during sedation or general

anesthesia) and rectal probes can deliver the best results, they are the most intrusive and are used only when very accurate readings are necessary.

The ear (aural temperature) is a very adequate site, but its accuracy depends on the operator completely occluding the aural membrane. The mouth should not be used in non-cooperative patients whereas the axilla requires a good skin contact, being less accurate if used in obese or very thin patients [85].

Table 2 presents the normal temperature ranges as a function of the site it is measured [212]. As expected, the core and rectal temperature ranges are approximately coincident. The oral temperature in healthy children can be as low as 35.5 °C. The axillary temperature, most of the times, is lower than the others and, in healthy adults, should not reach 37 °C.

Table 2 – Normal temperature (°C) as a function of the body site and age (from Welch Allyn, [212]).

Body site/Age	0 - 2 years	3 -10 years	11 – 65 years	> 65 years
Core	36.6 – 37.8	36.4 – 37.8	36.8 – 37.9	35.9 – 37.1
Rectal	36.6 – 38.0	36.6 – 38.0	37.0 – 38.1	36.2 – 37.3
Ear	36.4 – 38.0	36.1 – 37.8	35.9 – 37.6	35.8 – 37.5
Oral	---	35.5 – 37.5	36.4 – 37.6	35.8 – 36.9
Axillary	34.7 – 37.3	35.9 – 36.7	35.2 – 36.9	35.6 – 36.3

Since the heart rate can be estimated directly by artery palpation or using a stethoscope, it can be precisely measured using an electrocardiograph, an instrument that records the electrical activity of the heart, that is, the electrocardiogram (ECG). An adult resting heart rate ranges from 60 – 100 bpm (beats per minute) though trained athletes tend to have lower heart rates (< 60 bpm). On the hand, healthy teenagers, children and infants may have higher heart rates (> 100 bpm). Arrhythmias are abnormalities of the heart rate. During an arrhythmia the heart can beat very fast, very slow, or assume an irregular rhythm. On adults, a heart rate above 100 bpm is referred to as tachycardia, whereas heart rates below 60 bpm are referred to as bradycardia [155]. The heart rate measurement using the ECG will be discussed in the next section.

The blood pressure (BP) is the pressure applied on the artery walls as blood is forced through them by the heart [85]. Two pressures are measured. The higher or systolic is the pressure applied to the arteries while the heart contracts to pump blood to the body. The lower or diastolic occurs when the heart relaxes between beats [10]. A

BP of 120/80 mmHg is considered optimal for adults. An adult individual is considered to have a high blood pressure if the systolic BP is higher or equal to 140 mmHg or if the diastolic BP is over 90 mm Hg [156].

The BP can be measured using invasive and non-invasive methods. Invasive methods involve penetrating an artery wall and are restricted to the hospital environment. The auscultatory and oscillometric methods are the most common non-invasive methods. The first one requires obstructing the brachial artery (in the upper arm) with an inflated cuff. When the pressure is released from the cuff and the blood starts to pass through the artery, the turbulence provoked by the cuff can be heard (Korotkoff sound) by an operator using a stethoscope placed at the elbow, over the artery. The pressure applied to the cuff at this moment corresponds to the systolic BP. If the operator continues to release the cuff pressure, the sound will eventually disappear. The pressure applied by the cuff at this moment corresponds to the diastolic BP [175]. Devices based on oscillometry measure the oscillations (pressure variation in the cuff) caused by the blood flow during cuff deflation. The pressure at which the oscillations are maximal is defined as mean arterial pressure (MAP). Systolic and diastolic BP values are estimated from the successive pressure measurements made by the pressure sensor [119]. Other non-invasive methods exist, but all commercial solutions use occlusive cuffs.

Occlusive cuff-based methods are not adequate for continuous long-term BP measurements. Over the time, they get uncomfortable and may produce unreliable readings [185]. However, non-invasive continuous long-term BP measurements can be beneficial in many settings including surgical operation, intensive care units, and home health care [30]. Cuffless continuous BP measurement is a very active research area. The majority of the systems proposed so far are based on simultaneously processing two biosignals: ECG and photoplethysmogram (PPG) to determine the pulse transit time (PTT)¹ [28, 194]; ECG and impedance cardiogram² [217]; or PPG and the pressure

¹ The PPG reflects the changes in blood perfusion (delivery) in limbs and tissues. The PTT is the time it takes the pressure pulse to transfer from the aorta to the peripheral arteries.

² The impedance cardiogram is the derivative of the thorax impedance involving the measurement of cardiac-related impedance change.

measured in a ring positioned around a finger [185]. Although promising, these methods possess serious limitations: the PPG signal is easily distorted by motion artifacts; the PTT evaluation loses accuracy when the extremity (finger, earlobe) used to measure the PPG signal is exposed to cold; and all require frequent calibrations [184].

The respiration rate is the number of breaths a living subject takes within a certain amount of time. Oxygen (O_2) and carbon dioxide (CO_2) are the two gases exchanged during respiration. An increased CO_2 level in the blood is the primary stimulus for breathing [85]. Besides assessing if the subject is breathing, the respiration rate and pattern are important diagnostic tools that can be used to treat several illnesses such as sleep apnea syndrome, lower respiratory tract infections, chronic obstructive pulmonary disease, tuberculosis, and lung cancer [148]. Moreover, it is an important tool in the detection of life threatening situations for bedridden patients [65].

Invasive methods are exclusively used in the hospital environment and use probes that sense the air flow. The most common non-invasive methods are the impedance pneumography and the inductive plethysmography. Impedance pneumographic devices measure the voltage drop across thoracic electrodes, which increases during inspiration and decreases during expiration. Inductive plethysmography employs arrays of sinusoidally arranged copper wires excited by a low-current, high-frequency electrical oscillator circuit. Movement of the thorax causes variations in the magnetic fields, which are measured as voltage changes over time. Other methods for assessing the respiration rate include non-contact approaches. For example, EMFIT, a company from Finland, produces a non-contact monitoring system capable of detecting heart beating and breathing rate of in-bed patients. It is based on elastic, permanently charged ferro-electret film that converts mechanical stress into proportionate electrical energy [50]. Other non-contact technologies under development include the use of reflection radio waves, capacitive and optical sensors, microwave Doppler radar and thermal image processing [5, 120, 162, 180, 207].

Table 3 contains normal vital sign average and range values for various ages [169].

Table 3 – Normal vital sign values for various ages (from A.G. Perry and P.A. Potter [169]).

Age	Temperature Average (°C)	Heart Rate Average and Range (beats per min)	Blood Pressure Average (mm Hg)	Respiration Rate Range (breaths per min)
Newborn	36.8 (axillary)	120 (70 – 170)	80/40	40 – 90
1 – 3 years	37.7 (rectal)	110 (80 – 130)	98/64	20 – 40
6 – 8 years	37.0 (oral)	95 (70 – 110)	102/56	20 – 25
10 years	37.0 (oral)	90 (70 – 100)	110/58	17 – 22
Teen	37.0 (oral)	80 (55 – 105)	110/70	15 – 20
Adult	37.0 (oral)	80 (60 – 100)	<120/80	12 – 20
> 70 years	36.0 (oral)	80 (60 – 100)	Up to 160/95	12 – 20

1.3 Acquisition of vital signs

The energy produced by physiological processes, including heat, electrical, mechanical, or chemical energies, is measured for acquiring a vital sign. Some physiological processes involve electrical activity that can be directly read using electrodes attached to the surface of the body. Alternatively, a transducer must be used when non-electrical biological parameters are measured using electronic devices. An adequate transducer senses the presence, magnitude or variation of a specific form of energy produced by a physiological process and provides an electrical output that can be accurately correlated to the form of energy sensed. The transducer selection is critical because it significantly contributes to the precision, cost, invasiveness, and complexity of the system [183]. Examples of transducers include thermistors and thermocouples, used to measure temperature; strain gauges, used to measure pressure and force; and linear variable differential transformers, which are commonly used to measure displacement.

1.3.1 ECG acquisition

Bioelectrical signals originate in the ionic voltages produced by groups of excitable cells. When stimulated, a cardiac muscle cell goes through a cycle called action potential³. In its resting state, this cell is electrically negative with respect to the outside. However, when stimulated, its membrane turns more permeable to sodium (Na^+) ions, which results in a rapid influx of Na^+ ions. During this period, called depolarization, the

³ Apart from cardiac muscle cells, action potentials occur in several types of cells, which include neurons, other muscle cells, and endocrine cells.

electrical polarization of the cell in relation to the outside rises. When the cell reaches the depolarized state, the cell membrane begins to return to its original permeability. However, rather than Na^+ ions flowing out the cell, K^+ ions diffuse out, causing the cell to contract. At the end of this cycle, the cell returns to its resting state. This flux of ions during an action potential can be observed as potentials on the body surface [114, 175].

A stylized ECG waveform is shown in Figure 2. The ECG is recorded to continuously measure the heart rate, detect abnormal rhythms and identify if the heart muscle has been damaged in specific areas. The baseline of the ECG tracing is called the isoelectric line and denotes resting membrane potentials. Deflections from this point are lettered in alphabetical order; from P to U. The typical amplitude of the R wave is equal to 1 mV.




Figure 2 – ECG waveform (adapted from [25]).

The electrical conduction system of the human heart is depicted in Figure 3 [157]. During diastole, the atria are filled with blood. By the end of diastole, an electrical signal is generated by the sinuatrial (SA) node, located in the top of the right atrium. This signal is generated by specialized cells that have a property called automaticity, which reflects an ability to initiate electrical impulses spontaneously [18]. The signal spreads to the right and left atria causing them to contract and pump blood to the ventricles. This action pushes blood through the open valves from the atria into both ventricles. In the ECG waveform shown in Figure 2, this signal is observed as the P wave [157].

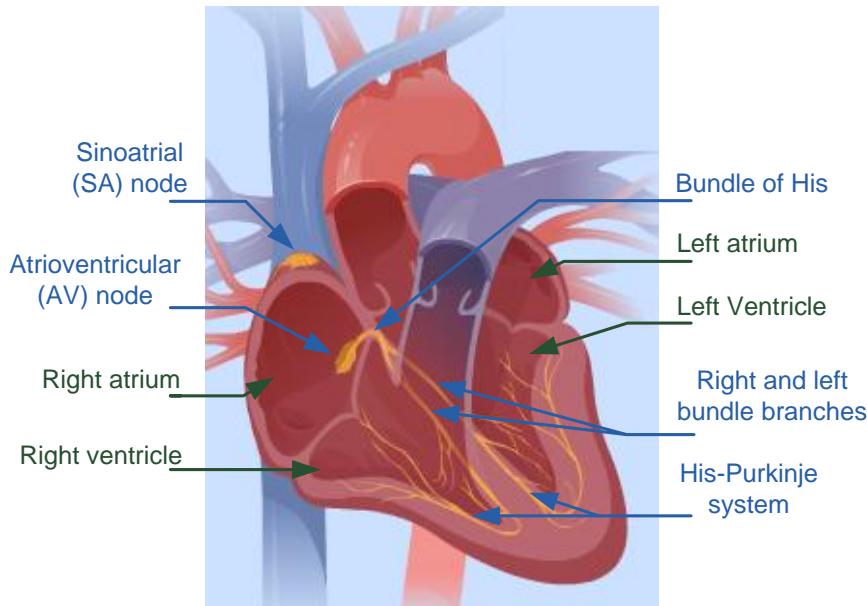


Figure 3 – Electrical conduction system of the heart (adapted from [157]).

After that, the electrical signal arrives at the atrioventricular (AV) node. At this point, it is slowed down to let the ventricles fill with blood. This period of time corresponds to the interval between the end of the P wave and the beginning of the Q wave. After that, the electrical signal moves through the bundle of His and divides on the right and left bundle branches until it reaches the bottom of the ventricles. In the ECG it corresponds to the Q wave. Then, the electrical signal spreads quickly across the His-Purkinje system, a set of nerve cells distributed across the ventricle walls, causing the ventricles to contract. The left ventricle, however, contracts a few milliseconds before the right ventricle. In the ECG, the R wave marks the contraction of the left ventricle whereas the S wave marks the contraction of the right ventricle. As the stimuli ceases, the ventricles relax. On the ECG, the T wave marks the point at which the ventricles are relaxing. The U wave most likely represents an electromechanical phenomenon that occurs after repolarization is completed. However, this low-amplitude low-frequency wave is frequently absent in the limb leads [176]. Finally, the heart muscle is relaxed and, in a few instants, the process is again repeated.

Table 4 shows a summary of the primary ECG waveform components, the electrical activity related to each one and their normal duration. The T wave and the ST segment durations are not listed because their durations are not commonly measured [121].

Table 4 – Primary ECG components and their durations [121].

Primary ECG component	Electrical activity related	Normal duration (ms)
P wave	Atrial depolarization	80 – 100
QRS complex	Ventricular depolarization	60 – 100
T wave	Ventricular repolarization	---
P-R interval	Atrial depolarization and AV nodal delay	120 – 200
ST segment	Isoelectric period of depolarized ventricles	---
Q-T interval	Action potential duration	200 – 400

A simplified block diagram of a single channel ECG sensor is shown in Figure 4. The instrumentation amplifier (IA) amplifies the potential difference between the leads applied to its inputs (in this case, the RA and LA leads) while rejecting large values of common mode noise. The low pass filter (LPF) and high pass filters (HPF) follow the IA and are used to attenuate unwanted frequency components. Most electrocardiographs can operate either on diagnostic or on monitoring filter mode. The diagnostic filter mode allows physicians to observe time relations that can be altered if the ECG signal bandwidth is limited. Generally, several leads are used, which allows the exploration of more than one reference plane. In contrast, the monitoring mode is typically used to detect abnormal heart rhythms (arrhythmias). In this case, just one lead (usually lead II, I or one of the chest leads) may be sufficient and several frequency components can be attenuated in order to limit artifacts. In diagnostic mode, the high pass filters are set at 0.05 Hz whereas the low pass filters are set at 40 Hz, 150 Hz or higher (up to 1 kHz). In monitoring mode, the high pass filters are set at either 0.5 Hz or 1 Hz whereas the low-pass filters are set at 40 Hz [138]. A 50 Hz notch filter can be used to attenuate the interference from the electrical power system that was not rejected by the IA.

Then, an analog-to-digital converter (ADC) samples and converts the analog signal into discrete samples. Generally, ADCs with resolution equal or greater than 12 bits are used. The digital signal from the ADC is then processed by a microcontroller, a microprocessor or a digital signal processor (DSP). The digital processing done may include signal conditioning, heart rate evaluation, arrhythmia detection, ECG components measurement (amplitude and duration), ectopic beats detection and episodes onset (ischemia) detection and classification. Depending on the electrocardiograph, the processed ECG signal may be displayed or printed. Additionally, it may be sent to a monitoring station or to the hospital information system (HIS) using a wired connection or a wireless channel.




Figure 4 – Simplified single-channel ECG.

Apart from the components already presented, electrocardiography devices include circuits that provide electrical isolation from the mains power and defibrillation protection. The electrical isolation limits the possibility of the passage of any leakage current from the instrument to the patient. Specifically, leakage currents must be limited below the safety standard limit of $10 \mu\text{A}$ [199]. Wearable devices powered by low-power batteries may not require an isolation circuit, but it is mandatory for all main powered units. Additionally, electrocardiographs must be protected against very high voltages from electric defibrillators that can damage the instrument.

Although ECG acquisition is well-developed and commercial solutions are available, the design of ECG front-ends offer several challenges due to the presence of the large DC offset and various interference signals. Whereas ECG signals have amplitudes around 1 mV , electrical fields can typically create potentials around 20 mV , too intense to be rejected by the IA. Additionally, unbalance in the electrode-skin interface impedance creates higher common mode potentials at one input than in the other. Hence, this difference is seen by the IA as differential voltage and is amplified causing distortions in the ECG waveform. Motion artifacts can saturate the amplifiers and muscle signals (EMG potentials) can appear as interference. Additionally, magnetic induced voltages can be formed in the lead wires producing additional interference [208]. Finally, the design of small wearable devices, such as the ECG sensor developed under this work, involves positioning electrodes in close proximity. In this situation, the ECG waveform amplitude is considerably reduced, which may force the addition of a subsequent amplification stage that can introduce noise.

A prototype of the developed ECG sensor is shown in Figure 5. As shown, no lead wires are necessary since disposable electrodes are connected directly to sensor connectors. It continuously measures a modified projection of the bipolar lead I vector⁴, evaluates heart rate, detects abnormal rhythms, namely tachycardia, bradycardia, asystole and background arrhythmia⁵, and sends data through the wireless channel.




Figure 5 – The wearable ECG sensor developed: (a) printed circuit board front side; (b) printed circuit board rear side; and (c) user case top side.

1.3.2 Body temperature measurement

Body temperature is typically measured using thermoresistive temperature transducers. These components exhibit a change in resistance with a change in the temperature they are submitted to. Thermistors are the most commonly used transducer in medical applications. They are made of ceramic semiconductors (metal oxides) and can either have a positive temperature coefficient of resistance (PTC) or a negative temperature coefficient of resistance (NTC). In PTC devices, an increase in temperature corresponds to an increase in resistance, whereas NTC devices behave the opposite way. The PTC is best suited for switching applications, whereas the NTC is used for

⁴ The ECG sensor uses non-standard electrode positioning on the torso. The resulting waveform is adequate for heart rate measurement, but cannot be used for diagnostic purposes.

⁵ Asystole refers to the state of no cardiac electrical activity. A background arrhythmia (variable rhythm) occurs if in one minute period, the sensor detects an R-R interval that corresponds to a heart rate 20% higher or 20% lower than the average R-R interval followed by another R-R interval that corresponds, respectively, to a heart rate 20% lower or higher than the average R-R interval.

precision temperature measurements, including biomedical instrumentation. Thermistors can be manufactured in a wide variety of configurations and protective coatings.

As will be further discussed in Chapter 4, thermistors exhibit a non-linear resistance versus temperature (R-T) characteristic. Moreover, each component has a particular characteristic, so calibration is required when using thermistors. Vital signs monitoring devices use pre-calibrated medical probes which are based on NTC thermistors that possess well-defined R-T characteristics. This ensures interchangeability, repeatability and accuracy equal or better than ± 0.1 °C between 32 °C and 42 °C [218]. Besides, the thermistor excitation current level must be limited to avoid self-heating, which causes the thermistor resistance to decrease and give a higher reading than the actual temperature.

Figure 6 shows the simplified block diagram of a wireless body temperature sensor based on an interchangeable thermistor. The thermistor excitation is supplied by a constant current source. The resistor in series with the thermistor is used to increase the voltage on the thermistor terminal that is connected to the difference amplifier (AMP). The analog voltage supplied by the amplifier is then converted into a digital value by an ADC. Then, this digital value is read by the microcontroller and the temperature of the thermistor is estimated. Periodically, temperature values are sent through the wireless channel.




Figure 6 – Simplified block diagram of a thermistor-based temperature sensor.

A prototype of the body temperature sensor developed in this work is shown in Figure 7. As shown, an armband holds the sensor to the upper arm and a temperature probe extends from the sensor to the axillary region.




Figure 7 – The temperature sensor prototype version.

1.4 Vital signs monitors

Vital signs monitors can be classified into two broad categories: bedside monitors and biotelemetry systems. Whereas bedside monitors are capable of accurately acquiring several vital signs from bedridden patients, biotelemetry systems provide the opportunity of continuously monitoring a patient without restricting his or her mobility. Moreover, biotelemetry systems can be used to monitor out-patients who can benefit from continuous monitoring during recovery or throughout their normal activities.

1.4.1 Bedside monitors

Presently, a variety of monitoring systems are used in hospitals, where bedside and portable monitors are the most common ones. Bedside monitors as the 1500 Patient Monitor, from Welch Allyn (Figure 8), are the most expensive, biggest and highest performance monitors. They are mainly used during anesthesia and in intensive care units (ICUs) to monitor several vital signs, including heart rate, respiratory rate, blood pressure and temperature. In addition, special functions such as capnography, oximetry, electroencephalography and pulmonary artery catheter measurements can be done [213]. Most of these monitors allow data to be transferred to monitoring stations or to specific electronic medical record (EMR) database applications. In specific cases, data can be transmitted using a wireless channel.

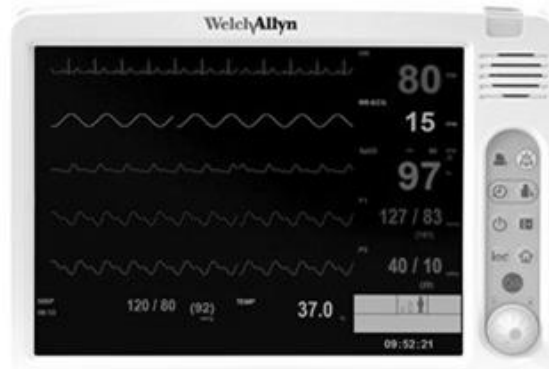


Figure 8 – The 1500 Patient Monitor (adapted from [213]).

Portable bedside monitors, such as the Dinamap Carescape V100 [72], shown in Figure 9, are battery powered units that can be carried from one patient to the other, allowing nurses to quickly check several patients' vital signs or used as a bedside monitor to continuously monitor patients. In general, they are smaller and simpler than bedside monitors, but some of these units can also transfer vital sign data.



Figure 9 – The Dinamap Carescape V100 vital signs monitor (adapted from [72]).

1.4.2 Biotelemetry systems

Biotelemetry systems are used in hospitals to continuously monitor patients whose recovery benefits from mobility. These systems are comprised of portable patient units, wireless infrastructure and a central monitoring station. The biotelemetry systems manufactured by the major health care device manufacturers have many resemblances and their portable devices have similar form factors. Figure 10 (a) shows a Micropaq patient unit worn by a patient in the hospital environment and Figure 10 (b) shows a snapshot of its user interface and external features. This system, which is manufactured by Welch Allyn, is capable of measuring the heart activity and the oxygen saturation

level (SpO₂) [211]. As can be observed, despite being portable, it is relatively large and obtrusive.



Figure 10 – Micropaq patient unit: a) A patient unit worn in the hospital environment (adapted from [210]); b) User interface snapshot and external appearance (reproduced courtesy of Welch Allyn).

Just recently, biotelemetry manufacturers began using standard-based wireless communication protocols, which may become a trend [17]. In 2005, Welch Allyn biotelemetry systems started to use the IEEE 802.11 protocol. On the same year, Philips Healthcare announced a biotelemetry system designed to work in the 2.4 GHz ISM band but, only recently, a new version based on the IEEE 802.11 protocol was released. On the other hand, some manufacturers still adopt custom-designed communication technologies. GE Healthcare's biotelemetry solution, Apex Pro, operates in the 420-460 MHz radio band and use a proprietary protocol [70].

1.5 Motivation

According to the Sixth Annual Health Grades Patient Safety in American Hospitals Study, failure to rescue⁶ events are among the medical errors of highest incidence rates in American hospitals, accounting for a total of 14,903 deaths between 2005 and 2007 [84]. Possibly, some of these events could have been prevented if early signs of degradation in patients' conditions were available to health care providers in due time.

⁶ According to the American Agency for Health care Research and Quality (<http://www.psnet.ahrq.gov/glossary.aspx>), "failure to rescue" refers to the inability of conducting actions that could have prevented a clinically important deterioration, such as death or permanent disability, from a complication of an underlying illness or a complication of medical care that developed on their watch. It may reflect the quality of monitoring, the effectiveness of actions taken once early complications are recognized, or both.

The assessment of patients' conditions generally includes vital signs monitoring since they convey important information about a patient's condition and response to treatments. Ideally, all hospital patients should have one or more vital signs continually monitored to reduce the care response delay in case of an adverse event [59]. Still, this scenario is the opposite of what presently occurs in most hospitals where typically few patients out of the intensive and intermediate care units have their vital signs continuously monitored. In case an emergency situation occurs, it is possible that several minutes, or even hours, elapse before care is delivered, which can have serious consequences to patients.

Out-patients and elderly can also take advantage from remote vital signs monitoring. In addition to avoiding some trips to hospitals and other health care centers, care can be greatly improved by maintaining accessible files that contain accurate and updated vital signs recordings. Early detection of problems results in early intervention, which can prevent deterioration and improve quality of life often through relatively minor, inexpensive interventions such as a change in lifestyle⁷. Possible settings include domestic, assisted living facilities, nursing homes, and mobile monitoring. In all these settings, wireless communication technologies such as wireless personal networks (WPANs), wireless local area networks (WLANs), cellular and GPS technologies can play a fundamental role. Nevertheless, out-patients and elderly lack devices capable of continuously monitoring their health. Actually, there are not many commercial devices tailored for this purpose, although some exist, such as the Ericsson Mobile Health [51].

The benefits depicted in the above scenarios, in conjunction with the emergence of wireless sensor network (WSN) technologies and low-data rate standard communication

⁷ The impact of ageing in health care systems is known for several years. In 1982, the United Nations endorsed the Vienna International Plan of Action on Ageing. It acknowledged the increasing population ageing and the related trend towards rising costs in health care. In 2003, the Centers for Disease Control and Prevention (CDC) published a study where it estimated a high increase in the health care costs due to the increase of the number of persons aged ≥ 65 years (available at <http://www.cdc.gov/mmwr/preview/mmwrhtml/mm5206a2.htm>). It also warned about the sustainability of the health care systems assuming the projected growth in the elderly support ratio (i.e., the number of persons aged ≥ 65 years per 100 persons aged 20-64 years).

protocols, namely the IEEE 802.15.4 and ZigBee protocols, provided the motivation for developing this work.

1.6 Approach and contributions

The purpose of this research is to investigate the feasibility of employing WSN technologies to monitor non-critical in-patients and chronically-ill out-patients, and determine scenarios under which those technologies can be deployed. The focus of this research is on the evaluation of the performance of a vital signs monitoring system based on the IEEE 802.15.4 and ZigBee protocols to remotely monitor patients.

The choice of the communication technologies above mentioned is based on the several benefits they can offer, such as:

- Standard protocol stack implementations are reliable and can considerably reduce the development costs;
- Standard-based radios and integrated communication modules are cheaper than customized solutions;
- In case an open health care profile is employed, medical sensors from a variety of manufacturers can be used in the same project without modification [62].

When this work started, on October 2007, few standard wireless protocols were available, such as IEEE 802.11 and Bluetooth. Standard low data rate protocols and WSN were still emerging technologies. The power consumption requirements of IEEE 802.11-based communication modules prevented their use in wearable devices; whereas Bluetooth had constraints in terms of topology, scalability and range. These restrictions made both protocols unsuitable options for the purpose of this work [22, 90]. Alternatively, low data rate WPAN standard protocols, such as IEEE 802.15.4 and ZigBee, were promising candidates because both were designed to address the need of low-cost, small footprint and low-power wireless devices and provided security services [91]. Moreover, ZigBee provided extended range network range through the use of routers [220] and means for device interoperability through a standard health care profile.

Whereas these protocols have appealing characteristics, it is not clear if they are able to satisfy the QoS requirements of health monitoring applications. The main difficulty arises from the fact that some sensors must be sampled quite often, generating a large amount of data and, consequently, requiring the network to operate under high load, which is not common in typical WSN scenarios.

Consequently, the following research questions apply:

- How stakeholders assess available vital signs monitoring systems and what can be done to improve these systems?
- To which health monitoring scenarios can WSN technologies be applied and what relative benefits can be accomplished?
- Are WSN systems based on the IEEE 802.15.4 and ZigBee protocols capable of guaranteeing the QoS requirements typical of vital signs monitoring systems?
- Which mechanisms provided by the IEEE 802.15.4 and ZigBee protocols should be used to optimize the proposed application performance?

The contributions of this research are listed as follows:

- A review of present remote vital signs monitoring systems followed by suggestions to increase their acceptability and performance;
- An overview of WSNs based on the IEEE 802.15.4 and ZigBee protocols and their applications in health and wellness monitoring;
- The prototyping of HM4All, which stands for Health Monitoring for All, a remote monitoring system which consists of one or more WSNs that gather vital signs from patients and send this data to monitoring stations through gateways. In contrast with the commercial monitoring systems available, the ZigBee protocol is used which, despite the limitations in terms of data rate, promises great reduction in power consumption, footprint, and cost, with direct consequences in the sensors' power autonomy, size and cost.
- Development of test programs used to evaluate HM4All's performance;
- Performance evaluation of HM4All based on simulation, laboratory tests and on-site, in a hospital scenario;
- The development of wearable ECG and axillary temperature sensors, including the hardware and the software.

Figure 11 presents a monitoring window displayed by HM4All. It shows vital signs, namely ECG waveform, heart rate and axillary temperature, and sensor battery levels. HM4All is introduced in Chapter 4.

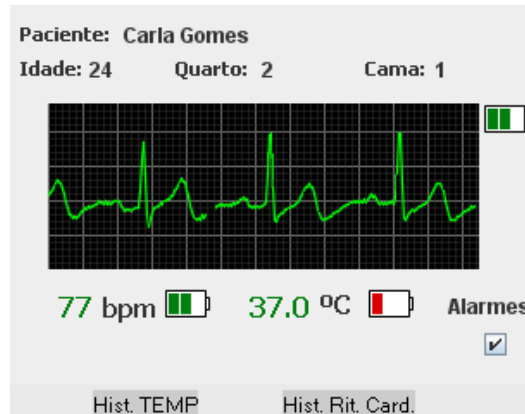


Figure 11 – Individual patient monitoring window.

The contributions described above have been succinctly presented in international conferences [60-64].

1.7 Thesis organization

This thesis is divided into seven chapters and six appendixes, which are described as follows:

Chapter 2 introduces important wireless protocols and describes the main concepts on WSNs. Special focus is given to the protocols used to develop the system prototype, IEEE 802.15.4 and ZigBee.

Chapter 3 reviews state-of-the-art research on vital signs monitoring. In addition, it briefly presents some commercial remote patient monitoring systems and enumerates general issues reported by health care providers and engineers, as well as suggestions to increase their acceptability and performance.

Chapter 4 describes the wireless monitoring system prototyped, including the developed hardware and software.

Chapter 5 analyzes and discusses the performance expected by vital signs monitoring systems based on the IEEE 8.2.15.4 and ZigBee protocols based on simulations and laboratory tests.

Chapter 6 analyzes and discusses the prototype system's performance based on experimental tests performed in the hospital environment.

Chapter 7 includes concluding remarks and indicates possible future directions for this research topic.

Appendix A contains the schematic diagram of the electrocardiogram sensor's board.

Appendix B contains the schematic diagram of the temperature sensor's board.

Appendix C presents the calculation performed to estimate the temperature sensor battery lifetime.

Appendix D contains the schematic diagram of coordinators' and routers' board.

Appendix E includes the questionnaires used to evaluate the developed system user acceptance and obtain critics and suggestions for future developments.

Appendix F describes the results of range tests executed in an in-patient floor of Hospital Privado de Guimarães.

Chapter 2

Wireless sensor networks and protocols

Initially, this chapter introduces fundamental concepts of wireless sensor networks. Then, it describes the protocols used to develop this work, IEEE 802.15.4 and ZigBee. Finally, it overviews other wireless protocols that can be used to remotely monitor patients' health and compares these protocols with IEEE 802.15.4 and ZigBee in different health monitoring scenarios.

2.1 Wireless sensor networks: definition and applications

Wireless sensor networks (WSNs) are comprised of a large number of spatially distributed devices with sensing, processing and radio communication capabilities. The main features of these devices include [96, 118]:

- Self-organizing capabilities and multi-hopping routing;
- Topology adaptation in response to changes in propagation conditions or node failures;
- Small size and low cost;
- Limitations in power consumption and processing capabilities.

Several physical parameters can be measured with small and low consumption sensors capable of being integrated in sensing devices. These include, among several others, temperature, humidity, pressure, touch, acceleration, motion and vibration. Additionally, devices can contain actuators, such as mechanical switches and piezoelectric actuators.

Numerous protocols have been designed to support different requirements applied to WSNs [118, 129, 192]. From an application point of view, protocols can aim, for instance, to minimize the energy consumption of individual nodes, to maximize the network lifetime, or to provide real-time quality of service (QoS) guarantees. From a layered view, medium access control (MAC) protocols provide channel access guarantees, whereas in the network layer, protocols provide end-to-end (multi-hop) communication services including routing through intermediate nodes. Some protocols can cover more than one or all protocol layers including the support to applications, which lay in the highest hierarchy layer [129].

This work is based on the IEEE 802.15.4 and the ZigBee protocols, two standard-based communication protocols described in the next sections. These protocols were designed to address the basic requirements of WSNs and have been implemented by several chip manufacturers, like Texas Instruments, Ember or Jennic. The adoption of one of these protocols can considerably reduce the final cost of a product. Additionally, in case of the ZigBee protocol, it is possible to create interoperable products by following an application profile published by the ZigBee Alliance.

2.2 The IEEE 802.15.4 protocol

The first version of the IEEE 802.15.4 protocol was published in 2003. Since then, it has undergone a revision in 2006 and three amendments had been approved. Additionally, at the time this work is being written, three task groups are active.

The IEEE 802.15.4 protocol is the basis of open-standard based protocols, such as ZigBee [220], RF4CE [8], and WirelessHART [82], as well as other proprietary network stacks [15, 108]. In contrast, simple applications may be built directly on top of it.

2.2.1 IEEE 802.15.4 protocol overview

The IEEE 802.15.4 protocol defines the PHY layer and MAC sublayer specifications for low-rate (limited to 250 kbps) wireless personal area networks (LR-WPANs). It is a low-complexity and flexible protocol targeted to provide wireless connectivity to resource-constrained devices.

Two types of devices can participate in an IEEE 802.15.4-based network: a full-function device (FFD) and a reduced-function device (RFD). FFDs implement the complete protocol set and can act as coordinators. RFDs implement part of the protocol and, consequently, can be implemented on simpler devices. FFDs can communicate with other FFDs and with RFDs, whereas RFDs, as simpler devices, can talk only to its FFD parent. Additionally, the protocol defines that entities participating in a network can assume the following possible roles [91, 92]:

- Personal area network (PAN) coordinator: A FFD that is the principal controller of a PAN.
- Coordinator: A FFD that provides synchronization services through the transmission of beacons.
- Alternate PAN coordinator: A coordinator that is capable of replacing the PAN coordinator. A PAN can have zero or more alternate PAN coordinators.
- Device: Any entity (FFD or a RFD) that contain an implementation of the IEEE 802.15.4 protocol.

The protocol supports the star and peer-to-peer topologies depicted in Figure 12. These topologies contain only one coordinator, the PAN coordinator. In the star topology, all devices are connected to the PAN coordinator, whereas in the peer-to-peer topology each device is capable of communicating with any other device within its radio sphere of influence. Additionally, the standard mentions the cluster-tree topology shown in Figure 13 as an example of a peer-to-peer network. In this topology, any of the FFDs may act as a coordinator or cluster head. A RFD can join any of the FFDs, which shall act as a coordinator to provide synchronization services. As routing is implemented as part of the network layer, the IEEE 802.15.4 protocol alone cannot support this topology.




Figure 12 – Star (left) and peer-to-peer (right) topologies [91, 92].




Figure 13 – The cluster-tree topology [91, 92].

Four data frames structures are defined by the IEEE 802.15.4 protocol: a) a data frame used to transfer data between devices; b) an acknowledgement frame used to confirm the successful reception of a data frame; c) a beacon frame used by the coordinator to synchronize devices and disseminate information; and d) MAC control frames [91, 92].

Devices contend to access the wireless channel using the carrier sense multiple access with a collision avoidance medium access mechanism. If an optional superframe structure is used, time slots can be allocated to low-latency applications or applications requiring specific data bandwidth [91, 92]. A superframe is limited by successive beacon frames transmitted by the coordinator, as shown in Figure 14. The period between successive beacons is called Beacon Interval (BI). It includes the Superframe Duration (SD), a period used by the coordinator to interact with the network devices, and the Inactive Period in which the coordinator switches off its radio and transmissions are not allowed. During the Inactive Period, devices enter in sleep mode, a state in which little energy is consumed.

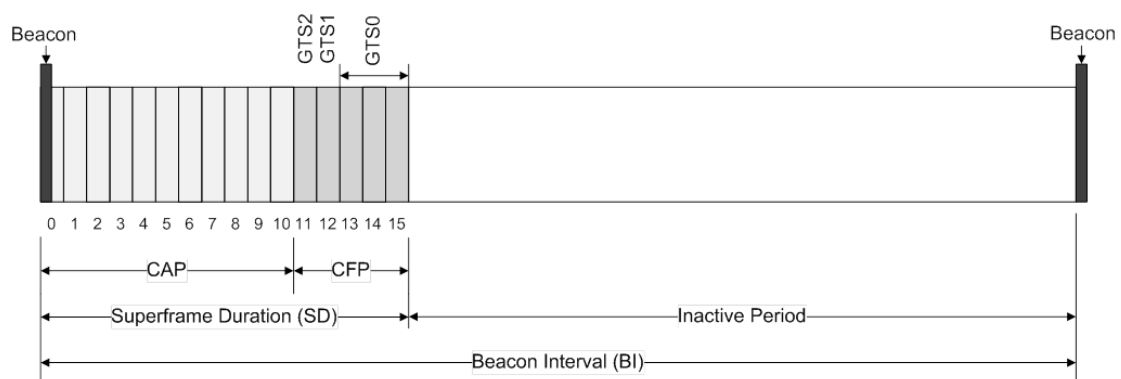


Figure 14 – Superframe structure [91, 92].

Each superframe contains 16 equally sized time slots. Up to seven time slots can be allocated to devices that contain strict bandwidth or latency requirements. These time

slots are called guaranteed time slots. The remaining time slots constitute the contention access period (CAP). During this period, devices contend to access the wireless channel using the slotted version of the CSMA-CA mechanism. A network that supports the use of beacons is called a beacon-enabled network, whereas nonbeacon-enabled networks do not support the transmission of beacons and, consequently, do not employ a superframe structure.

A PAN may choose not to transmit beacons. In this case, the coordinator is permanently active and, when a device wishes to transfer data to the coordinator, it simply transmits its data frame using the unslotted version of the CSMA-CA, and then switches off its radio. The CSMA-CA mechanism is presented in Section 2.2.3.

2.2.2 Physical layer

The original Physical (PHY) options introduced by the first version of the protocol and the optional PHY options included in the 2006 version are shown in Table 5 [47, 91, 92]. The alternative PHY extensions introduced by the IEEE 802.15.4a (adds 14 overlapping chirp spread spectrum (CSS) channels in the 2450 MHz band, and 16 channels in three UWB bands), IEEE 802.15.4c (support to Chinese bands) and IEEE 802.15.4d (support to Japanese bands) amendments are not included [93-95].

Table 5 – PHY options introduced by the IEEE 802.15.4-2006 revision [92].

Channel page	Channel number	Optional?	Frequency band	Modulation	Spreading method	Data rate (kbps)	Symbol rate (ksymbol/s)
0	0	No	868 MHz	BPSK	DSSS	20	20
	1 – 10	No	915 MHz	BPSK	DSSS	40	40
	11 - 26	No	2450 MHz	O-QPSK	DSSS	250	62.5
1	0	Yes	868 MHz	ASK	PSSS	250	12.5
	1 – 10	Yes	915 MHz	ASK	PSSS	250	50
	11 - 26	Reserved	---	---	---	---	---
2	0	Yes	868 MHz	O-QPSK	DSSS	100	25
	1 – 10	Yes	915 MHz	O-QPSK	DSSS	250	62.5
	11 - 26	Reserved	---	---	---	---	---

The 2006 version of the protocol added the concept of page, in which frequency channels are defined by a combination of channel numbers and channel pages. The PHY options on channel page 0 are the ones originally defined in the 2003 original version of the protocol. The PHY options on channels 1 and 2 were included in the 2006 protocol version and are optional.

The adoption of a specific PHY option must consider specific country regulations. For instance, in Europe, the European Standard Institute recommends the use of the 868 MHz, whereas in the United States, Australia and New Zealand, the 915 MHz is recommended by local regulation agencies. On the other hand, the 2450 MHz is allocated worldwide to ISM applications [47].

The advantages of the 868 MHz and 915 MHz bands are that they offer a longer range and are less crowded. Alternatively, the 2450 MHz band is available worldwide, offers more channels and requires potentially smaller antennas.

The center frequency (F_c) of each channel is determined using the following Equations (1), (2) and (3) [91, 92]:

$$F_c = 868.3 \text{ in megahertz, for } k = 0 \quad (1)$$

$$F_c = 906 + 2(k - 1) \text{ in megahertz, for } k = 1, 2, \dots, 10 \quad (2)$$

$$F_c = 2405 + 5(k - 11) \text{ in megahertz, for } k = 11, 12, \dots, 26, \text{ where } k \text{ is the channel number.} \quad (3)$$

Figure 15 shows the alignment of IEEE 802.15.4 (2450 band) and IEEE 802.11b PHY channels in the United States and Europe⁸. Four channels are between or above the guard bands of IEEE 802.11b channels. This is the case of channels 15, 20, 25 and 26 in relation to channels used in the United States and channels 15, 16, 21 and 22 in relation to channels used in Europe. The IEEE 802.15.4 specification recommends the operation in one of these channels to enhance coexistence with IEEE 802.11b networks [91, 92].

⁸ The IEEE 802.11b standard recommends the channel selection shown in Figure 15, though channel use depends on the restrictions imposed by each country regulations and user preferences.




Figure 15 – Alignment between IEEE 802.11b and IEEE 802.15.4 PHY channels [91, 92].

IEEE 802.11g and IEEE 802.11n also operate on the 2450 MHz band, but have distinct PHY channels characteristics. Whereas the receiver bandwidth of a IEEE 802.11b signal is 22 MHz, that of IEEE 802.11g and IEEE 802.11n signals are 16.25 MHz and 33.75 MHz, respectively. However, typically, IEEE 802.11g uses the same channel structure as IEEE 802.11b (that is, three channels share the same area), which cannot occur in the case of IEEE 802.11n [134]. Before setting up an IEEE 802.15.4 network, it is recommended to make a wireless survey to determine which channels are being used by IEEE 802.11-based networks and restrict the use of IEEE 802.15.4 operating channels to the ones with the least interference.

Figure 16 shows the PHY layer reference model. As shown, it provides two services to the higher layer, accessed by two service access points (SAPs).




Figure 16 – The PHY layer model [91, 92].

The physical data SAP (PD-SAP) supports the transport of data between two MAC sublayers [91, 92].

The physical layer management entity SAP (PLME-SAP) provides service interfaces to layer management functions. The services provided by the PLME-SAP includes [91, 92]:

- Energy detection (ED) within the current channel;
- Link Quality Indicator (LQI) for received packets;
- Clear channel assessment (CCA) for carrier sense multiple access with collision avoidance (CSMA-CA); and
- Channel frequency selection.

Additionally, the PLME-SAP maintains a database, the PHY PAN information base (PHY PIB), which contains objects that belong to the PHY layer. The PHY PIB attributes are described in Table 6 [91, 92].

Table 6 – PHY PIB attributes [91, 92].

PIB attribute	Description
phyCurrentChannel	Channel to use for transmission and reception.
phyChannelsSupported	The list of channels to be scanned in case a higher layer requests any kind of scan. Scans include passive, orphan, ED and active scans.
phyTransmitPower	The transmit power and tolerance.
phyCCAMode	The CCA mode to be used (see description in the text).

Before transmitting, devices perform the CCA according to three different modes defined in the standard. The following CCA modes are defined and, at least one should be implemented [91, 92]:

- CCA Mode 1 - Energy above threshold. The PHY reports a busy channel if the energy detected is above the ED threshold.
- CCA Mode 2 - Carrier sense only. The PHY reports a busy channel if a signal compliant with the standard is detected regardless if the signal is above or below the ED threshold.
- CCA Mode 3 - Carrier sense with energy above threshold. Combines both methods. The PHY reports a busy channel if the complaint signal detected is above the ED threshold.

PHY PIB attributes can be retrieved and updated using get and set primitives provided by the PLME-SAP.

Two important characteristics of the PHY layer are defined by two constants. The *aMaxPHYPacketSize* defines the maximum size of any frame (PHY service data unit or PSDU) the PHY can transmit or receive, which is 127 octets. The PHY constant *aTurnaroundTime* defines the maximum time required by the transceiver to change from transmit mode (TX) to receiver mode (RX) and vice-versa, which is equal to 12 symbol periods.

2.2.3 Medium access control layer

Figure 17 shows the MAC sub-layer reference model [91, 92]. The MAC provides two services, accessed through two SAPs. The MAC data service, accessed through the MAC common part sublayer-service access point (MCPS-SAP), supports the transport of data. The MAC management service, accessed through the MAC layer management entity-service access point (MLME-SAP), provides management functions. Similarly to the PLME, it maintains a database, the MAC PIB [91, 92].




Figure 17 – The MAC model [91, 92].

Three types of data transfer can occur: a) from a device to the coordinator; b) from the coordinator to a device; and c) between peer devices. The two first methods can occur in star and peer-to-peer networks, whereas the third method can only occur in a peer-to-peer network [91, 92].

The MAC parameters and constants used and their descriptions are presented in Table 7. For each transmission attempt, the device maintains three variables, NB, CW and BE. NB is the number of backoffs the device can make before it declares a channel access failure. CW is the contention window length and is used only in the slotted version of the algorithm. It defines the number of backoff periods that need to be clear

of channel activity before the transmission can start. BE is the backoff exponent. It is used to determine the number of backoff periods the device shall wait before performing the CCA [91, 92].

Table 7 – CSMA-CA parameters and constants [91, 92].

Parameter	Duration or range and default value		Description
	2003 version	2006 version	
<i>aUnitBackoffPeriod</i>	20 symbols (0.32 ms for 2450 MHz PHY)		The number of symbols forming the basic time period used by the CSMA-CA algorithm.
<i>macMinBE</i>	0 – 3 (default = 3)	0 – <i>macMaxBE</i> (default = 3)	The minimum value of the backoff exponent.
<i>aMaxBE</i>	5	---	The maximum value of the backoff exponent.
<i>macMaxBE</i>	---	3 – 8 (default = 5)	
<i>macMaxCSMAbackoffs</i>	0 – 5 (default = 4)	0 – 5 (default = 4)	The maximum number of backoff periods.
<i>aMaxFrameRetries</i>	3	---	The maximum number of retries allowed after a transmission failure.
<i>macMaxFrameRetries</i>	---	0 – 7 (default = 3)	

The unslotted version of the CSMA-CA algorithm defined in the 2003 version of the protocol is shown in Figure 18 [91, 92]. In the 2006 version, the constant *aMaxBE* should be substituted by the *macMaxBE* parameter in step (1).

Initially, the backoff exponent (BE) takes the value *macMinBE* and the number of transmission attempts for the current packet (NB) is set to zero. After that, the device waits for a random backoff interval defined in the range from 0 to $(2^{\text{BE}} - 1)$ unit backoff periods, where one unit backoff period is equal to 20 symbols (0.32 ms, in case of the 2450 MHz band). Next, if the CCA function indicates that the channel is idle, the device starts its transmission immediately after its transceiver changes from receiver to transmit mode. On the other hand, if the channel is busy, the device defers its transmission, increments the number of transmission attempts for the current packet, NB, and also increments the BE value, if BE has not exceeded *aMaxBE*. Then, if the maximum number of transmission attempts has not been exceeded, it calculates a new random backoff interval it must wait before assessing the channel again. The device can try to access the channel a maximum number of times defined by the *macMaxCSMAbackoffs* parameter. When this limit is reached, the MAC layer discards the data and declares a channel access failure.




Figure 18 – Unslotted version of the CSMA-CA mechanism defined in the 2003 version of the protocol [91, 92].

If beacons are being used in the PAN, the MAC sublayer shall employ the slotted version of the CSMA-CA algorithm for transmissions in the CAP of the superframe. The slotted version of the CSMA-CA algorithm is presented in Figure 19. In the 2006 version, the constant *aMaxBE* was substituted by the *macMaxBE* parameter in step (1).

In the slotted version, transmissions should start in the beginning of a backoff period, where each backoff period is equal to 20 symbols and the first backoff period is aligned with the start of the beacon. First, the MAC layer initializes the NB, the CW and the BE, which depends if the battery life extension field is set or not. Then, it locates the beginning of the next backoff period boundary, delays for a random number of backoff periods and request that the PHY perform the CCA in the current superframe. If the channel is found busy, the MAC layer updates the values of NB and BE. If the number of backoffs exceeds *macMaxCSMABackoffs*, it shall declare a channel access failure. On the other hand, if the channel is found idle, the MAC layer decrements the contention window and verify if it has expired. If it has not expired, the device should perform another CCA and repeat this process; otherwise, the algorithm declares success and the device shall transmit just after its transceiver changes to transmit mode.




Figure 19 – Slotted version of the CSMA-CA mechanism defined in the 2003 version of the protocol [91, 92].

Despite the CSMA-CA mechanism, even when the channel can be accessed and messages get transmitted, they might not reach the destination due to several factors, such as collisions, fading and interference. Additionally, the CSMA-CA mechanism has no specific means to avoid the hidden-node and the exposed-node problems.

A generic hidden-node scenario is illustrated in Figure 20 (a), where the transmission range of nodes A and B are represented by circles drawn around the nodes. If A is transmitting to C, and B starts to transmit to D (B’s carrier sense fails), A’s and B’s packets collide at C. In this case, the transmission from A, which is hidden from B, is corrupted by B’s transmission. A more severe problem happens if C and D happen to

be the same node, that is, if A and B transmit to a common receiver, C, as seen in Figure 20 (b). In this case, if B starts to transmit while A is transmitting (or vice-versa), both packets are corrupted.




Figure 20 – Hidden-node scenarios.

The exposed-node problem is depicted in Figure 21. Assume that C's transmissions to A do not collide with B's transmissions to D, but B can sense transmissions done by C. Suppose that after C starts a transmission to A, B decides to transmit a message to D. As B can sense C's transmission, the CSMA-CA algorithm reports a busy channel and, consequently, B defers its transmission and bandwidth is wasted [13, 118].




Figure 21 – Exposed-node scenario.

2.3 The ZigBee protocol

ZigBee is a standard-based commercial protocol developed by the ZigBee Alliance, a non-profit association of companies, governmental regulatory groups and universities [221]. In opposition to other related protocols (for instance, RF4CE was designed as a replacement technology to IR-controlled products, whereas WirelessHART is targeted at industrial automation), ZigBee was designed as a general purpose protocol that can fit several applications. Additionally, it was designed to support multi-application environments and interoperability between devices from various manufacturers.

By the time this work is being written, home automation and smart energy were the main application areas of the protocol. Home automation products include home security systems [7], electronic door locks [127] and light control systems [29]. Smart energy products include energy meters that support wireless data communication [46, 48] and energy management tools that enable consumers to track energy consumption

[32]. Other interesting applications include a greenhouse monitoring implementation [83] and a parking lot automation system [145].

2.3.1 ZigBee protocol overview

ZigBee stands on top of the version 2003 of the IEEE 802.15.4 protocol and defines the Network and the Application layers. Alike the IEEE 802.15.4, it was designed for low battery consumption, low cost, low data rate (250 kbps maximum) and easy installation. It adds routing capabilities to support tree and mesh topologies and provides enhanced security services.

The first version of the ZigBee protocol, ZigBee version 1.0, was released in December 2004. This work is based on this version. The second version, released in December 2006, was followed by the ZigBee 2007/PRO specification. Each new release adds to and improves functionality provided in previous versions of the specification [42]. The ZigBee 2007/PRO is compatible with the ZigBee 2006 version, whereas the back compatibility with the 2004 version is not assured by the ZigBee Alliance.

The following types of devices are defined [220, 224]:

- A ZigBee coordinator is equivalent to the IEEE 802.15.4 PAN coordinator.
- A ZigBee router is an IEEE 802.15.4 FFD, which can work as an IEEE 802.15.4 coordinator and is capable of routing messages and accepting associations.
- A ZigBee end device is an IEEE 802.15.4-2003 RFD or FFD participating in a ZigBee network, which is neither the ZigBee coordinator nor a ZigBee router.

The ZigBee layer model is shown in Figure 22. Both the Security Services Provider and the ZigBee Device Object (ZDO) offer services to the NWK and APL layers. The application objects 1 – 240 are developed by manufacturers (user users). They use the Application Framework and share Application Support Sublayer (APS) and security services with the ZDO [220, 224].

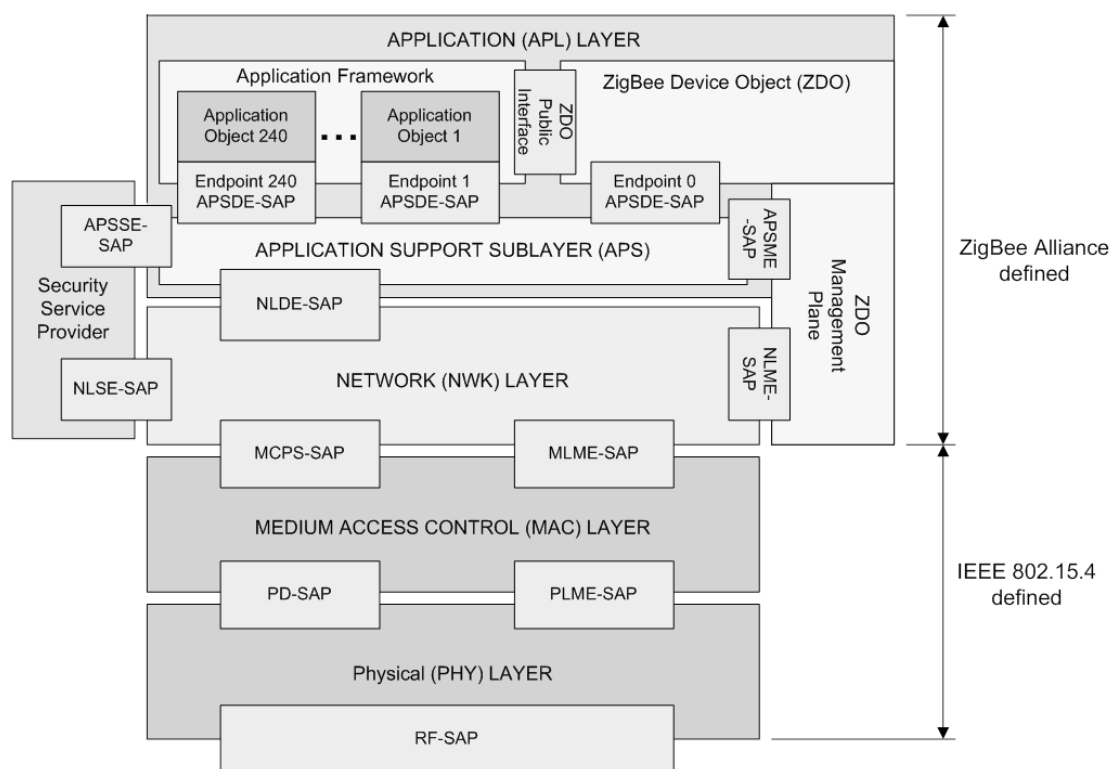


Figure 22 – ZigBee model [220, 224].

2.3.2 The Network layer

The Network (NWK) layer is the lower layer defined by the ZigBee protocol. This layer includes mechanisms used to join and leave a network; to apply security to frames and to route frames to their intended destinations. In addition, it is responsible for discovery and maintenance of routes between devices; the discovery of one-hop neighbors; and the storing of pertinent neighbor information are done at the NWK layer. The NWK layer of a ZigBee coordinator is responsible for starting a new network, when appropriate, and assigning addresses to newly associated devices [224].

The NWK layer includes two service entities to provide services to the Application layer. The NWK layer data entity (NLDE) provides the data transmission service through its associated SAP, the NLDE-SAP, whereas the NWK layer management entity (NLME) provides the management service via its associated SAP, the NLME-SAP. The NLME utilizes the NLDE to achieve some of its management tasks and maintains a database of managed objects known as the network information base (NIB) [224].

The NLDE is responsible for generating the network level protocol data unit (PDU) and for assigning to it the correct address (final destination or the next hop towards the final destination).

Three general communication mechanisms are available: unicast, broadcast and multicast. Unicast is used to send a message to a single device, whereas broadcast messages are sent to all devices within a given radius. Broadcast transmissions are not acknowledged. Instead, a passive acknowledgement mechanism may be used. Passive acknowledgement means that every ZigBee router and ZigBee coordinator keeps track of which of its neighboring devices have successfully relayed the broadcast transmission [220]. Multicast transmissions are used to send a message to devices that belong to a specific multicast group, and within a given transmission radius measured in hops.

The NLME supports network start, join and leave requests; new device configuration; address assignment; one-hop neighbor and route discovery; and receiver synchronization.

The ZigBee protocol allows two types of routing: tree and mesh. The tree routing allows devices to relay messages based on the address of the destination, without using a routing table, because the addresses are assigned to devices in a special way (using the Cskip algorithm, described in [224]). In mesh networks, convenient routes to the destination are established on demand using the Ad hoc On-Demand Distance Vector (AODV) routing protocol [118].

Tree routing is simpler than mesh routing, but is less reliable. In the first case, there are no options, but to follow tree links to reach the destination; while on mesh routing devices can establish new routes, picking the best links, in case a new route must be established or an old one fails. Tree routing is not allowed by the ZigBee PRO feature set.

Two other routing establishment mechanisms commonly referred as route aggregation were introduced in the ZigBee PRO feature set: many-to-one and source routing. These mechanisms are tailored to specific situations where the establishment of individual routes using the AODV routing protocol is less efficient.

The many-to-one routing is optimal for networks where most nodes transmit data to a data collector node or gateway. The data collector broadcasts a single many-to-one

routing request message to establish reverse routes on all devices. The many-to-one route request should be sent periodically to update and refresh the reverse routes in the network [56, 220].

In contrast to many-to-one routing, source routing addresses the situation where messages are sent from a data source node out to multiple remote nodes. The source node sends periodically many-to-one routes to establish a route to it on each device. When a node sends a message using a many-to-one route, it first sends a route record transmission. The route record transmission is unicast along the many-to-one route until it reaches the central node. As the route record crosses the many-to-one route, it appends the 16-bit address of each device in the route into the message payload. When the route record reaches the source node, it contains the address of the sender, and the 16-bit address of each hop in the route. The source node can store the routing information and retrieve it later to send a source routed packet to the remote node [56, 220].

2.3.3 The Application layer

The Application (APP) layer consists of the Application Support (APS) sub-layer, the application framework and the ZDO [224].

The Application Support sub-layer

The APS sub-layer is responsible for data transmission of application PDUs between ZigBee devices services. Additionally, all versions provide discovery and binding of devices, and maintain a database of objects, called APS information base (AIB). Earlier versions support APS-level encryption, duplicate frame rejection and fragmentation.

There are two types of discovery: device discovery and service discovery. Once a device joins the network, it may decide to discover a device it can talk to. Device discovery involves interrogating a node to discover its MAC address or its network

address⁹. Service discovery is the process whereby services available on a receiving device are discovered by other devices. It uses the information contained on a descriptor [224].

Four addressing modes are defined by the protocol: direct, indirect, broadcast, and group¹⁰. Direct addressing (or unicast addressing) is used by one device to send a message to another one by indicating its full address and endpoint. Indirect addressing is used by resource-constrained devices to communicate to bound devices without having to know its address. The message is sent to the coordinator, which looks up the required addressing fields and retransmits the message to each corresponding destination. Broadcast addressing (or application broadcast) is used to send a message to all endpoints on a destination device. Group addressing is used to selectively address a specific group of endpoints [220].

The Application Framework

The application framework contains up to 240 user-defined application objects, each one identified by endpoints 1 to 240. An endpoint is similar to a TCP port; it defines a communication entity that provides means to access a specific application running on a device and are identified by an 8-bit field [41]. Endpoints also allows for different devices to be implemented on a single node [73]. For instance, a device may function as an on/off light switch and a heating/cooling unit. Endpoint 0 is used to address the device profile (that is, all node descriptors), whereas endpoint 255 is used to address all active endpoints. Endpoints 241-254 are reserved for future use [220, 224].

Applications are developed based on application profiles. Application profiles corresponds to application domains (for instance, home automation) and defines the data exchange form for the application functions of a ZigBee device [57]. Application profiles issued by the ZigBee Alliance are called public profiles, whereas profiles developed by vendors are called private or manufacturer specific profiles. Application profile identifiers (IDs) are 16-bit values that range from 0x0000 to 0x7fff for public

⁹ The MAC address or IEEE extended address is a unique 64-bit address attributed to each device during manufacturing. The network address is a 16-bit address assigned during network association and used by the NWK layer for routing messages between devices.

¹⁰ Group addressing is not available in the ZigBee v1.0 (2004) protocol version.

profiles and 0xbf00 to 0xffff for manufacturer specific profiles. The following public application profiles have been released: Home Automation, Building Automation, Remote Control, Smart Energy, Health Care, Input Device, Telecom Services, Retail Services and 3D Sync [221]. The use of a public application profile allows interoperability between products from different vendors.

Each profile defines an enumeration of device descriptions and cluster identifiers, as shown in Figure 23. Both device and cluster identifiers are described by a 16-bit value.

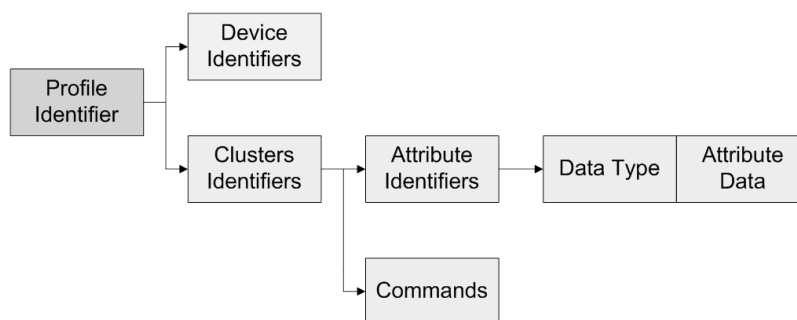


Figure 23 – Profile definition.

A device identifier specifies the device description supported on an endpoint. Table 8 shows few device identifiers defined by the Home Automation profile [223]. Device identifiers are mainly used for commissioning tools. ZigBee performs service discovery based on profile and cluster identifiers, not on device identifiers [73].

Table 8 – Some device identifiers defined by the Home Automation profile [223].

Device Name	Device Identifier
On/Off Light	0x0100
Dimmable Light	0x0101
On/Off Light Switch	0x0103
Dimmer Switch	0x104

Clusters are application objects and encapsulate both commands and data. For instance, in the Home Automation application profile, cluster ID 0x0006 corresponds to the cluster name On/Off. This cluster contains the attribute On/Off and defines the commands On, Off and Toggle. If an application needs to know the state of a light, it may read the value of the attribute associated to it. Additionally, if a bound light switch needs to change the state of a light, it should send the command Toggle defined in the cluster On/Off [73].

Clusters have direction, that is, they are listed as input or output. This is used for the purpose of service discovery. For instance, a switch, which implements cluster 0x0006 as an output, can find a light, which implements cluster 0x0006 as an input. On the other hand, a light shall not bind to another light as both implement cluster 0x0006 as an input [73]. The ZigBee Cluster Library is a document maintained by the ZigBee Alliance that contains cluster functionality. It was first released on October 2007, and is updated with new functionality as these are approved by the Alliance.

Commands are identified by an 8-bit value and their meaning depend on the cluster they belong to. For instance, 0x00 is “off” command in the OnOff Cluster, whereas 0x00 is the “move-to-level” command in the LevelControl Cluster [73].

Attributes encapsulate the data type and the attribute data and store the current state associated to a cluster. For instance, it stores the state of the light (“on” or “off”) controlled by an application that runs on an endpoint that implements the OnOffCluster from Home Automation profile.

Service discovery involves retrieving simple descriptors from peer devices. The simple descriptor contains information specific to each endpoint contained in the node, such as the profile identifier and the input and output clusters supported by the endpoint [220].

Apart from the simple descriptor, there are four additional devices descriptors: the node descriptor, the node power descriptor, the complex descriptor and the user descriptor. These descriptors apply to the complete node, whereas each simple descriptor corresponds to one specific endpoint defined in the node. The node, node power and simple descriptors are mandatory, whereas the complex and user descriptors are optional [220].

The node descriptor provides information about the capability of the node. For instance, it contains the node type (coordinator, router or end device), the frequency bands supported by the device and if it is an alternate PAN coordinator. The node power descriptor informs whether the device is battery powered and the current level of the battery. The complex descriptor is presented in XML form using compressed XML tags which contains, for instance, the device serial number and its model name. The user descriptor contains up to 16 ASCII characters used to identify the device (for instance, porch light or hall switch) [57].

Binding is the procedure used to create a unidirectional logical link between applications running in two or more devices (specifically, between a source endpoint/cluster identifier pair and a destination endpoint). In the 1.0 version (2004), resource constrained devices were allowed to send a message omitting the address of the bound destination device and other necessary information (endpoint, cluster identifier and attribute identifier). The message was addressed to the ZigBee coordinator, which maintained a binding table where it recorded the information required to relay the message to the destination [224].

Since the ZigBee 2006 release, the binding mechanism is called source binding. It involves the use of a binding table that is maintained by the APS layer of the node and allows applications to operate without having to manage recipient address information for the frames they generate. This information can be input at commissioning time without the main application on the device even being aware of it. The APS layer determines the destination address from its binding table, and then forwards the message to the destination application (or multiple applications) or group [220].

The ZDO

Endpoint 0 is reserved for the data interface to the ZDO. It is responsible for initializing the APS layer, the NWK layer and the Security Service Provider (SSP) and implementing discovery, security management, network management and binding management. Similar to the application profile defined in the application framework, the ZDO contain a profile, which is called ZigBee Device Profile (ZDP). The ZDP contains one device description and mandatory and optional clusters [57].

The device profile can be configured as a client and/or a server. In the client-server model, the server provides services to the client, which initiates requests for these services [57]. The standard also refers to clients as local nodes and servers as remote nodes [220]. ZDP commands are optional for clients. ZDP commands are divided in three groups: device and service discovery, binding management, and network management. Some ZDP commands and their descriptions are shown in Table 9 [220]. It is also shown the category of the device which originates the command, client or server, and if the server response to this command is mandatory or optional.

Table 9 – Some ZDP commands [220].

Application group	Command issued by the Client (C) or Server (S)	Mandatory (M) or Optional (O) processing by the server	Command	Description
Device and Service Discovery	C	M	NWK_addr_req	This command is used to request the network address of a remote device.
	S		NWK_addr_rsp	This command is issued in response to an NWK_addr_req command.
Binding	C	O	End_Device_Bind_req	Command used to request bind with a remote device. It is unicast to the ZigBee coordinator.
	S		End_Device_Bind_rsp	This command provides the status (SUCCESS, NO_MATCH, etc.) in response to an End_Device_Bind_req command.
Network Management	C	O	Mgmt_Rtg_req	The local device unicasts this command to a remote routing enabled device to request the content of its routing table.
	S		Mgmt_Rtg_rsp	This command is generated in response to an Mgmt_Rtg_req command.

Security

ZigBee security adds to the security model provided by the IEEE 802.15.4 protocol. It provides message integrity, authentication, freshness and privacy. The following definitions apply [47]:

- Link key: is a key used by the APS to secure application data between two devices.
- Network key: key used by the NWK layer for broadcast messages and other messages generated by this layer.
- Master key: is an optional key and is used for the generation of the link key.
- Trust center: is a device (the coordinator or a dedicated network device) that stores a list of allowed devices and all keys. Additionally, it performs key update and distribution and authenticates devices that request to join the network.

ZigBee security services include methods for key establishment and transport, device management, and frame protection. The encryption and authentication services are based on a standard security specification called counter mode encryption plus

cipher block chaining message authentication code (CCM) with the 128-bit Advanced Encryption Standard (AES) protocol [90, 189].

Security can apply to network and application layers. The layer that generates the message is responsible for securing it.

There are two security modes: a standard mode used by all ZigBee versions and a high-security mode used by ZigBee PRO. In the standard mode, APS communication between any two devices may optionally be encrypted with a link key. In the high-security mode, communication between any two devices may be encrypted with a link key. Additionally, in the standard mode, the Trust Center can update the network key by broadcasting a new key, which is not allowed in high-security mode. In this mode, link keys are derived using the symmetric-key key establishment (SKKE) [93] and no keys are exchanged unencrypted. Also, network key update is unicast from the Trust Center to each device, encrypted with the proper link key [47, 49].

In standard security mode, the list of devices, master keys, link keys and network keys can be maintained by either the Trust Center or by the devices themselves. In high-security mode, the Trust Center stores and maintains this information. Consequently, in high-secure mode, the Trust Center resource requirements grow with the number of network devices. Standard security requires less memory and bandwidth than high security, however high security affords greater security because all messages may be encrypted with link keys. Further, it allows devices to be forcibly removed from the network by excluding them from a network key update [39].

2.4 ZigBee versions comparison

Table 10 compares the different versions of the ZigBee protocol. Although they have many common features, recent versions contain improvements that are important to certain applications.

Table 10 – ZigBee versions feature comparison.

Feature	ZigBee v1.0 (2004)	ZigBee 2006	ZigBee 2007	ZigBee PRO
Stack size	(1)	Smaller	Small	Bigger
Frequency agility	No	Yes	Yes	Improved
Fragmentation	No	No	Yes	Yes
Addressing	Tree-based	Tree-based	Tree-based	Stochastic
Multicast	No	No	No	Yes
Routing	Tree and mesh	Tree and mesh	Tree and mesh	Mesh
Route aggregation (many-to-one and source routing)	No	No	No	Yes
Asymmetric link handling	No	No	No	Yes
Standard Security (AES 128 bit)	Yes	Yes	Yes	Yes
High Security mode (SKKE)	No	No	No	Yes

(1) Deprecated stack version. Only used in legacy systems.

The ZigBee PRO feature set requires the larger amount of memory (ROM and RAM). On the other hand, it provides the most features. For instance, ZigBee PRO establishes a procedure for tracking channel failures and reporting them to a special device, the Network Channel Manager, which can take frequency agility measures in case interference is detected.

Fragmentation is the ability to handle data transfers that are larger than the maximum payload size that a frame can handle. This functionality is specified in the 2007/PRO version of the standard.

The addressing mechanism in all ZigBee versions, except the ZigBee PRO feature set, is tree-based. ZigBee PRO devices attribute stochastic addresses to children nodes, which they can maintain as long as no conflict is detected. Though this addressing scheme presents advantages [216], a ZigBee PRO-based network requires more time to be set up because, when a device joins the network, it broadcasts the address it has chosen, to check for any conflict. If a conflict is detected, the device picks a new random address and makes a new broadcast to announce its new address. This procedure is also executed when a device rejoins the network.

ZigBee version 2006 has included group addressing, which is also supported by the 2007 version. Devices can be assigned to groups, and whole groups can be addressed with a single frame; thereby reducing network traffic for packets destined for groups [43].

Broadcasting involves the transmission of a message to every device in the network, within a given transmission radius measured in hops. Unlike broadcasting, multicasting is used to transmit a message to nodes that belong to a multicast group, despite the shape of the multicast group, in a more efficient way.

All multicast messages received by devices that belong to the multicast group are broadcasted again. On the other hand, messages received by devices that do not belong to the multicast group are unicasted if the value of the non-member radius field is not zero. Every time a message is unicasted, the non-member radius field is decremented. A disconnected multicast group is shown in Figure 24. M1 originates a message addressed to the multicast group that consists on devices M1-M5. The message is initially received by M2 and then broadcasted again. After that, the message is received by A1. A1 confirms that the value of the *non-member radius* field attached to the received message is greater than zero and unicasts it. Eventually, the message reaches all multicast group members [73].




Figure 24 – Multicasting in ZigBee PRO.

The many-to-one and source routing mechanisms employed by the ZigBee PRO feature set were already described in Section 2.3.2. They are more efficient than the routing based on the AODV mechanism for situations that involve several remote nodes transmitting messages to a data collector node or vice-versa.

All versions except the ZigBee PRO feature set consider unidirectional links. Assuming that A and B are neither neighbors nor have a parent-child relationship, a route from A to B must exist before A sends a message to B. Similarly, a route from B to A must exist to allow B to send a message to A. ZigBee PRO, on the other hand, establishes bidirectional routes and avoids asymmetric low quality links [73].

The high security mode improves on the standard security offered by the ZigBee stack and is required by certain application profiles such as the ZigBee Smart Energy.

The choice of one stack over the other may consider the application requirements and the development platform. For instance, the smaller ZigBee stacks are adequate for devices that implement the Home Automation profile. Conversely, a deeper network, with a high quantity of mobile devices, benefits from the stochastic addressing mechanism implemented by the ZigBee PRO feature set [73].

2.5 Jennic's programming environment

This work is based on the JN5139 modules, from Jennic, which implement the ZigBee version 1.0 (2004) protocol. These modules contain a simple operating system called Basic Operating System (BOS), designed to be used in conjunction with the ZigBee stack implemented. The vendor offers application programming interfaces (APIs) that simplify the development of application programs.

BOS is a non-preemptive task scheduler, and tasks have the same level of priority. The lack of context switching between tasks removes the need for individual stacks and saves on memory space. BOS controls the execution of the ZigBee stack and user tasks. The ZigBee stack is implemented as a single task. Event mechanisms are provided to allow communication between tasks. The passing of data between tasks is achieved using a messaging mechanism. There is a default user task that is created as part of the BOS initialization. The PHY layer and MAC sub-layer run in interrupt context, as do events generated by the on-chip peripherals [103].

Application programs are developed in C language. Code can be build either using a makefile supplied by Jennic or using the Code::Blocks Integrated Development Environment (IDE) [105].

The flow of a generic program is shown in Figure 25. It consists of several functions that should be used by the vendor to develop application programs. The function `AppColdStart` is the application program starting point. It is used for system initialization. Before leaving this function, BOS must be initialized and started. The control goes back to BOS that, after some internal procedures, returns the control to the function `JZA_vAppDefineTasks`. This function is used to register an additional (and optional) user task. After executed, BOS assumes the control and, after executing internal actions, executes the function `JZA_boAppStart`. It is used to place calls to the application framework to register any ZigBee descriptors for the node's endpoints. Once

this is complete, a call should be made to start the ZigBee protocol stack before terminating the function and returning control to the BOS. At this point, the initialization is complete [113].

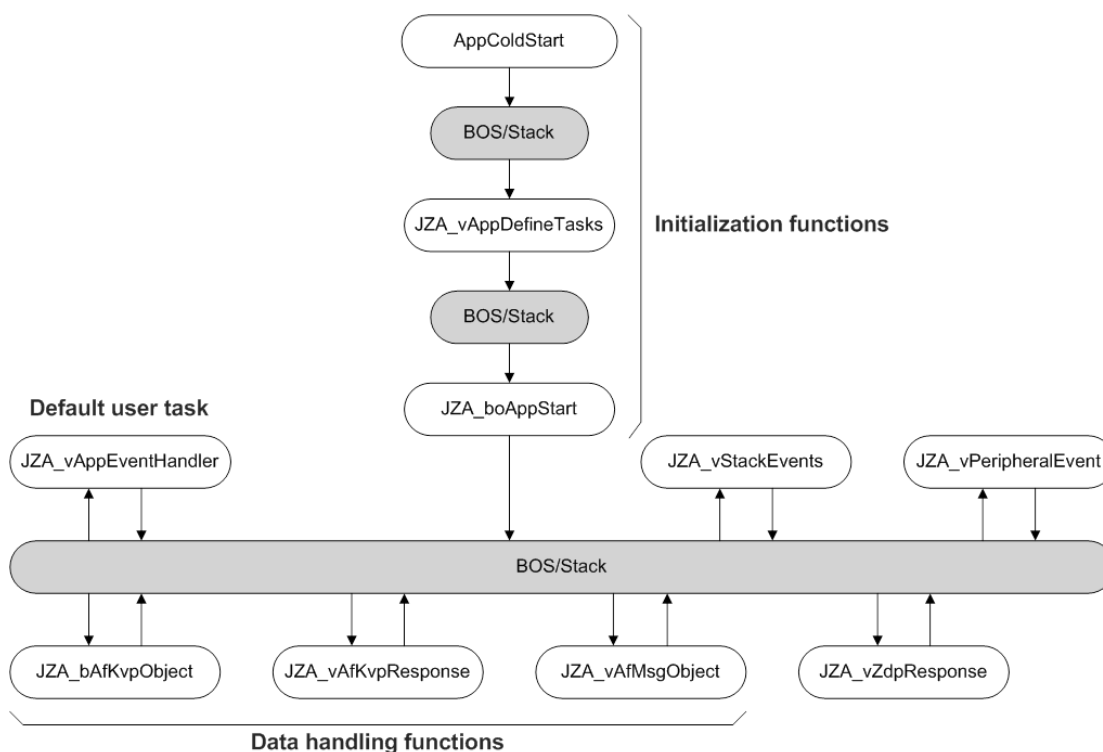


Figure 25 – Flow diagram of a generic application developed for the JN5139 module [113].

Once BOS and the ZigBee stack have been initialized, the control is passed to the user application through a number of functions. The function `JZA_vAppEventHandler` is called regularly by the BOS (the default user task is executed by this function). All tasks that should be executed regularly should be included in this function. The function `JZA_vStackEvents` is called when an event from the stack is received. For instance, it is called when the device joins or leaves the network and when the MAC layer reports a failure to access the channel. The function `JZA_vPeripheralEvent` is called, within the interrupt context, when a hardware interruption occurs. Hardware interruptions include interruptions generated by any hardware peripheral [113].

The functions `JZA_bAfKvpObject`, `JZA_vAfKvpResponse` and `JZA_vAfMsgObject` are called when a frame addressed to one of the endpoints running in the device is received. The function `JZA_bAfKvpObject` is called when a KVP frame is received, whereas the function `JZA_vAfMsgObject` is called when a MSG frame is received. The function `JZA_vAfKvpResponse` is called when an incoming

KVP response frame is received from another node. This may be an acknowledgment of an earlier command frame send to that node, or a response to a request for data [113].

The JZA_vZdpResponse is called when a response is received from the ZigBee Device Profile object (for instance, a response to a binding request) [113].

2.6 Other standard-based wireless communication protocols of interest

By the time this work started, the IETF (Internet Engineering Task Force) released the 6LoWPAN specification. ZigBee and IEEE 802.15.4 were emerging low-power wireless technologies; whereas IEEE 802.11 and Bluetooth were already established protocols. Only recently, Bluetooth Low Energy specification was released and ANT, a wireless protocol used on high-end fitness and sports high-end products, was considered for wellness management and home health monitoring applications. In the subsequent paragraphs these technologies will be briefly presented and compared in terms of their suitability for remote health monitoring applications.

An IEEE 802.11 standard network consists of stations (mobile or not) and wireless access points. A station is a computing device equipped with a wireless LAN network adapter. A wireless access point is a networking device equipped with a wireless LAN network adapter that acts as a bridge between stations and a traditional wired network, such as the Ethernet backbone. Two operating modes are defined by the standard: the ad-hoc mode where stations communicate directly and the infrastructure mode where there is at least one wireless access point connecting stations to a wired network.

The IEEE 802.11 family of standards defines the physical medium and MAC protocol specification devoted to wireless local area networks (WLANs). The technology associated to this family of standards is called Wi-Fi. There are four Wi-Fi technologies available, which are based on IEEE standards, as shown in Table 11. The IEEE 802.11n introduces a high throughput PHY based on orthogonal frequency division multiplexing (OFDM) modulation and MIMO (Multiple Input Multiple Output) technology with extensibility up to four spatial streams, resulting in a raw data rate of up to 600 Mbps. The IEEE 802.11-2007 includes QoS and high security services. Additionally, fast handover procedures are included in the IEEE 802.11r-2008 amendment.

Table 11 – Basic characteristics of Wi-Fi technologies available.

IEEE Standard	Frequency band ¹¹	Maximum data rate
802.11a	5 GHz	54 Mbps
802.11b	2.4 GHz	11 Mbps
802.11g	2.4 GHz	54 Mbps
802.11n	2.4, 5 GHz, 2.4 or 5 GHz (selectable) and 2.4 GHz and 5 GHz (simultaneously)	600 Mbps

The Bluetooth protocol, defined by the Bluetooth Special Interest Group (SIG), was developed to provide a universal radio interface for ad-hoc wireless connectivity, interconnecting computer and peripherals, PDAs, cell phones and other devices while maintaining high levels of security. It operates in the 2.4 GHz band and has a gross over the air data rate that varies from 1 Mbps to 3 Mbps, depending on the version, supporting video streaming, voice and data transmission applications. Three device classes are defined according to the range: class 3 devices have approximately 1 m range while class 2 devices can achieve 10 m and, finally, class 1 devices have an approximate range of 100 m¹². Frequency hopping and adaptive frequency hopping techniques are used to improve coexistence with other networks which operate in the same band [22].

Devices must join a piconet to communicate. A piconet which consists of a group of devices arranged in a star topology and synchronized in a specific fashion. In a piconet, one device provides the synchronization reference and is known as the master. All other devices are known as slaves. Up to seven slaves can actively participate in a piconet. Multiple piconets can join together to form a larger network known as a scatternet, but this topology is seldom used and is not defined by the core protocol. Likewise ZigBee, Bluetooth specification also contains several application profiles, which include the Health Device Profile [22].

By the end of 2009, the Bluetooth SIG released the Bluetooth v4.0 specification which includes the Bluetooth Low Energy definition and, in July 2010, this organization announced the formal adoption of the new specification. This new Bluetooth version

¹¹ The 5 GHz band extends beyond the 5800 MHz ISM band (see Table 1). The channels used are regulated per country.

¹² Bluetooth class 2 devices are the most common ones.

was designed to optimize the power consumption of low data rate network devices that transmit small packets. Bluetooth Low Energy enabled-devices will be able to operate on coin cell batteries which, depending on the application, will last more than one year [195].

The ANT protocol, defined by Dynastream, a Garmin owned company, is a proprietary wireless sensor network protocol running in the 2.4 GHz ISM band. It was initially conceived to be used by fitness devices, but recently it has being presented as an alternative for wellness applications [159]. The specification includes the four lower layers of the Open Systems Interconnection (OSI) model. ANT uses GFSK modulation and 78 radiofrequency (RF) channels. ANT devices transmit data packets that contain an 8-byte payload and use a raw data rate of 1 Mbps, four times higher than IEEE 802.15.4-based devices. The ranges, however, are typically smaller [198]. ANT handles peer-to-peer, star topologies and other more complex topologies. The range of an ANT network can be extended by using relays nodes, as shown in Figure 26.




Figure 26 – Star and cluster tree topologies in ANT.

6LoWPAN aims to enable IP version 6 (IPv6) packets to be transmitted over low power wireless networks, specifically IEEE 802.15.4. Considering that an IPv6 packet is much larger than an IEEE 802.15.4 one, it was necessary to define encapsulation and header compression mechanisms that allow IPv6 packets to be sent to and received from IEEE 802.15.4-based networks [151]. This approach has many advantages because low power devices can access directly IP-based devices without the need of complex gateways. Additionally, it is possible to directly adopt IP-based standards developed to provide security, configuration and other services. Moreover, the IP-based tools, procedures and techniques used to incorporate and manage devices can be readily utilized [35]. ZigBee Alliance has created an IP Task Group that is currently working

on a new stack that incorporates the support to communication over IPv6 and mechanisms to discover and communicate with IETF 6LoWPAN-based devices [219].

The discussed protocols are compared in Table 12. As shown, each has its own advantages and disadvantages. IEEE 802.11 is already used in hospitals to monitor patients. Typically, multi-lead ECG signals and SpO₂ values are monitored. Patient units are portable, but obtrusive. Newer units are powered by AA batteries that must be recharged every 25 hours. However, new low-power communication modules based on IEEE 802.11b can be powered by batteries and promise to bring Wi-Fi connectivity to portable medical devices.

IEEE 802.15.4, ZigBee and 6LoWPAN are valid options to IEEE 802.11 in the hospital environment to monitor low-acuity patients, in eldercare facilities to monitor residents and in residential environments to monitor recovering patients. Systems based on these protocols are expected to be pervasive and to include unobtrusive sensors.

Bluetooth and Bluetooth Low Energy are well fitted for single-hop communication links. Two examples are the communication between a blood glucose meter and a personal insulin pump and the transmission of heart rate values to a PDA. Bluetooth Low Energy is much more efficient than classic Bluetooth if only small packets are used and it is expected that this new protocol will enable a large number of new applications.

ANT is a proprietary protocol mainly used in fitness. Its simplicity and relatively high data rate are advantages over other low data rate protocols for health monitoring applications such as ZigBee and 6LoWPAN.

2.7 Summary

Initially, this chapter presents some of the main characteristics of WSNs. Then, it describes the standard-based wireless protocols used to develop this work: IEEE 802.15.4 and ZigBee. All protocol versions are addressed and the differences between them are pointed out. Finally, it overviews and compares wireless protocols that can be employed to develop health monitoring applications, including personal health and wellness monitoring, out-patient home monitoring, and vital signs monitoring in the hospital scenario.

Table 12 – Wireless protocols that can be used to implement patient remote monitoring applications.

Protocol	Major applications	Frequency band	Over the air raw data rate	Typical range	Typical autonomy on batteries	Main relative advantages	Main relative disadvantages
IEEE 802.11x	WLAN	2.4 GHz ISM band and 5 GHz	Up to 600 Mbps	Several meters	Hours or days	<ul style="list-style-type: none"> • Pervasive. • High throughput: IEEE 802.11a/g or very high throughput: IEEE 802.11n. • QoS, fast handover and high security support. 	<ul style="list-style-type: none"> • Higher cost. • Higher power consumption.
IEEE 802.15.4	Low data rate applications. P2P and star topologies.	2.4 GHz ISM band and sub-1 GHz ISM Band	Up to 250 kbps	Indoors: 10 – 30 m. Better in the sub-1 GHz band. Outdoors: up to 6,000 m.	Months or years	<ul style="list-style-type: none"> • Simple and small stack. • Low cost. • Supports low latency devices and mesh networking. 	<ul style="list-style-type: none"> • Lower data rate than IEEE 802.11x, Bluetooth, Bluetooth Low Energy and ANT. • Supports only P2P and star topologies (only PHY and MAC layers are defined).
ZigBee	Low data rate applications. Star, tree and mesh topologies.	Same as IEEE 802.15.4	Same as IEEE 802.15.4	Same as IEEE 802.15.4	Months or years	<ul style="list-style-type: none"> • Supports large networks with thousands of nodes. • Interoperability based on public application profiles. • High-security mode. 	<ul style="list-style-type: none"> • Lower data rate than IEEE 802.11x, Bluetooth, Bluetooth Low Energy and ANT. • Supports no specific QoS mechanism.
6LoWPAN	Low data rate applications with direct connection to IP based networks.	Same as IEEE 802.15.4	Same as IEEE 802.15.4	Same as IEEE 802.15.4	Months or years	Enables direct connection to IP-based networks.	Interoperability not already considered. No application profiles have already been issued.
Bluetooth	Cable replacement. PAN. Video, voice and data.	2.4 GHz ISM band	1 - 3 Mbps, depending on the version	1m, 10m or 100m, depending on the version.	Days or weeks	<ul style="list-style-type: none"> • Present in smart phones, PDAs, computers, etc. • Higher data rate than IEEE 802.15.4, ZigBee and 6LowPAN. 	<ul style="list-style-type: none"> • A Bluetooth device can communicate to only 7 active slave devices. Scatternets are not defined in the core protocol. • More complex stack than IEEE 802.15.4-based technologies and ANT.

Bluetooth Low Energy	PAN. Optimized for small packets and low data rate applications.	2.4 GHz ISM band	1 Mbps	10 m	Months	<ul style="list-style-type: none"> • Some Bluetooth devices will support Bluetooth Low Energy too. • Bluetooth Low Energy is up to 17 times more efficient than classic Bluetooth. • The energy usage is much lower than classic Bluetooth if only small packets are transmitted. 	<ul style="list-style-type: none"> • A Bluetooth device can communicate to only 7 active slave devices. • Scatternets are not defined in the core protocol.
ANT	WSN and PAN. Fitness and wellness devices	2.4 GHz ISM band	1 Mbps	10 m	Months or years	<ul style="list-style-type: none"> • Higher data rate than ZigBee. • Typically, ANT-based devices consume less power than IEEE 802.15.4 based devices (relatively smaller on-time). 	<ul style="list-style-type: none"> • Proprietary protocol. • Typically, ANT devices have a smaller range than IEEE 802.15.4, ZigBee and 6LowPAN devices.

Chapter 3

Health monitoring systems based on wireless technologies

Wireless communication technologies have been employed to monitor in-patients since the 70's. Through the years, Wireless Medical Telemetry Systems (WMTS) have undergone several changes. Systems' reliability has been considerably enhanced from a dropout of 50 min/day to a few seconds per day [17, 210]. Other major improvements include the increase in systems' capacity and coverage area. Despite these facts and the importance attributed by physicians to vital signs monitoring [59], only a small number of patients out of intensive and intermediate care units have their vital signs monitored.

The raise in health care costs and life expectancy combined with recent advances in microelectronics, smart textiles and wireless and sensor technologies have motivated academia and companies to propose innovative patient monitoring systems, which will have a positive impact on the way health care is provided [164]. Apart from hospital applications, other scenarios include:

- Monitoring of recovering patients at home;
- Chronic active patients health monitoring;
- Medication intake and specific behavior recording systems used to record medication adherence level or monitor specific behaviors (for instance, exercise level or food intake);
- Elderly monitoring at home, in assisted living facilities, in retirement communities, or in isolated areas; and
- Victims in disaster scenarios.

This chapter starts with a review of general requirements of patient monitoring systems, which are referenced in subsequent sections. Then, regulations, standards and

profiles applied to medical devices are introduced. Finally, some telemetry systems used to monitor patients are described. The limitations of those systems are discussed and possible ways to improve them are suggested. These findings are based not only on a thorough literature review, but also on several interviews conducted with physicians, nurses and engineers during a research period at Massachusetts Institute of Technology (MIT).

This chapter initially presents the most important requirements that must be fulfilled by remote monitoring systems and reviews regulation, standards and profiles that apply to medical devices. Then, recent research on remote health monitoring is reviewed. Finally, the views of physicians and informatics professionals on important aspects of remote patient monitoring systems are presented.

3.1 Patient monitoring general requirements

The general requirements of patient monitoring systems include the following:

Identification of emergency situations: It is important that individual sensors or network coordinator devices identify a potential life-threatening condition as it emerges and make all possible efforts to report it [206].

End-to-end reliability: Both routine and emergency messages should have a high probability of being correctly received by the health care provider who is intended to interpret the data and may have to take actions to assist the patient [206].

End-to-end delay: Messages must be delivered in reasonable time, determined by the level of emergency and according to specific normative documents.

Security: In telecommunications, security involves the ability to conserve or prove that the following attributes apply to devices and to their communications [53]:

- **Authenticity:** Ensures that device A is really A and that a packet received from device A was not transmitted by device B masquerading as A.
- **Authority:** Ensures that an entity is allowed to perform a requested task.
- **Integrity:** Ensures that the data received is the same data transmitted.
- **Confidentiality:** Ensures that data transmitted by device A to device B is not revealed to other devices capable of receiving A's transmissions.

Additionally, the IEEE P11073-00101 standard [87] recommends the implementation of general network security actions to prevent malware and denial of service attacks to ensure the physical security of devices.

Small form-factor and unobtrusiveness: Ideally, devices employed to continuously monitor mobile in-patients and out-patients shall be wearable and shall have no cables. However, several devices used to monitor patients, despite being portable, are relatively obtrusive.

Power efficiency: It is required that battery replacements are not very frequent so that patients can be monitored during their stay in a certain service (for instance, while recovering from an ambulatory surgery) and do not interrupt other actions that should be taken by health care providers (for instance, it is expected that a telemetry cardiac monitor can work for an amount of hours that exceeds a nurse's shift). Ideally, batteries should last for several days or even months without replacement.

Scalability: The communication network should be able to scale well in terms of the number of monitored patients. Additionally, it should be able to monitor if, at any time, the number of devices reaches its capacity and take measures to ensure that it operates as specified.

3.2 Regulation, standards and profiles

In the broader sense, a wireless medical device might include the device itself and the infrastructure necessary to transport data. The IEEE 11073-00101 standard [87] describes a wireless medical device as being comprised of four stages, as shown in Figure 27. The first stage includes the medical device, which can be an external device (e.g., a blood pressure monitor with wireless connectivity), a wearable or an implantable device. Data generated by devices are transported through downstream stages until reaching the patient or a health care provider.




Figure 27 – Stages of a wireless medical device [87].

Medical devices (stage 1 of Figure 27) are regulated by federal government bodies. A medical device may be classified as Class I (including Classes Is and Im), IIa, IIb or III; with Class III covering the highest risk products. Class I medical devices can be put on market by self-certification, whereas medical devices that are classified on the other classes must receive a Certificate of Conformity issued by a Notified Body [215]. Examples of Class I devices include elastic bandages, examination gloves, and hand-held surgical instruments.

Classes IIa, IIb and III medical devices are regulated in European Economic Area (EEA) by three approach directives, depending on the classification of the device: Active Implantable Medical Device Directive, AIMDD (90/383/EEC); (General) Medical Device Directive, MDD (93/42/EEC); or In Vitro Diagnostic Medical Device Directive, IVDMD (98/79/EC). These directives are in convergence to standards issued by the International Organization for Standardization (ISO), where the most relevant is the ISO 13485:2003, which define the international quality system standards for medical devices [99]. These devices should also comply with product risk management, electromagnetic interference and compatibility (EMI/EMC), and usability regulations and must conform to local country laws on personal data privacy and hardware disposal (Restriction of Hazardous Substances Directive or RoHS).

Whereas stage 1 medical devices are strictly regulated, the other stages are not. Current standardized technologies included on stage 2 were not designed to transport medical data and to support the QoS requirements that this transport involves. According to [87], such systems might not be considered a medical device when operating under typical conditions. Connectivity, including wireless connectivity based

on standard-based technologies, is not considered on medical regulatory documents, but instead in nonmedical standards (e.g., IEEE standards) and nonmedical consortium agreements (e.g., Wi-Fi Alliance and ITU). The regulation within stage 2 would be complicated because data transport involves complex combinations of distinct technologies that include WPANs, WLANs, wireless metropolitan area networks (WMAN) and wireless wide area networks (WWAN). As pointed out in [87], data storage in stage 3 share the same difficulties as stage 2 as it is defined by several nonmedical standards. Additionally, in several cases, stage 4 merges with stage 3 as it is based on applications that include data storage and analysis.

The IEEE 11073 family of standards, which was adopted as ISO standard through ISO TC215, and was developed in coordination with other standards development organizations, including IEEE 802 committee, IHTSDO¹³, HL7¹⁴ and DICOM¹⁵, aims on providing real-time, plug-and-play¹⁶ interoperability between point-of-care medical devices. Additionally, it aims to promote the efficient exchange of care device data, where it assumes that information that is captured at the devices “can be archived, retrieved, and processed by many different types of applications without extensive software and equipment support, and without needless loss of information” [101].

Among the IEEE 11073 family of standards, it is important to highlight the IEEE 11073-00101 [87], which provides guidance for the use of RF wireless communication technologies for point-of-care (PoC) medical devices that exchange vital signs and other medical device information using shared information technology infrastructure. This standard covers several application use cases and considers potential

¹³ The International Health Terminology Standards Development Organization (IHTSDO) is a not-for-profit association that develops and promotes use of SNOMED CT, a multilingual health care clinical terminology.

¹⁴ The Health Level Seven International (HL7) is a global authority on standards for interoperability of health information technology.

¹⁵ The Digital Imaging and Communications in Medicine (DICOM) is a standard designed to ensure the interoperability of systems that deal with medical images and derived structured documents as well as to manage related workflow.

¹⁶ Plug-and-play interoperability means that the user does not need to do any action, apart from connecting the device, to allow it to communicate data as defined.

applications for standard-based communication technologies, including IEEE 802.15.4 and ZigBee protocols.

The IEEE 11073-00101 standard presents QoS parameters (reliability, latency, priority, and bandwidth requirements) of typical patient monitoring applications, however; whereas maximum latency values are specified for each category of device data (real-time alarms and alerts, real-time waveform transmission, non-real time events, etc.), reliability is only qualitatively specified (highest/essential, high, medium and low). Table 13 presents the QoS attributes defined for some data classes of interest (bandwidth requirement is omitted). These parameters may be used in prioritizing classes of data generated by devices that share a network system. As shown, real-time alarms and alerts should have the highest reliability and priority. On the other hand, patient state changes and real-time reminders shall have high reliability, but medium priority. Real-time waveform data shall have high reliability and priority [87].

Table 13 – QoS requirements for different categories of medical device data [87].

Data type	Reliability	Latency	Priority
Alerts/alarms (real-time)	++++	< 500 ms from the wireless sensor to the gateway to the wired network < 3 s communication latency	++++
Patient state change (real-time)	++++	< 3 s communication latency	++
Reminder (real-time)	++++	< 3 s communication latency	++
Waveforms (real-time)	+++	< 3 s to central station < 7 s for telemetry to in-room monitor	+++
Physiologic parameters (real-time) (e.g., episodic BP, HR, SpO ₂ , ETCO ₂ , temperature)	+++	< 10 s to central station < 3 s communication latency from monitor to clinician	+++

Personal health device standards are a group of standards addressing the interoperability of personal health devices such as weighing scales, blood pressure monitors and blood glucose monitors. These standards differ from PoC standards due to an emphasis on devices for personal use (rather than hospital use) and the usage of a simpler communications model. This family of standards is based around a framework defined the IEEE 11073-20601-2008 standard [100] and its amendment, IEEE 11073-20601a-2010 [89]. At present, the following device specialization standards have been released [88]:

- IEEE 11073-10404 - Device specialization - Pulse Oximeter;
- IEEE 11073-10407 - Device specialization - Blood Pressure Monitor;

- IEEE 11073-10408 - Device specialization - Thermometer;
- IEEE 11073-10415 - Device specialization - Weighing Scale;
- IEEE 11073-10417 - Device specialization - Glucose Meter;
- IEEE 11073-10420 - Device specialization - Body composition analyzer;
- IEEE 11073-10421 - Device specialization - Peak flow;
- IEEE 11073-10441 - Device specialization - Cardiovascular fitness and activity monitor;
- IEEE 11073-10442 - Device specialization - Strength fitness equipment;
- IEEE 11073-10471 - Device specialization - Independent living activity hub;
- IEEE 11073-10472 - Device specialization - Medication monitor.

Additionally, other device specialization standards are being developed. They include the following [88]:

- IEEE P11073-10406 - Device specialization - Basic ECG (1 to 3-lead);
- IEEE P11073-10413 - Device specialization - Respiration rate monitor;
- IEEE P11073-10418 - Device specialization - INR (blood coagulation);
- IEEE P11073-10419 - Device specialization - Insulin pump.

The Personal, Home and Hospital Care (PHHC) profile developed by the ZigBee Alliance relies on the on-going work being developed by IEEE 11073 to allow the interoperability with medical devices [222]. Instead of developing its own protocol format and data encodings, this profile uses the IEEE 11073-20601 standard as the basis of the data exchanges between the devices supported by the PHHC profile. Additionally, it adopts the encoding formats defined by each device specialization standard. Consequently, the last released version of the ZigBee PHHC does not include the encoding format definition for devices that do not have a corresponding device specialization, such as electrocardiograms and respiration rate monitors.

Continua Health Alliance, a non-profit organization created to improve the quality of personal health care by promoting interoperability between health care devices and enterprise services, also relies on the use of industry standards, which include wire and wireless technologies such as USB, Bluetooth and ZigBee and the data model defined by the IEEE 11073 family of standards.

3.3 Wireless medical systems review

This section reviews wireless medical systems (both commercial and systems proposed by academia). Non-wearable and wearable personal monitoring systems intended for chronic diseases management and wellness are addressed on Sections 3.3.1 and 3.3.2. Systems specifically designed for elderly care, including personal monitoring systems, are addressed on Section 3.3.3. Finally, systems for hospital use are reviewed on Section 3.3.4.

3.3.1 Non-wearable personal monitoring systems

Some of the present medical devices used for personal health monitoring (e.g., blood pressure monitors and glucose meters) can supply data to applications running on personal devices (e.g., a portable dedicated patient unit or a PDA). Among these devices, some employ proprietary communication protocols, whereas others are interoperable. Interoperable devices are advantageous to users who can, for instance, measure a physiological parameter using medical devices from different companies and have the devices readings uploaded to an on-line application not created by any of the devices' manufacturers.

The Continua Health Alliance, which promotes the interoperability of personal health care devices, initially selected two communication technologies: USB (wired) and Bluetooth. In 2009, ZigBee and Bluetooth Low Energy were included; however, at the present time, no medical device based on these two wireless protocols has already been certified by the Alliance.

WellDoc is a commercial system for chronic disease management [214]. It was initially developed to help patients manage Type 2 diabetes, but it can be configured to support the management of other chronic conditions. The results of blood glucose measurements are sent wirelessly to a cell phone which runs an application that interacts with the patient, helping him attain his goals (see Figure 28). Additionally, the application running in the cell phone sends data to a central processing station where data are analyzed and added to the patient's logbook. Depending on the results of the data analysis, the patient's primary caregiver is contacted.



Figure 28 – WellDoc remote application is asking .a patient to confirm the blood glucose value sent by the blood glucose meter using a Bluetooth connection (adapted from [214]).

3.3.2 Wearable personal monitoring systems

Wearable systems may contain sensors that directly connect to a data aggregator (e.g., a Wi-Fi residential router); often sensors integrate a body area network (BAN), a collection of wearable sensors directly connected (using wired or wireless links) to a network controller node (also called base station).

Despite the apparent simplicity, the design of wearable systems for health monitoring offers several challenges. Sensors must satisfy strict medical criteria, including suppression of noise caused by the displacement of electrodes or sensing parts due to user movement. Additionally, sensors should be small and unobtrusive. Considering that wireless communication technologies are used, reliability, end-to-end latency and data security should be assured. Both sensors and the network controller node should minimize energy consumption to increase operational lifetime. Finally, these systems need to be affordable to ensure wide public availability [164].

Mobihealth was one of the first personal health monitoring systems proposed by academia [125]. It integrated sensors and actuators to a wireless BAN based on the Bluetooth protocol. Other systems based on this communication technology include MyHeart [81, 137], WEALTHY [165], and MagIC [177], which integrate sensors and smart textiles.

With the introduction of low-power technologies, several applications proposed by academia moved from Bluetooth-based systems to systems based on small embedded systems with a wireless low-power transceiver, usually termed motes. Some motes manufactured by Crossbow (now Moog Crossbow Inc.) are shown in Figure 29. Both CodeBlue [135] and AID-N [201] were developed based on motes and address emergency scenarios where overwhelming quantities of patients need to be monitored, tracked and triaged. CodeBlue is a common protocol and software framework, which

allows wireless monitoring and tracking of patients and health care providers. It consists of physiological sensors, a developed 1-lead ECG sensor and a commercial SpO₂ sensor mounted as daughter boards that attach on Mica2 motes [202]. These sensors communicate with handheld commercial devices used by first care responders where an application displays measured parameters and waveforms. It also includes wearable sensors and RF beacon nodes used to track patients. AID-N extends upon CodeBlue by adding electronic triage tags based on MicaZ motes, which contain an IEEE 802.15.4 compliant RF transceiver [202].



Figure 29 – Cricket, Telos, Mica2 and MicaZ motes (adapted from [204]).

With the emergence of the IEEE 802.15.4 protocol, several researchers proposed its use for personal health monitoring based on BANs consisting of small wearable sensors. Some examples are the Ubimon system [133], the health monitoring systems developed as part of the European project Sensation [152], and the physiological sensing applications developed under the Human++ research program, from IMEC [167].

In environments where several patients are being monitored using BAN-based systems, it is possible that a sensor, instead of associating to the BAN coordinator belonging to the person that is wearing it, incorrectly reassociates to another network in the vicinity. Researchers from Philips Research Center have proposed the use of body-coupled communication to allow a sensor to identify the person it belongs to and, consequently, be able to discriminate all other sensors attached to the same person [55]. The inconvenience of this approach is the need of a specific transceiver and associated hardware and software. On the other hand, it avoids the need of confirming the correct association of sensors after network establishment. Additionally, it allows sensors to be included to an established BAN without human intervention.

An interesting technology to automatically record dietary habits is being developed by researchers from TU Eindhoven and ETH Zurich, under the project MyHealth [11]. Although the work focuses on the sensors, it can be applied in a typical BAN scenario. Authors have employed body sensors (intake gestures, swallowing and chewing) to automatically record intake timings, food category and amount. A good recognition performance of categorical food (such as fruit and vegetables) intake has been achieved, which is an improvement of incorrect and biased self-recoding. However, as sensors are very obtrusive, their usability needs yet to be improved.

CardioNet and Nuvant are examples of commercial wearable systems specifically designed to monitor non-lethal cardiac arrhythmias. CardioNet [26] allows patients to enroll in a program in which they are continuously monitored using a collar-type wearable 3-lead ECG. Data are transmitted by an ECG sensor, using a Bluetooth link, to a patient device (a special PDA). When the patient device identifies an abnormality, it automatically sends data to a remote patient station, using a cellular connection, where specialized nurses analyze data, respond to events and report data to physicians.



Figure 30 – CardioNet system patient devices (adapted from [26]).

Nuvant [33], from Corventis, improves from CardioNet because it uses PiiX, a wearable, band-aid like, single-lead ECG sensor (shown in Figure 31). Like CardioNet, when an arrhythmia is detected by PiiX, it sends ECG data (15 sec pre- and 30 sec post-event) to zLink Mobile, a mobile PDA-like transmitter, through a wireless link. Then, data are transmitted using a cellular link, by the zLink Mobile to the company's monitoring center where emergencies are handled and data are recorded to be sent to physicians.




Figure 31 – Nuvant system, from Corventis (adapted from [33]).

The form factor of PiiX has many resemblances with the wireless ECG patch-like sensor developed by IMEC under the Human++ project [158]. It is based on the nRF24L01 low power radio, from Nordic, and is mounted on a flexible polyimide substrate, which increases the user comfort. The sensor without a case is shown in Figure 32.




Figure 32 – A patch-like ECG sensor developed by IMEC/Holst Center (adapted from [158]).

3.3.3 Systems based on wireless technologies for elderly care

In this section, the importance of technologies dedicated to help seniors live independently, safer and get help in case of emergency situations is discussed. Then, wireless systems designed for elderly care are briefly reviewed. Finally, concerns that might prevent or hamper the use of these technologies are described.

Importance of technologies dedicated to elderly care

The worldwide population of those over 65 is predicted to reach 761 million by 2025, more than double that it was on 1990 [178]. The proportion of elderly is growing in most developed countries. Japan currently leads the world with nearly one-quarter of

its population at age 65 or older, followed closely by Italy and Germany. Today, 40 million people in the United States are age 65 and older, but this number is projected to more than double to 89 million by 2050 [102]. In proportion of children, aging numbers are growing too in the United States. In 2011, the proportion of children has dropped to about one-fourth, whereas the share of elderly has risen to 13 percent. By 2050, one-fifth of the United States population will be age 65 or older [102].

Meeting the needs of elderly, many dealing with chronic diseases, is a labor-intensive chore most governments may not be able to afford. General practitioners and other health care agents are overwhelmed [178]. Whereas successful aging shall be addressed as a complex system, technology certainly will be a big part of the problem solution. Mobile communication, sensing and networking technologies are promising enabling tools for assisting the elderly. These technologies will allow seniors to manage diseases with less intervention and live independently for more time.

Personal monitoring and persuasion systems

Falls by elderly have become a major public health problem in modern society. It is both the first cause of accidental death, but also the beginning of the loss of autonomy from these people [142]. Presently, among various methods offered to assist seniors in case an emergency situation occurs, there are commercial systems based on badges or pendants, which can establish a wireless connection to a remote station installed inside the user's house. Once aware of the emergency, the remote station places a call to an operator that handles the situation.

Personal alarm systems are relatively popular in the United States and in the United Kingdom, where several companies offer this service [3, 149]. However, these systems require the intervention of the person (by pressing a button), which may be impractical, for instance, in the event of a serious fall. Automatic fall detection systems, such as Medical Alert, from Wellcore [150], and Lifeline, from Philips [171] try to solve this problem by automatically detecting falls and connecting the operator.

Because a fall is a complex event (there are several types of falls and each has particular features), fall detection algorithms can fail. These and other questions are addressed by several academic works. The studies [23, 24, 52, 116] present fall detection algorithms based on kinematic sensors, whereas Kangas et al. compare several fall detection algorithms in [115]. Another study [19] presents a fall detection system

based on the ZigBee protocol, whereas in [20] the authors presents a platform consisting of fall detection and physiologic sensors. In [117, 136, 173], the authors present algorithms capable of recognizing several human activities, including falls. An intelligent walking-aid is presented in [128]. It consists of a smart cane capable of detecting a fall and automatically sending an alarm to a smart phone through a Bluetooth connection.

Casas et al. [27] have developed an alarm and automatic fall detection system based on ZigBee that aims on monitoring independently-living seniors in an isolated Spanish village. By using ZigBee routers, the authors were able to establish a wireless network that covers the exterior spaces of the village and the interior of the houses. In the case of an emergency situation, users can press a button to alert a person in charge of handling alarms. Additionally, sensors can automatically detect falls and connect to the operator. As pointed out by the authors, the use of the ZigBee protocol endows the village with a wireless infrastructure that enables the integration with many other systems that might potentially improve the users' quality of life.

Elders should be stimulated to regularly exercise. Context-aware concepts were used to design Flowie, a virtual coach designed to persuade seniors to walk regularly [6]. The design concept combines a pedometer with wireless connectivity with a touch screen photo frame. The user interface runs on the digital photo frame with attractive animations and easy to visualize graphs, as those shown in Figure 33. The user interface was developed considering the opinion of several users who considered it encouraging and stimulating.



Figure 33 – Flowie user interface (adapted from [6]).

New wireless sensor technologies applied to home automation can replace commodity fire alarms and security systems with pervasive systems capable of recording routines and detecting a risky behavior. These systems will be comprised of several motes installed, for instance, on medication packaging, chairs, beds, glasses and clothes. In [78], the authors propose a system to determine a person's activities and

make decisions about his/her health status. The system is based on a wireless sensor network distributed throughout the house and a wearable sensor that interact to permit tracking the user.

Systems designed for assisted living facilities

The work developed by Fraile et al. [66] proposes a multi-agent system that uses smart wearable devices and mobile technology for the care of patients in a geriatric home care facility. The system is based on a ZigBee WSN and includes location and identification microchips installed in patient clothing and caregiver uniforms.

Two recent works propose fall prevention (instead of fall detection) systems designed to help assisted living facilities' care providers avoid patient falls or estimate the risk of a patient fall. In [122], the authors present GRiT Chair Alarm, which aims to prevent patient falls from chairs and wheelchairs by recognizing the gesture of a patient attempting to stand. In case the system detects an attempt to stand up, it tries to persuade the patient to remain seated. This system requires the use of capacitance sensors on the chair or wheel-chair and dispenses the use of wearable sensors. Najafi et al. propose a system capable of estimating the risk of a fall based on the characteristics of postural transition (sit-to-stand and stand-to-sit movements). Transitions and their duration are measured using a miniature gyroscope attached to the chest and a portable recorder placed on the waist [154].

The ZigBee protocol was used by Gamboa et al. as the basis of a vital signs monitoring, fall detection and tracking system specifically designed for patients in assisted living facilities [69]. The system is highly scalable and is based on wearable sensors that communicate directly with routers distributed all over the facility. Andrushevich et al. proposes three solutions for assisted living facilities while discusses the benefits of incorporating IPv6 support to the ZigBee protocol, as announced by the ZigBee Alliance [14].

Usability issues

Usability issues can discourage user from wearing devices. Steele et al. report findings of a qualitative study on the perception of elderly people on the use of current sensor networking technology for assistive care [190]. According to the authors, the two most important factors for elderly acceptance are cost and control. Various levels of

control were desired, but all participants expressed a certain desire of control over the system. The study in [181] presents two case studies and concludes that user-acceptance is inhibited if privacy is violated (e.g., by the use of video cameras) and that the perceived usefulness of the system is a crucial factor of acceptance.

Patients with dementia may not tolerate the use of wearable devices, creating additional constraints. For this group, it should be considered embedding sensors on devices and clothes such as belts, shoes and watches. An example is Actibelt, a system developed by the Sylvia Lawry Centre for Multiple Sclerosis Research to measure gait and activity parameters and detect falls. The patient unit shown in Figure 34 measures accelerations and communicates using a ZigBee link with a remote unit that collects and analyzes data [1].



Figure 34 – Actibelt patient unit embedded in a belt buckle (used with permission, adapted from [1]).

3.3.4 Wireless hospital monitoring systems

Wireless hospital monitoring systems, also called hospital telemetry systems, analyze physiological signals remotely collected and are mostly used to monitor patients' heart activity. One of their main advantages, as it is used in hospitals, is that it allows the patient to get up and move around, at least within the device's transmitting range, which is highly valuable for some patients, particularly those going through a physical rehabilitation process. It allows a patient whose recovery depends on movement to pursue that recovery effectively, while keeping the physicians ability to monitor his or her cardiac condition.

Commercial hospital telemetry systems operate either in specific telemetry bands or in one of the ISM bands [17]. ApexPro is a patient telemetry system manufactured by GE Healthcare [70]. The system consists of wireless patient units, Ethernet access points and network switches, application and database servers, and a monitoring station

for patient data visualization. An ApexPro patient unit (see Figure 35) contains a 6-lead ECG and an optional SpO₂ sensor. This unit operates in the frequency band from 420 MHz to 460 MHz [71], which is regulated and, consequently, less prone to interference than the 2450 MHz ISM band (used by WLANs). In Portugal, it may be configured to operate in one of the following frequency intervals defined by ANACOM, the Portuguese communications authority: a) 433.05 MHz – 434.79 MHz; b) 434.04 MHz – 434.79 MHz; or c) 458.1125 MHz – 458.1500 MHz [12]. Since 2009, the system included support for remote viewing of telemetry information on portable devices, such as PDAs and cell phones [70].



Figure 35 – ApexPro patient unit (adapted from [70]).

Instead of operating in one of the regulated frequency bands, which do not offer a large bandwidth, some vendors like Welch Allyn and Phillips Healthcare have opted for the use of ISM bands. Wireless patient monitoring devices from Welch Allyn, including the Micropaq wearable monitor (3-lead ECG, heart rate and SpO₂ sensors, shown in Figure 10, Chapter 1), were designed to share the hospital IEEE 802.11 a/b/g network with other hospital applications¹⁷. This design option lowers network management and infrastructure costs. On the other hand, it increases the complexity of network management because hospital technicians must ensure that the critical messages, such as alarms, are transmitted successfully [17].

IntelliVue, from Philips Healthcare, is a wireless patient monitoring system that also includes patient portable units. It operates in the ISM band, but it cannot share the spectrum with another IEEE 802.11 network. For instance, in case an enterprise IEEE 802.11b/g/n network is used in the same local area, IntelliVue's wireless infrastructure

¹⁷ Depending on the bandwidth availability, the system can run on a dedicated 802.11a network.

shall be set to use IEEE 802.11a. The system functionalities and network infrastructure are similar to ApexPro [170].

Aingeal is a hospital monitoring system presented by Intelesens in 2010 [98]. To the best of our knowledge, it is the first commercial remote monitoring system based on a wearable sensor designed to monitor ambulatory in-patients, the same scenario considered in our work. Patient respiration, ECG signals, skin temperature and activity (3-axis accelerometer) are constantly monitored. Any out of range heart or respiration rates, or irregular heart events are transmitted via Wi-Fi (IEEE 802.11 b/g) to a central processing station for further analysis by clinicians. The wearable sensor (see Figure 36) has autonomy of 24 hours.

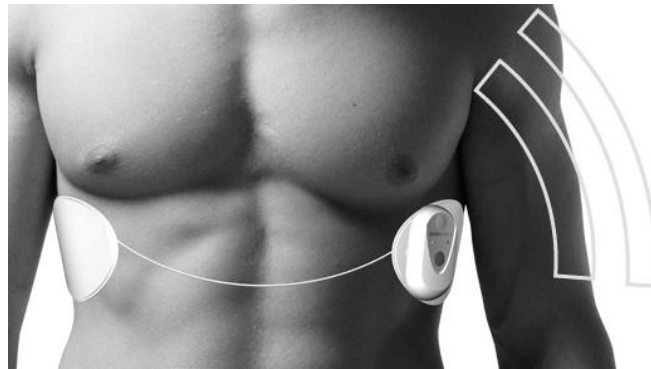


Figure 36 – Aingeal electrodes and clip-on patient device (adapted from [98]).

Few systems for in-patient remote monitoring were proposed by academia. One of the first systems based on WSN technologies to monitor patients in the hospital environment was presented by researchers from University of Texas (UT), in 2006, and was based on a ZigBee multi-hop network [80]. Patient wearable units consisted of a MicaZ mote [34] that was interfaced with a commercial blood pressure and heart-rate monitor. Routers were also based on MicaZ motes and were powered by an energy harvesting circuit that included solar panels. The system was tested in the laboratory environment using three patient units and no data was lost. In opposition to this study, our work considers the system scalability and restricts the network depth as it can seriously affect the amount of data that can be relayed.

In 2009, researchers from John Hopkins University temporarily deployed and analyzed the performance of MEDiSN (an acronym for Medical Sensing Application), a developed remote vital sign monitoring system, in the waiting areas of the emergency room (ER) of John Hopkins Hospital [123]. MEDiSN is comprised of a gateway, a

variable number of patient monitoring units and a wireless backbone of relay points. All wireless devices are based on Telos motes, which incorporate a CC2420 802.15.4 radio. The software running on devices uses TinyOS 2.0. Over the IEEE 802.15.4 protocol stack, the devices run the Collection Tree Protocol (CTP) provided by TinyOS [74]. Patient monitoring units collect ECG signals and oxygen saturation in blood (SpO_2) and generate a 111-byte data packet every 500 ms, almost the same traffic generated by our ECG sensors. The ER is covered by a WLAN operating on channels 1, 6 and 11. The system was tested while being used to monitor an average of three patient monitoring units. A backbone of eight relay points was employed. Two 24-hour tests were executed, one with the system operating on channel 22 and the other on channel 26. The authors report an end-to-end delivery ratio greater than 99.9% on both channels, despite the quality of the links, because CTP was able to find high quality links. Contrary to our study, where several causes for packet losses are found and explained, authors restrict their analysis to the quality of the wireless links and the performance of the routing protocol. Additionally, despite having patients use the developed system, they do not report any mobility issue.

The study presented on [31] is based on the same hardware and software components used by MEDiSN, except that authors have developed a CTP companion routing mechanism called Dynamic Relay Association Protocol (DRAP), which is deployed on patient nodes to discover and select relay nodes as the patient moves. During the system evaluation, sensor nodes were programmed to transmit a data packet every 30 s or 60 s. Data collected from 32 patients over a total of 31 days of monitoring shows that the median network and sensing reliabilities¹⁸ per patient were 99.92% and 80.55%, respectively. For one patient, network reliability was about 95% and, for other 6 patients, it was smaller than 98%, even considering that patient mobility was quite low. 80% and 90% of the services outages (period of time the service is unavailable) were less than 0.86 and 1.41 minutes, respectively. The median time between services outages was equal to 17.7 minutes. In this study, authors include neither the number of patients being concurrently monitored nor the size of the packets transmitted by sensors,

¹⁸ Sensing reliability, as used by the authors, refers to the correction of measurements made by sensors. As reported, the sensing reliability was negatively affected by incorrect measurements made by the SpO_2 sensor used, which was very susceptible to patient movements.

which prevents readers from making any judgment about the system reliability presented.

Researchers from MIT developed SMART (an acronym for Scalable Medical Alert and Response Technology) [45]. This system can monitor SpO₂, ECG, and the location of multiple patients. A commercial SpO₂ sensor was used, whereas the ECG sensor was developed as a Cricket mote daughter board. Figure 37 shows a patient unit. On the left side, the sensor box, ECG leads and the SpO₂ probe are shown; whereas, on the right side, it is shown the PDA responsible for connecting to the dedicated WLAN. Vital signs and location data are sent to a PDA using a serial wired connection. All information gathered is transmitted using Wi-Fi to a central unit with signal processing and data analysis capability. Alarms can be sent to caregivers' PDAs.




Figure 37 – The patient unit designed by MIT researchers under the SMART project (adapted from [36]).

In [36], the authors qualitatively evaluated a temporary deployment of SMART in the waiting areas of an emergency department of Brigham and Women's Hospital. The evaluation involved patients being monitored, where a maximum number of four patients were simultaneously monitored. The authors conclude that the system was well accepted by patients and caregivers. Additionally, they mention the occurrence of false alarms that should be further investigated. Similarly to our approach, raw ECG data are transmitted. In our approach, however, wired BANs are not used, but fully wireless sensors connected to an infrastructure distributed throughout the hospital.

3.4 Interviews

Several interviews were conducted with physicians and engineers to better understand what concerns them about present wireless vital signs monitoring systems and what can be done to improve these systems. The cardiologist that contributed to this work, Dr. Otílio Rodrigues, from Hospital Privado de Guimarães, was interviewed a

few times. Additionally, professionals from the Boston area were interviewed on November 2009, while some follow-ups occurred during 2010.

Dr. Warren Sandberg¹⁹ stated that, ideally, all in-patients should be monitored, which, according to him, will happen in the near future. He added that wireless technologies are very attractive because they free patients from their beds and eliminate wires which might transport resistant microorganisms that can survive cleaning processes used by hospitals. He also pointed out that the cost, including the cost of disposable items and maintenance costs, is a decisive factor for the adoption of new hospital systems.

Additionally, he pointed out that health monitoring systems should improve the context of alerts and trend reports. Also, these alerts and trends reports should consider physicians and nurses in a different way. While both alerts and trends concern stakeholders, physicians are more interested in the first ones. On the other hand, nurses, who spend longer periods watching over patients, carefully examine trends, in order to discuss with physicians the patients' response to treatments.

According to Dr Sandberg, alerts should be sent to physicians' personal devices. This opinion is shared by Dr. Henry Feldman²⁰, who stated that doctors should not carry several warning and communication devices. Instead, he suggested alerts should be directed to cellular phones instead of dedicated communication devices.

Dr. Sandberg suggested remote monitoring systems should be integrated with other information systems that store patient data (for instance, laboratory, pharmacy, and historical data) to anticipate life-threatening events. A prize-awarded investigation, jointly conducted by scientists from University of Coimbra and University of Utah, integrated, in one place, data supplied by ICU monitoring systems, historical data and laboratory results with events, rules, and data mining models to identify, in real-time, possible future risks to the patient [77]. A team led by Professor Guttag, from MIT, was also investigating the possibility of predicting life-threatening events, such as a cardiac

¹⁹ Dr. Warren Sandberg is an anesthesiologist. He is a professor at Harvard Medical School, and works at MGH and at the Center for Integration of Medicine and Innovative Technology (CIMIT). He is one of the leaders of the project "Operating Rooms of the Future" that is being developed by CIMIT.

²⁰ Dr. Henry Feldman is a hospitalist and works as the Chief Information Architect at Beth Israel Deaconess Medical Center.

failure. However, the development of this research was conditioned to the gathering of sufficient data in local ICUs.

Dr. Feldman referred that remote monitoring systems generate frequent false alarms and have low autonomy, deficiencies that must be corrected by manufacturers. Nancy Foster, one of the nurses who work in the intermediate care unit of Beth Israel Deaconess Medical Center, confirmed that medical telemetry units²¹ used in that service generate frequent false alarms and are bothersome to patients. In fact, during the interview, some false alarms occurred, which forced her to check patients and reset alarms. One of the patients whose monitor unit generated a false alarm complained that the monitoring unit was too heavy (it weighs about 0.5 kg) and that the wires were bothering him.

Regarding the biomedical signals to monitor, all doctors interviewed agreed that all primary vital signs (body temperature, heart rate, blood pressure, and respiration rate) were essential. Dr. Nathaniel Sims²² highlighted the importance of monitoring the respiration rate, a primary vital sign frequently neglected. Apart from the primary vital signs, Dr. Sandberg highlighted the importance of monitoring oxygenation, whereas Dr. Otilio Rodrigues considered imperative to monitor ECG because its analysis allows cardiologists to detect important conditions, such as damages in the heart muscle and arrhythmias.

William Driscoll, anesthesia information manager at MGH, affirmed that human errors are a major concern in hospitals. His team was adapting bed-side vital signs monitors to transmit measured vital signs to a handheld computer using a Bluetooth-based connection. Using the handheld computer, a nurse scans a 2D tag on a patient's file placed near the patient's bed to read the personal data. Then, the nurse acquires vital signs that are transferred to the handheld computer. Finally, the nurse confirms that patient data and vital signs measurements are correct, and then data are sent through the

²¹ The units are part of an IntelliVue system bought several years before the date of the interview. New IntelliVue models have an increased autonomy of 25 hours and, according to the manufacturer, new algorithms that provide reduced false alarm incidence.

²² Dr. Nathaniel Sims is an assistant Professor of Anesthesia at the Harvard Medical School. He also works at MGH and is a researcher at the Center for Integration of Medicine and Innovative Technology (CIMIT).

WLAN to a medical records repository. Another interesting project in early stages of development at MGH aimed at remotely programming infusion pumps. For this new project, they were considering the ZigBee protocol, which they were not familiar with.

Rickey Hampton, wireless communications manager at MGH and researcher at Partners Health care, stated that telemetry was vital at MGH, which had two telemetry systems. One of the systems operated at the WMTS bands and the other on the 2450 MHz ISM band (this one was being upgraded to operate on the 5800 MHz ISM band). By the time he was interviewed, MGH had 17 floors covered by telemetry systems and managers had plans to further increase the coverage area. According to him, systems that operate on WMTS bands are more reliable than technologies that share the ISM spectrum with enterprise WLAN. Nevertheless, he considers that several technologies are required to fulfill all patient monitoring needs, including low-power, low data-rate communication protocols.

When questioned about problems concerning monitoring systems, he stated that one of the main problems is poor systems integration. According to him, hospitals have many interoperability issues that must be solved.

3.5 Summary

Wireless medical devices consist of the medical devices themselves and the data transport, storage and analysis stages. Whereas medical devices should conform to strict regulation, other stages are not heavily regulated, though they should meet guidelines issued by standard and regulatory organisms, such as ISO, IEEE and local spectrum regulation entities.

Academic works and commercial products are reviewed and classified into three broad categories: personal health monitoring systems, monitoring systems for elderly care, and hospital monitoring systems. Most of the commercial products were released after this work started, and just a few are based on low power low data-rate protocols.

Finally, this chapter presents the views of physicians and informatics professionals on important aspects of remote patient monitoring systems. Preferably, all in-patients should be monitored. However, to achieve this ideal situation, systems should become more pervasive, better integrated, less expensive, and should generate more meaningful and reliable alerts and alarms.

Chapter 4

The developed monitoring system

As discussed in Section 3.4, despite the concerns of health care providers, few patients have their vital signs constantly monitored. HM4All, which stands for Health Monitoring for All, was developed to solve some of the issues that prevent vital signs monitoring systems from becoming widespread in hospitals, nursing homes and residential environments.

HM4All was designed to allow patients to be constantly monitored using unobtrusive and low cost sensors based on the ZigBee protocol, which, as presented in Chapter 2, was developed for low battery consumption, low cost and easy deployment. Contrary to several previous works, which have opted for an architecture based on body area networks (BANs) [133, 161], HM4All was developed considering that sensors connect directly to the ZigBee network infrastructure, thus avoiding the use of a patient unit that acts as a concentrator.

This chapter describes HM4All, the developed vital signs monitoring system, which comprises wireless medical sensors, ZigBee networking devices, a ZigBee-to-IP gateway and Web-based applications.

4.1 HM4All overview

HM4All high-level architecture is shown in Figure 38. Data generated by a wearable sensor are transported by ZigBee routers and coordinator to a ZigBee-to-IP gateway. Then, data are stored in the data server and made available to monitoring centers or wireless portable devices carried by health care providers.




Figure 38 – HM4All high-level system architecture.

Two wireless sensors were developed: a single channel (3-lead) electrocardiogram (ECG) sensor and an axillary temperature sensor. Both sensors are based on the JN5139-M00 wireless module (with integrated antenna) [106]. The amount of data generated by each sensor is shown in Table 14. Although ECG waveform and heart rate (HR) measurement are transmitted by the same sensor and data are aggregated, the amount of data generated by each function is shown separately. Battery level information is also sent by both sensors.

Table 14 – Amount of data generated by sensors.

Sensor	Sampling rate or period	Sampling size	Data generated
ECG (single channel, modified bipolar limb lead I)	200 Hz	12 bits per sample	2400 bps (raw data) or 1200 bps (2:1 compressed data)
Heart rate	10 s	1 byte	1 byte every 10 s
Axillary temperature	1 min	1 byte	1 byte every minute
Battery level	3 min	1 byte	1 byte every 3 minutes

ZigBee coordinators and routers are based on the JN5139-M02 high-power module [106] and use the same electronic printed circuit board (PCB) and case. These devices have two light emitter diodes (LEDs) that allow users to verify when they are powered on and when they are actively communicating. Additionally, they contain a reset button, a Universal Asynchronous Receiver / Transmitter (UART) serial communication

interface port connection and allow channel selection through a set of dip switches (channel selection can alternatively be done by code).

The ZigBee-to-IP gateway is a graphical user interface (GUI) based application developed in C# language that can execute in any computer with the .Net Framework installed. It validates and processes data frames received from a ZigBee coordinator and sends processed data to the Application Server application through an HTTP (Hypertext Transfer Protocol) connection. Additionally, it contains a user interface where sensor data are exhibited and recorded and connections are established and monitored.

The Application Server software comprises a Web server application based on Java servlets²³ and uses the Apache Web server [200] and the MySQL database [153]. This application collects sensor data from ZigBee-to-IP applications and sends data to remote clients, such as monitoring stations.

Applications that run on clients provide a user interface that allows care givers to execute the following tasks: a) visualize patients' vital signs in real-time; b) access historical patient data records; c) configure individual alarms; and d) execute management functions, such as patient registration, sensor insertion, and association between patients and sensors. Figure 39 shows HM4All conceptual diagram, where all remote connections are represented.

²³ Java servlets extend and enhance Web servers and can be thought of as an applet that runs on the server side. Servlets provide a component-based, platform-independent method for building Web-based applications, without the performance limitations of Common Gateway Interface (CGI) programs.

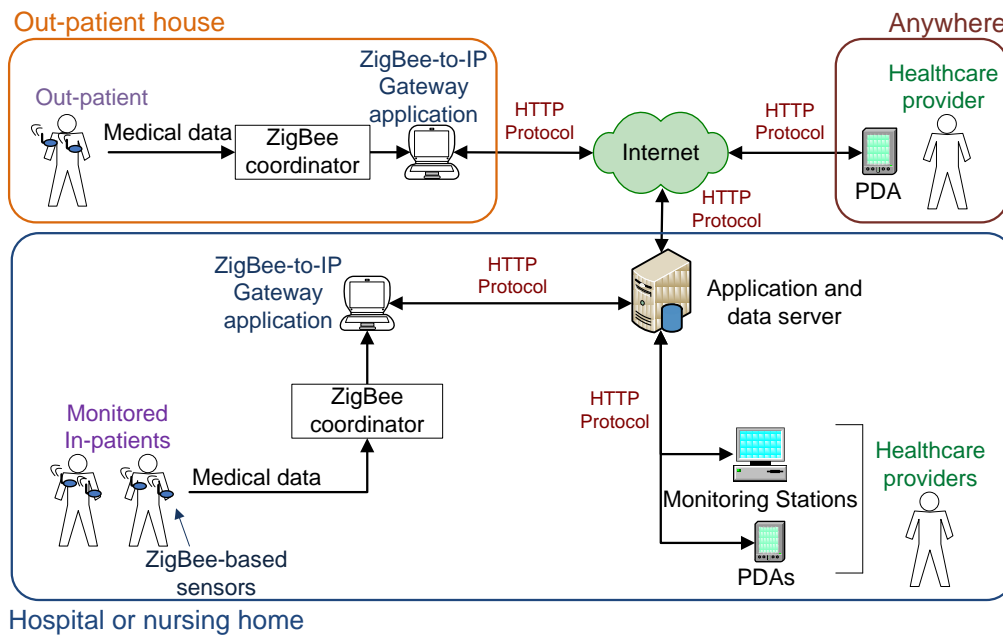


Figure 39 – HM4All conceptual diagram.

The Monitoring Station application consists of several Web pages. Among these, the most important is the page that presents patients' vital signs, as shown in Figure 40. Each small window presents real-time ECG waveform, heart rate and temperature values collected from sensors worn by patients. Additionally, it presents patient information (name, age, room and bed) and the battery level of each sensor. Finally, users can activate or deactivate alarms and access historical temperature and heart rate information displayed in graphical form.




Figure 40 – Monitoring page presenting data from several sensors (all data presented in this page was previously recorded for testing purposes).

The Application Server software and the Monitoring Station application were developed by Duarte Pereira, during his Master's Program [168].

4.2 Developed HM4All components

This section describes, in detail, HM4All components developed under this work. All components were developed from scratch, including the design of PCBs and their assembly. The design of sensors and network components includes the development of hardware and application software.

4.2.1 ECG sensor

The ECG sensor hardware was developed together with Ana Carolina Matos, during her Master Program [140]. It consists of a one-channel, three-lead ECG sensor designed to continuously monitor a modified projection of the lead I vector²⁴ and to determine the heart rate.

A simplified block diagram of the ECG sensor hardware is shown in Figure 41. The instrumentation amplifier (IA) amplifies the difference between the right arm and left arm signals while rejecting large values of common mode noise. A band-pass filter (BPF) and a 50 Hz notch filter follow and are used to attenuate unwanted frequency components not rejected by the IA. A non-inverting amplifier follows and is used to adjust the desired gain. Finally, a low-pass filter (LPF) is used to attenuate any DC level before the signal is buffered and sampled by the 12-bit analog-to-digital converter (ADC) internal to the JN5139-M00 module. The right-leg electrode is connected to the reference potential.

²⁴ Lead I or bipolar limb lead I corresponds to the voltage between the (positive) left arm (LA) electrode and (negative) right arm (RA) electrode. The developed ECG sensor uses non-standard electrode positioning on the torso. The resulting waveform is adequate for heart rate measurement, but cannot be used for diagnostic purposes.




Figure 41 – ECG sensor block diagram.

The distance between electrodes is 20 mm, which results in QRS amplitudes of, approximately, 0.33 mV^{25} . The ECG sensor has a passband between 0.5 Hz and 40 Hz and a gain of 4160 V/V, which results in an output signal that does not exceeds 1.37 V for normal ECG amplitudes. It is operated by a CR2 lithium battery with power holding capacity of 850 mAh and a nominal voltage of 3.0 V. The average current consumption is equal to 12.3 mA, resulting in a battery life of, approximately, 70 hours, which was confirmed experimentally. The schematics diagram of the ECG sensor is included in Appendix A.

The sensor's electronic PCB and its final prototype case are shown in Figure 42. The PCB was designed using the PADS tools [143]. The sensor case was designed using the SolidWorks computer aided design (CAD) software [187] and was manufactured in polyurethane, by SolidTech, a Portuguese company, using rapid prototyping techniques. The silicone molds used to manufacture the sensor cases are shown in Figure 43 (a), whereas a case just extracted from the mold is shown in Figure 43 (b).




Figure 42 – ECG sensor case and electronic PCB.

²⁵ The normal limb lead I QRS amplitude measured for electrodes positioned in the arms and hip or near the shoulders and hip is around 1 mV. However, if electrodes are placed closely on the torso, the QRS amplitude drops to approximately one third of this value.




Figure 43 – ECG sensor cases: a) the silicone mold used to produce the final prototype cases; and b) an ECG case just extracted from the mold.

An innovative characteristic of the developed ECG sensor is the absence of cables between the sensor and the disposable electrodes. It was achieved by connecting female ECG electrode connectors directly to the PCB, as shown in Figure 44.




Figure 44 – Rear side of the ECG sensor.

Real-time QRS detection was implemented using a slightly modified version of the algorithm developed by Pan and Tompkins [163]. It is a high-performance algorithm with a very low error rate of 0.68% on the MIT-BIH Arrhythmia database²⁶ [124]. It includes a pre-processing stage and rules for QRS detection. The pre-processor consists of a series of filters that perform low-pass (LP) and high-pass (HP) filtering, derivative, squaring and integration operations, as shown in Figure 45. Equations (4) to (8) presents the difference equations of each filter or operation, respectively, the low-pass filter, the high-pass filter, the differentiator, the squaring and the moving-window average integrator defined for a sampling frequency of 200 Hz [174].

²⁶ The MIT-BIH database is maintained by the Laboratory for Computational Physiology, which is part of the Harvard-MIT Division of Health Sciences and Technology. Several databases, including the Arrhythmia database, are accessible at <http://www.physionet.org/physiobank/database/#ecg>.




Figure 45 – Pan-Tompkins pre-processor.

$$y(nT) = 2y(nT - T) - y(nT - 2T) + \frac{1}{32} [x(nT) - 2x(nT - 6T) + x(nT - 12T)] \quad (4)$$

$$\begin{aligned} y(nT) &= y(nT - T) + x(nT) - x(nT - 32T) \\ p(n) &= x(nT - 16T) - \frac{1}{32} y(nT) \end{aligned} \quad (5)$$

$$y(nT) = \frac{1}{8} [2x(nT) + x(nT - T) - x(nT - 3T) - 2x(nT - 4T)] \quad (6)$$

$$y(nT) = [x(nT)]^2 \quad (7)$$

$$y(nT) = \frac{1}{N} [2x(nT - (N - 1)T) + x(nT - (N - 2)T) + \dots + x(nT)], N = 30 \quad (8)$$

Both the LP and HP filters were designed to maximize the QRS energy and contain only integer coefficients to reduce the computational complexity²⁷. The derivative filter provides QRS slope information and is nearly linear (and, consequently, nearly ideal) from DC to 30 Hz. After the derivative, a squaring operation emphasizes higher frequencies (especially from QRS complexes) and, consequently, reduces P and T waves. Finally, a moving-window integrator performs smoothing of the output of the squaring operation. The width of the window W is chosen not to merge the QRS complexes and T waves and to avoid the generation of several peaks from a single QRS [174].

Figure 46 shows the result of the pre-processing of an ECG waveform using the pre-processor implemented in MATLAB²⁸. The top-most trace shows an unfiltered ECG signal collected with the developed ECG sensor. The output of the filters results in a signal where the energy is concentrated on the band-pass frequencies. The output of

²⁷ More complex filters may improve the performance of the QRS detector. However, when Pan and Tompkins published their work (1985), it was critical to reduce the complexity of the real-time signal processing algorithms.

²⁸ MATLAB is a high-level technical computing language and interactive environment for algorithm development, data visualization, data analysis, and numerical computation. MATLAB is developed by MathWorks.

the derivative operation attenuates the P and T waves and enhances the QRS complex. The squared operator further enhances high frequencies, but the result still possesses multiple peaks for each QRS. Finally, the output of the integrator is a large pulse for each QRS complex. Smaller additional pulses result from large T waves. The shift between the actual QRS location and the corresponding output pulse is due to the cumulative delay of the various filters.




Figure 46 – Results of ECG pre-processing. From top to bottom: ECG waveform; output of the band-pass filtering; output of the squaring filter; and output of the integrating moving-window filter.

After pre-processing the ECG waveform, the algorithm searches for peaks which are compared to running estimates of signal (QRS) and noise peaks. Every time a new peak is detected, it is classified as signal or as noise by comparing the peak's amplitude with a threshold level, THRESHOLD I1. Then, THRESHOLD I1 and both the signal and the peak levels are updated. Another threshold, THRESHOLD I2, is maintained by the algorithm to be used in the search back procedure. A peak that does not exceed THRESHOLD I1 but exceeds THRESHOLD I2 ($\text{THRESHOLD I1} > \text{THRESHOLD I2}$) is considered for search back. Whenever a QRS is not detected for a time interval within $1.5 * \text{R-R intervals}$, the peak considered for search back is taken as a QRS.

The pre-processing algorithm used by the ECG sensor employs the smaller 80 ms integration window proposed by Urrusti and Tompkins [205] instead of the original 150 ms integration window proposed by Pan and Tompkins because, according to those authors, the smaller window slightly increases the QRS detector performance. Also, the threshold factor was increased to 0.6 to further enhance the detection performance in

presence of high T waves that result in unwanted peaks at the output of the pre-processor²⁹ [79, 163]. Finally, the original HP filter was substituted by a 13th-order HP equiripple finite impulse response (FIR) filter with stopband frequency at 0.5 Hz and bandpass frequency at, approximately, 5 Hz. Figure 47 shows the output of the integrator for an ECG waveform acquired using the developed sensor. The top graph shows the output produced by the original pre-processor, whereas the bottom graph shows the output of the pre-processor that employs the designed HP FIR filter. As shown, the designed FIR filter reduces the amplitude of unwanted pulses that correspond to T waves, which decreases the probability of false detections. The disadvantage of this filter is the increased computational complexity. The coefficients of the HP filter designed are presented in Table 15.

Table 15 – Designed HP FIR filter coefficients.

Filter coefficients						
a ₀	a ₁	a ₂	a ₃	a ₄	a ₅	a ₆
-0.1348973	-0.05155303	-0.06514797	-0.08656943	-0.1241208	-0.2102601	-0.6359727
a ₇	a ₈	a ₉	a ₁₀	a ₁₁	a ₁₂	a ₁₃
0.6359727	0.2102601	0.1241208	0.08656943	0.06514797	0.05155303	0.1348973




Figure 47 – Results of the Pan-Tompkins pre-processing using different HP filters. Top: original HP filter. Bottom: HP FIR filter designed to further attenuate P and T waves.

Before porting the QRS detection algorithm to the JN5139 module, it was coded in C# language. This version was employed to test the performance of the algorithm implementation, where ECG waveforms acquired with the developed sensor were used.

²⁹ The QRS detector was tested using waveforms collected by the developed sensor. It was noticed that it achieves a better performance if the threshold factor is increased to 0.6 instead of 0.25, as originally proposed by Pan and Tompkins, and further increased to 0.3125 by Hamilton.

Figure 48 shows the C# implementation user interface. It processes an ECG signal from a text file and presents the number of detected peaks and the mean, maximum and minimum heart rate values. The mean heart rate value is obtained considering the average of all R-R intervals, whereas the maximum and minimum heart rate values are obtained from the shorter and longer R-R intervals, respectively. Additionally, it displays information needed to track the algorithm execution and generates a file that contains either the output of a pre-processing filter or a list of samples corresponding to location of the R waves detected.

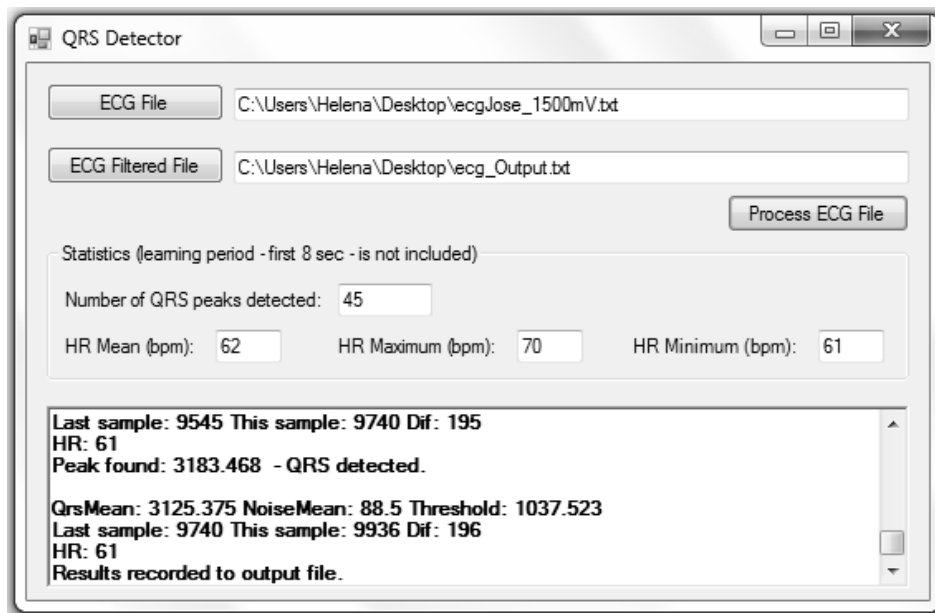


Figure 48 – User interface of the QRS detector implemented in C# language.

The ECG sensor application that runs in the target board was developed in C language and contains four modules: EcgSensorED, PreProc, QrsDet and Printf. The EcgSensorED is the main module, responsible for wireless communication and ECG data sampling and analysis. The other modules execute ECG signal pre-processing, QRS detection and communication tasks using the UART serial interface, respectively.

The JZA_vAppEventHandler is the most important function of the EcgSensorED module³⁰. Figure 49 shows a flowchart that describes the tasks executed by this function. Whenever JZA_vAppEventHandler executes, it verifies if one or more interruptions have occurred and treats the older one. Every 5 ms (sampling period), a

³⁰ As discussed in Section 2.5, it is the default user thread and is constantly called by the JN5139 module's scheduler.

timer generates an interruption that is captured by JZA_vAppEventHandler. In that case, it starts the ADC internal to the JN5139-M00 module to sample at ADC1 input and exits. When the ADC is ready to sample data, it generates an interruption. When JZA_vAppEventHandler detects this interruption, it calls a function that samples the ECG signal and stores the sampled value. After that, JZA_vAppEventHandler increments the variable u8EcgMeasurements, which is used to count the number of ECG measurements already done. If the value of u8EcgMeasurements is equal to the number of samples to be transmitted, the static boolean variable bSendECG is set and the transceiver is switched on. Additionally, if it is time to send a new heart rate value the static boolean variable bSendHR is set.




Figure 49 – JZA_vAppEventHandler flowchart.

If no interruption must be served by JZA_vAppEventHandler, it verifies if bSendECG is set. In that case, if the transceiver is already powered on, it sets off bSendECG. Then, depending on the value of bSendHR, it calls a function that requests the transfer of a data frame with the ECG samples or with the ECG samples plus the heart rate value.

Just after sampling the ECG signal, the ECG sensor application calls functions that filter the signal and search for a new QRS complex. If a QRS complex is detected, a new R-R interval value is determined and stored in a local variable. Heart rate values are determined as the arithmetic mean of a defined number of successive R-R intervals. Presently, the heart rate is determined from the previous ten successive R-R intervals.

The occurrence of tachycardia, bradycardia or asystole³¹ is verified by the sensor. The function responsible for asystole detection maintains a static local variable that is increased every time a new ECG sample is taken. This variable is initialized every time a new QRS is detected. However, if no cardiac activity is detected for 3 seconds, an asystole event is detected. Additionally, the ECG sensor detects the occurrence of background arrhythmia³².

The ECG signal is sampled at a frequency of 200 Hz. However, the sensor transmits 2:1 compressed ECG data. Data compression is done using the turning-point algorithm [203]. This algorithm is based on the concept that ECG signals are normally oversampled to easily visualize the higher-frequency attributes of the QRS complex. It provides a way to reduce the effective sampling rate by half to 100 Hz by selectively selecting important signal points.

4.3 Temperature sensor

The temperature sensor prototype is shown in Figure 50 (a). The case was designed using Solid Works CAD and prototyped using stereolithography. As shown in Figure 50 (b), it is held in place by a fabric arm band. The thermistor tip of the temperature probe is placed in the axillary region.

³¹ As presented in Section 1.2, on adults, a heart rate above 100 bpm is referred to as tachycardia, whereas heart rates below 60 bpm are referred to as bradycardia. Asystole refers to the state of no cardiac electrical activity.

³² As discussed in Section 1.2, a background arrhythmia is detected by the ECG sensor if, during a one minute period, the heart rate experiences a positive and a negative variation of 20% of its average value. In the current implementation, tachycardia, bradycardia and the absence of cardiac electrical activity are detected by the Monitoring Station application (presently, background arrhythmias are not detected). However, the system can be easily modified to identify and present alarms sent by the ECG sensor.

The temperature sensor complies with the technical characteristics listed in Table 16. The axillary temperature was chosen because, unlike other sites for which the temperature is correlated to the core temperature of the body, the placement of a probe in the axillary region causes only minor discomfort to the patient.




Figure 50 – Temperature sensor prototype: (a) assembled with arm band and probe and (b) being worn.

Table 16 – Temperature sensor technical characteristics.

Parameter	Value
Measuring site	Axillary
Temperature range	34 °C – 42 °C
Precision	± 0.2 °C
Resolution	0.1 °C

The temperature sensor uses a medical thermistor probe. Thermistors are one of the most accurate types of temperature transducers despite their non-linear resistance versus temperature (R-T) characteristic. This characteristic is approximately described by Equations (9) and (10), which are known as the Steinhart-Hart equations [191]. Typical values of the constants A_0 , A_1 , A_2 , B_0 , B_1 , and B_3 are supplied by manufacturers or are obtained by the user through calibration.

$$\frac{1}{T} = A_0 + A_1 \ln(R) + A_2 (\ln(R))^3 \quad (9)$$

$$\ln(R) = B_0 + \frac{B_1}{T} + \frac{B_2}{T} \quad (10)$$

If an ordinary thermistor is used in a medical device, calibration is required for each thermistor used. In order to avoid it, pre-calibrated medical thermistor probes with well-defined R-T characteristics are used in monitoring devices, such as bedside monitors.

This ensures interchangeability, repeatability and accuracy equal or better than ± 0.1 °C between 34 °C and 42 °C [218].

Besides non-linearity, another important concern is self-heating. Being a passive component, a thermistor must be excited to determine its resistance and, consequently, the temperature it is being subjected to. However, the excitation current level must be limited to avoid self-heating, which causes the thermistor resistance to decrease and give a higher reading than the actual temperature. As a rule of thumb, manufacturers recommend the excitation current does not exceed 100 μ A [16] or the dissipated power does not exceed 100 μ W [141], regardless of the value of the thermistor resistance.

Two Exacon 400-series interchangeable skin temperature medical probes were used: an adult D-S18 and an infant D-S10 [54]. Both probes have a resistance of 2252 Ω at 25 °C and are factory-calibrated to conform to a standard R-T characteristic with a deviation less than 0.02 °C. The probes' R-T characteristic is defined by the Steinhart-Hart constants supplied by the manufacturer and listed in Table 17. The R-T curve obtained from the constants supplied by the manufacturer is shown in Figure 51.

Table 17 – Exacon 400-series probes Steinhart-Hart constants.

Constant	Value
A_0	1470.196×10^{-6}
A_1	237.7907×10^{-6}
A_2	104.6897×10^{-9}




Figure 51 – Exacon 400-series probes R-T characteristic.

The circuit diagram of the electronic PCB is included in Appendix B, whereas a simplified block diagram of the sensor is shown in Figure 52. It consists of three modules: the power supply, the analog front-end and the wireless module [21].




Figure 52 – Simplified block diagram of the temperature sensor.

Figure 53 presents the schematic diagram of the power supply module. The sensor is powered by a CR2540 coin battery with power holding capacity of 610 mAh and a nominal voltage of 3.0 V. The average current consumption of the sensor is equal to 107 μ A, which corresponds to a battery lifetime of 241 days. The energy consumption and the battery lifetime calculations are presented in Appendix C.

The battery voltage, V_BATT, is permanently applied to the wireless module, but is only available to the front-end when a measurement should be done. The voltage VCC is controlled by the wireless module by driving the base terminal of Q1, a PMBT3906 switching transistor, through the digital input-output (DIO) line DIO1. When it is necessary to measure temperature, DIO1 is set OFF, Q1 saturates, and the voltage VCC is applied to the analog front-end. During saturation, the collector-emitter voltage is typically under 100 mV and does not exceed 250 mV.




Figure 53 – Schematics diagram of the power supply of the temperature sensor.

The voltage across the thermistor probe is measured by the analog-to-digital converter (ADC) internal to the JN5139-M00 module, which has 12-bit resolution. If the input range of the ADC is set to its maximum value of 2.4 V, the voltage value of each bit is equal to 0.59 mV/bit ($2.4 / (2^{12} - 1)$). In order to correctly measure temperatures in the range between 34 °C and 42 °C, the thermometer was designed to measure 0.1 °C below and above these minimum and maximum temperature values. For a resolution of 0.1 °C, a total of 83 possible temperature values should be measured,

consequently, 49 ADC codes are available per 0.1°C (2^{12} codes / 83 intervals), which is adequate because it exceeds the precision of the temperature probes acquired ($\pm 0.02^\circ\text{C}$).

Figure 54 shows the schematic diagram of the front-end. The thermistor probe is connected to J3 connector. REF200 [197] contains two $100\ \mu\text{A}$ current sources which are used to excite the thermistor and a reference load comprised by the fixed resistor R3 and the variable resistor R9. Both U2 and U4 are OPA336 operational amplifiers [196]. These components operate on a single supply as low as $2.1\ \text{V}$. The output is rail-to-rail and swings to within $3\ \text{mV}$ of the supplies with high impedance loads.




Figure 54 – Schematics diagram of the front-end of the temperature sensor.

The voltage developed between the thermistor terminals, V_{TH} , is filtered and is input to the non-inverting input of U4. The voltage developed across the reference load, V_{REF} , is buffered by U2 and is applied to the inverting input of U4. The output of U4, V_{OUT} , is then applied to the ADC1 input of the ADC. The output voltage of the non-inverting amplifier is defined by Equation (11), where the gain G is defined by Equation (12).

$$V_{OUT} = [V_{TH} * (1 + G)] - [V_{REF} * G] \quad (11)$$

$$G = (R6 + R10)/R5 \quad (12)$$

The thermistor probe resistance, R_{TH} , for temperatures equal to 33.9°C and 42.1°C are equal to $1539.6\ \Omega$ and $1102.9\ \Omega$, respectively. The values of the G and V_{REF} are obtained from the values of V_{OUT} and V_{TH} . For a temperature equal to 33.9°C , V_{TH} is

equal to 0.1540 V and V_{OUT} assumes a value of 2.4 V. Likewise, for a temperature of 42.1 °C, V_{TH} is equal to 0.1103 V and V_{OUT} is equal to 0. Substituting these values into Equation (11), Equations (13) and (14) are obtained. Solving both equations, it is possible to obtain $V_{REF} = 0.1123$ V and $G = 53.92$.

$$[0.1103 * (1 + G)] - [V_{REF} * G] = 0 \quad (13)$$

$$[0.1540 * (1 + G)] - [V_{REF} * G] = 2.4 \quad (14)$$

Using the V_{REF} value just determined, then the sum of R3 and R9 is found equal to 1123 Ω (approximately, the same value of R_{TH} for $T = 42.1$ °C). Considering R5 equal to 1 k Ω and the value of the gain G, the value of the sum of R6 and R10 is found equal to 53.92 k Ω . Then, choosing 1% values for all fixed resistors and an excursion of, approximately, $\pm 5\%$ for all variable resistors, it is possible to determine the value of all other fixed and variable resistors, as shown in Figure 54 and Appendix B.

A first-order LP filter is created by R6, R10 and C5. The value of C5 is determined using Equation (15) for a cutoff frequency F_c of 250 Hz. R4 and C2 create a first-order LP filter designed to remove high-frequency components generated in the temperature probe. The values of the R4 and C2 are also determined for the same cutoff frequency. C3 and C4 capacitors are used for decoupling.

$$F_c = \frac{1}{2\pi RC} \quad (15)$$

The temperature sensor application runs the state machine depicted in Figure 55. When the sensor is powered on, the application assumes the state INITIALIZING, initializes hardware peripherals and configures the stack. When initialization is complete, it assumes the state JOINING, starts the network association process and switches on the LED by setting off the digital input-output (DIO) output DIO16 (see Appendix B). After joining the network, it switches off the LED by setting on DIO16 output and moves to the ASLEEP state. Then, it starts a wake-up timer to generate an interruption when a certain period of time have elapsed and sends the device to sleep. The sleeping period corresponds to the thermistor probe settling time and is equal to 30 seconds. When the active wake-up timer fires, the application goes to the WAITING_TO_READ_TEMP state, switches on the front-end by setting off the DIO output DIO1 and configures the ADC internal to the JN5139-M00 module to start

sampling at the ADC1 input³³. When it captures the interruption generated by the ADC, it reads the ADC conversion voltage, switches off the front-end by setting on the DIO output DIO1, starts the ADC to sample the battery voltage and assumes the state `WAITING_TO_READ_BATT`. When the interruption generated by the ADC is captured, the application reads the ADC conversion result and requests the lower layers to send both ADC conversion results to its parent. Then, it goes to the state `WAITING_TO_SEND_DATA`. After notified of the result of its request to transmit, the application sets a wake-up timer, assumes the state `ASLEEP` and sends the device to sleep. When the application captures the interruption generated by the active wake-up timer, it starts again the ADC to sample data in the ADC1 input, switches on the front-end and assumes, once more, the state `WAITING_TO_READ_TEMP`.

In case the device is unable to communicate with its parent (if three consecutive transmission attempts fail), the application assumes the device has moved away from its parent. Then, the application switches ON LED1 to indicate communication problems, assumes the state `JOINING` and tries to find a new parent. The application will attempt to reassociate repeatedly until successful.

³³ As shown in Figure 54, the output of the front-end is connected to the ADC1 input.

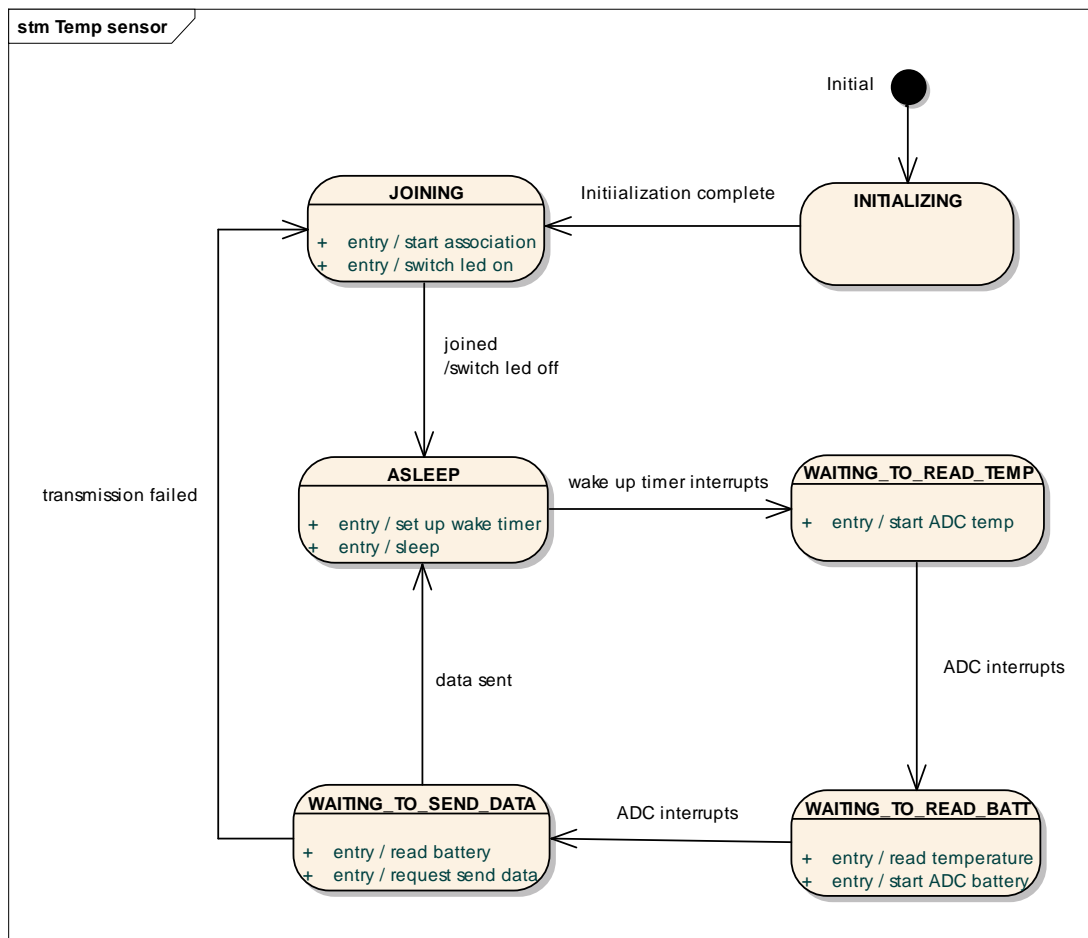


Figure 55 – State diagram of the temperature sensor application.

The front-end was tested using a variable resistor in substitution of the thermistor probe. Several values of resistance were adjusted and the voltage at pin 1 of U4 was measured. The module of the maximum difference found was equal to 29 mV, which corresponds to a deviation of 0.1 °C in the temperature measurement³⁴. The precision of the ADC conversion was also tested comparing the conversion values with pre-adjusted voltage values input to its ADC1 input. A maximum deviation of 5 mV was found for an input range between 0 to 2.4 V³⁵.

³⁴ Resistance and DC voltage were measured using a Fluke 115 multimeter. The resistance scales used have an accuracy of $\pm 0.9\% + 1$ ($\pm [(\% \text{ of reading}) + \text{counts}]$). The DC millivolts scale used has an accuracy of $\pm 0.5\% + 2$.

³⁵ The measured deviation is smaller than the specified typical values of the offset and gain errors, which are both equal to 20 mV. Consequently, it is possible that other modules present a worst performance.

The accuracy of the temperature sensor was tested comparing the sensor readings with the readings supplied by the GE Dinamap V100 vital signs monitor [72], which has an accuracy of ± 0.1 °C. The module of the maximum deviation observed when the D-S18 probe was used was equal to 0.3 °C, which is within the specified precision of the sensor.

4.3.1 Coordinators and routers

Coordinators and routers share the same hardware, but have different application programs. The circuit diagram of the electronic PCB is included in Appendix D. A mounted coordinator/router is shown in Figure 56. The antenna and the AC adapter are connected to the device.



Figure 56 – Coordinator/router (power on and activity LEDs not shown).

A USB-to-serial cable, such as the TTL-232R-3V3 [67], from FTDI, is required to program the devices (coordinators, routers and sensors) and to connect coordinators to computers that run the ZigBee-to-IP gateway application. The TTL-232R-3V3 cable operates at +3.3VDC levels (signals only, VCC is still +5 VDC on both ends).

Coordinators and routers can be powered by an external +5 VDC input applied to J1 connector or, in case of coordinators, by the +5 VDC input from the USB-to-serial cable connected to J2. Jumper J3 is used to select the power source. A RF Solutions ANT-24G-HL90-SMA 2.4 GHz compact helical antenna is used, however, several alternatives are enumerated by Jennic in the JN5139 module datasheet [110]. The EIA RS-232 UART serial interface operates in asynchronous mode and can be programmed to operate at a baud rate of up to 460,800 baud.

Switch SW1 is used to turn *on* or *off* the device. Switch SW2 is used to reset the device, whereas switch SW3 is used to program it. LED1 (ACTIVITY) blinks to indicate that the device is actively communicating, whereas LED2 (POWER ON) is permanently lit when the device is connected to a power supply.

Four slide dip-switches can be used to select the ZigBee channel to be used by the device. Each switch selects a binary value, being the right-most one the less significant bit. The decimal values 0 to 15 correspond to ZigBee channels 11 to 26. Alternatively, the operating ZigBee channel can be selected by the application program. A coordinator/router electronic PCB is shown in Figure 57.




Figure 57 – Coordinator/router electronic PCB: (a) top side and (b) rear side.

The router application is quite simple. At first, the device is initialized and then it searches for a specific network to join. After associated, it allows other devices to join and relays messages on behalf of its children.

The coordinator application program initially starts a network. Devices can then join and send data. When a data packet is received, its payload is framed before sending it through the UART serial interface to ensure resynchronization in case one or more bytes are lost. The framing procedure includes start and end flags, character count and byte stuffing, as shown in Figure 58.




Figure 58 – Data framing structure.

The start flag, the byte 0xA5, is followed by a byte that represents the size of the frame. The frame size includes data bytes, stuffing bytes and the end flag, the byte 0x5A. If one of the data bytes is equal to the start flag, an ESC flag (0x00) is inserted

just after it to inform the ZigBee-to-IP gateway application that the previous byte is not a start flag. Data bytes include two bytes that identify the sensor that collected the data and two bytes that classify the data into four possible categories: ECG, heart rate, temperature, and battery level.

4.3.2 ZigBee-to-IP gateway

This ZigBee-to-IP gateway application is responsible for processing data received from a ZigBee coordinator and sending data to the Application Server software. Additionally, it provides an interface where users can monitor the network activity, detect packet losses and record data. This application was developed in C# and contains seven classes, as shown in Figure 59.

Whenever data are received, the `portDataReceived` method from `SerialPortManager` class is invoked. All bytes stored in the `SerialPort` receive buffer³⁶ are read by this method and then each read byte is processed by the `processByte` method. This method runs the state machine depicted in Figure 60 to remove framing (START and END flags) and stuffing (ESC flag) bytes included by the coordinator application and to resynchronize in case a communication is detected.

The first time the `processByte` method executes, the state machine assumes the initial state `INIT_FRAME`. If the first byte received is a START flag, the state machine goes to the `READ_SIZE` state. Otherwise, it goes to the `SYNC` state and remains in this state until it receives a START flag.

³⁶ The `SerialPort` class belongs to the C# language and represents a serial port resource. The receive buffer includes the serial driver's receive buffer as well as internal buffering in the `SerialPort` object itself.




Figure 59 – Class diagram of the ZigBee-to-IP application.

While in the READ_SIZE state, the state machine goes to the READ_MSG1 state if it receives a valid message size. Any valid message should have a message size equal or greater than the minimum message size defined by the MIN_SIZE constant. Then, the state machine initializes two variables, dataSize and bytesRead. On the other hand, if the value of the size byte is invalid, it goes to the SYNC state.

If a START flag is received while processing message bytes, the state machine transitions to the READ_MSG2 state. Then, if an ESC flag is received, the state machine recognizes the START flag as a valid data byte, removes the ESC flag and goes back to the READ_MSG1 state. Otherwise, it goes to the SYNC state.

After reading all expected message bytes, the state machine transitions to the END_FRAME state. If the next received byte is an END flag, the message is stored and the state machine transitions to its initial state. Otherwise, it transitions to the SYNC state.




Figure 60 – State diagram of the procedure used to process bytes received by the serial port interface.

The SyncBuffer class implements a synchronized buffer³⁷ to which access is controlled by the Available_Slots semaphore. This class provides two access methods: insert and remove. The method insert is used by the class SerialPortManager to store a valid received message, whereas the method remove is used by the class Domain to retrieve a stored message. The critical sections from insert and remove methods are mutually exclusive, that is, these sections cannot be executed simultaneously.

The MainForm class is defined as a Form class and implements methods that provide services and display information to the user. The method connectMenuClick uses the method connectSP, from Domain class, to start a serial port connection. Once called, this method invokes the method connect, from SerialPortManager class, to start the connection. If a serial connection is successfully established, the method connectSP creates and starts the thread processThread. This thread runs the method process, also from Domain class. This method removes one message at a time from the synchronized buffer. Depending on the message data type (ECG, heart rate, temperature or battery

³⁷ A synchronized buffer is shared by a producer and a consumer thread. The access to this buffer is controlled to avoid simultaneous operations.

level) different methods are called to process it: processECG, processHR, processTemp or processBatt.

For each active sensor, the class Domain maintains an object from the class SensorList. Objects from this class are binding connected to objects from the class SensorData. Information stored by SensorData objects includes, for instance, the sequence number of the last message transmitted by a sensor, which can be used to count the number of messages lost by each sensor.

After processing a message, the relevant information is passed by to the sendData method, from Domain class, to the monitoring application using an HTTP connection. The post method is used and the string included in the request is constructed concatenating pre-defined characters ('&' and '=') and labels with corresponding data (previously converted to string format). Data description and values are presented in Table 18, whereas the code snippet presented in Figure 61 illustrates the concatenation procedure executed by the sendData method.

Table 18 – Data sent to the Monitoring station application. Data are converted to string format before being passed to the method that constructs the data string sent along with the HTTP post request.

Label	Description	Values
sensorType	Battery level	"01"
	ECG data	"02"
	Heart rate	"04"
	Temperature	"05"
sensorID	Sensor identification code	0x0000 - 0xffff (2 bytes)
timestamp	Time difference, in milliseconds, between the local time obtained when the message is received by the class SerialPortManager and a reference time (1 st January 1970, 1AM)	64-bit word
data	Battery level	0 – 3, where 0 corresponds to minimum battery level.
	ECG samples	50 floating-point values separated by “#” characters. ECG samples range from 0.6 to 2.4.
	Heart rate value	1 integer value (0 – 240)
	Temperature value	1 floating-point value (34.0 – 42.0)
seqNr	Packet sequence number generated by the sensor’s Application layer	0x0000 - 0xffff (2 bytes)

```
// Create POST data
string postData = "sensorType=" + sType;
postData = postData + "&sensorID=" + sID;
postData = postData + "&timestamp=" + sTime;
postData = postData + "&data=" + sData;
postData = postData + "&seqnr=" + sSNr;
```

Figure 61 – Code snippet used to construct the string used to send data along with the post request.

The ZigBee-to-IP gateway user interface is shown in Figure 62. The main menu contains four items: Serial Port, Data Server, Record and Exit. The Serial Port option allows the user to configure the serial port (COM port number, baud rate, number of data bits, parity and the number of stop bits), start a connection or disconnect. The Data Server option allows the user to set the Internet Protocol (IP) address and port used to connect to the remote computer that runs the Application Server application. The Record option can be used to define a folder where the application creates log files used to record sensor data received from each sensor. Alternatively, it is possible to record data sent from one specific sensor. It is also possible to start, interrupt and resume data recording. Finally, the option Exit is used to end the application.

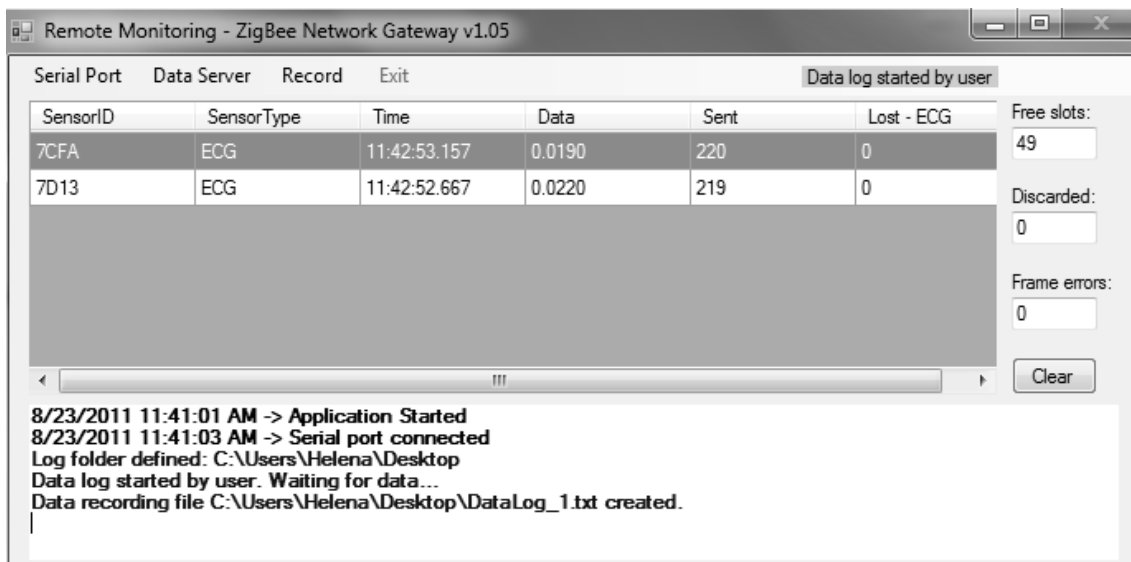


Figure 62 – ZigBee-to-IP gateway user interface.

The top window registers information about sensors and displays the value of the vital sign just measured. The information is displayed in tabular format where the columns contain the following information (from left to right): sensor identification number; the sensor type or information, where four options are defined: ECG, HR (heart rate), TEMP (temperature) or BATT (battery voltage); the time the last data packet arrived; the data (ECG sample is shown in millivolts, heart rate in beats per minute, temperature in Celsius degrees and battery level in volts; in case of ECG data, only the value of the first received sample is shown), the number of successfully received data packets; and the number of lost packets (presently, this functionality is implemented only for ECG packets).

On the right side of the user interface, it is shown the number of free slots in the synchronized buffer and the number of discarded packets³⁸. Additionally, the user interface presents the number of frames received with errors. The button Clear is used to clear all information presented to the user.

The bottom window presents log messages that correspond to events and errors. The events include the date and time the application was initiated, the date and time the serial connection was opened and closed and information regarding data logging actions. Errors include failure to open a serial port connection or error during an attempt to establish a connection with the remote application.

4.3.3 User acceptance

Questionnaires were used to gather information to evaluate the developed monitoring system. Information was collected in two sessions. The first session occurred on October 2009, when five patients admitted to Hospital Privado de Guimarães wore both the ECG and the temperature sensors. These patients and two nurses who accompanied the tests filled questionnaires. Additionally, on September 2011, eight health care professionals from Hospital Privado de Guimarães, including six nurses, one doctor and one biomedical engineer, anonymously answered a new questionnaire designed to evaluate the system. The small number of questionnaires is insufficient for statistical significance of results, but the goal of this task was merely to identify problems that were consensual among patients and care givers. The questionnaires are included in Appendix E.

The graph shown in Figure 63 presents the scaled responses to questions answered by patients that evaluated the ECG sensor. All patients recognized that the sensor was lightweight, unobtrusive and comfortable. However, one patient was unsure if the sensor was portable or small. Patients also answered yes/no questions. All patients stated that the ECG sensor did not drop during the use; that they did not have any of fear

³⁸ The synchronized buffer can fill if the gateway is unable to send processed data to the Application Server software (e.g., in case the connection is broken). If the buffer is full and a new message should be stored, the oldest stored message is removed from the buffer and the new one is included. In that case, the gateway application increases by one unit the number of discarded messages.

to use the sensor; and that they considered the ECG sensor could contribute positively to the care provided by the hospital.




Figure 63 – Evaluation of the ECG sensor made by patients.

The graph shown in Figure 64 presents the scaled responses of health care professionals about the ECG sensor. As shown, most of the health care professionals agreed that the sensor is easy to put and remove from patients. Also, most agree that it is easy to find the correct sensor positioning, clean the sensor and change the battery. However, four respondents were unsure if it is easy to switch the sensor on or off, whereas three were unsure if it was easily cleaned.




Figure 64 – Evaluation of the ECG sensor made by health care professionals.

By analyzing the data gathered, it is possible to conclude that it is important to reduce the size of the ECG sensor and ease the way it is switched on and off and cleaned. The size of the ECG sensor can be decreased by reducing the size of the electronic PCB and by substituting the CR2 lithium battery by a small rechargeable battery. By increasing the number of copper layers and using very small surface mount technology (SMT) components (most of the components used have a 1206 case style), it is possible to reduce the size of the electronic PCB. The LED hole was appointed by one respondent as a difficult area to clean, whereas some respondents have pointed out that the on-off switch is not easily accessible. These points should be improved in a new sensor version.

Figure 65 presents the scaled responses to questions answered by patients that evaluated the temperature sensor. Despite the overall good impression, it is clear that its size should be reduced. Patients who answered to yes/no questions said that the temperature sensor did not drop any time during use; that they had no fear to use the sensor; and that they considered that it could contribute positively to the care provided by the hospital.




Figure 65 – Evaluation of the temperature sensor made by patients.

Figure 66 presents the result of the evaluation of the temperature sensor made by health care professionals. Most of them acknowledged that it is easy to switch on and off and that it is easy to put and remove from a patient. However, three respondents were unsure if it is easy to change the battery and if it is easy to clean the sensor.




Figure 66 – Evaluation of the temperature sensor made by health care professionals.

From the analysis of the questionnaires used to evaluate the temperature sensor, it is possible to conclude that its size should be reduced and that it is necessary to improve the way it is cleaned and the way the battery is exchanged. As well as for the ECG sensor, the size of the temperature sensor's PCB could be reduced by using more copper layers and smaller components. The sensor case should be smoothed and made water-resistant to ease the way the sensor is cleaned. The way the battery is exchanged seems simple, though it could not be clearly demonstrated because small parts of the prototype sensor box, including the retaining clips of the battery cover, could not be correctly manufactured.

Figure 67 ranks the opinions of eight healthcare professionals about general questions. All of the respondents considered the developed system useful. However, one respondent considered it is useless out of the ICU and intermediate cardiac unit. Additionally, two respondents were not sure of the importance of a remote vital signs monitoring in areas such as emergency and other waiting areas. Respondents also included general comments about the system. Several respondents considered it can improve in-patients' comfort and streamline nurses' work. Additionally, two respondents considered it would be an ideal solution for palliative care and other hospital environments where patients' vital signs should be monitored quite often.




Figure 67 – System and concept evaluation.

Figure 68 ranks the functions health care professionals considered that should be included in the developed system. The most cited functions are oxygen saturation and the blood pressure monitoring. Other missing functions included by respondents are: respiration rate and glucose monitoring; issuance of a patient report which lists the values of the vital signs monitored; issuance of a trend report which contains all vital signs monitored in graphical form; and the ability to simultaneously monitor all in-patients admitted to one floor.




Figure 68 – Missing functions.

Figure 69 list the negative points and the improvement suggestions proposed by the health care professionals interviewed. Four respondents suggested reducing the dimensions of the sensors. Additional suggestions pointed out by one respondent were: including monitoring blood pressure and oxygen saturation; improving the temperature sensor exterior design and simplifying the access to historical recordings.

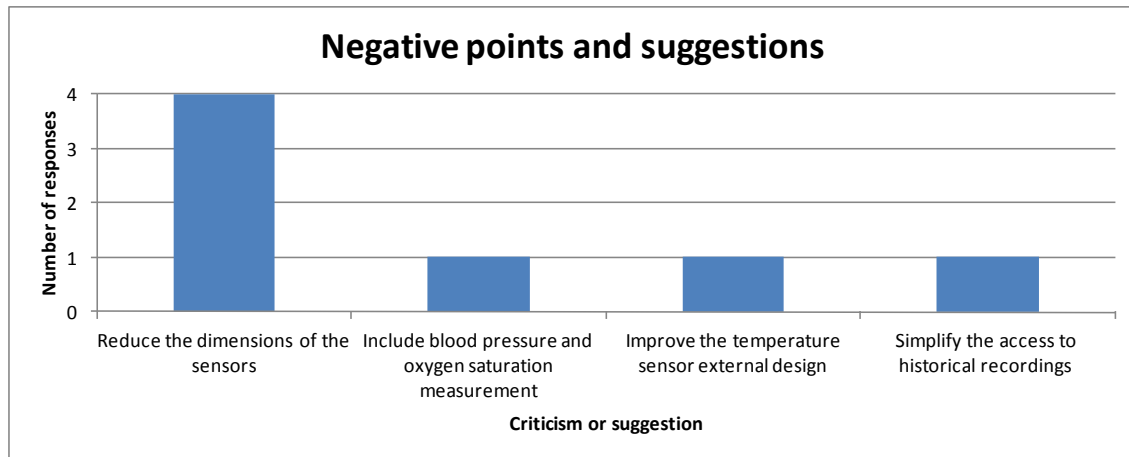


Figure 69 – Negative points and suggestions.

By analyzing the results of the questionnaires and the respondents' verbal opinions, it is possible to conclude that patients accepted the sensors well and that health care providers considered the developed system useful. Also, most health care providers considered that a remote vital signs monitoring can improve the quality of care offered to patients and streamline their work. However, it was clear that these professionals consider essential monitoring the oxygen saturation and the blood pressure.

4.4 Summary

This chapter presents the architecture and the components of HM4All, a remote vital signs monitoring system consisting of wireless sensors based on the ZigBee protocol and Web-based technologies. The system can be used in hospitals; nursing homes; and residential scenarios, provided an Internet connection is available.

Two distinct sensors were developed: an ECG and an axillary temperature sensor. The ECG sensor continuously monitors a modified projection of the bipolar lead I vector and determines the heart rate. The distance between electrodes is 20 mm, resulting in a small form-factor. Moreover, it does not require a chest band and cables between the sensor and the disposable electrodes, which is attractive to hospitals. Based on the heart rate values measured, the system can detect heart rhythm abnormalities,

such as tachycardia, bradycardia, asystole, and background arrhythmia. The temperature sensor is based on a medical thermistor probe and has a precision equal to $\pm 0.2^\circ\text{C}$. It measures the temperature in the axillary area, which is accessible and clinically acceptable.

Networking devices were also developed. These devices are based on high-power modules, which offer an extended range.

The ZigBee-to-IP gateway processes data received from ZigBee coordinators and sends processed data, through a remote connection, to the Application Server software. A serial communications protocol adopted by the coordinator and the ZigBee-to-IP gateway assures that this application can detect serial communication errors and resynchronize. Additionally, it employs a synchronized buffer to temporarily store messages before they are processed. It ensures that the application can continue processing data even in the presence of temporary remote connection problems. Finally, the ZigBee-to-IP gateway has a user interface that presents data and sensor information.

The system acceptance was evaluated using questionnaires answered by patients and health care providers from Hospital Privado de Guimarães. Patients have considered that the developed sensors can improve the care being provided by the hospital. Additionally, most patients considered the sensors lightweight and unobtrusive. All health care providers who responded the questionnaires considered the developed system useful. However, they have suggested including oxygen saturation and blood pressure monitoring. Other suggestions include monitoring the respiration rate, providing patient reports and making it easier to access historical information.

Chapter 5

HM4All evaluation based on simulation and laboratory tests

Although other authors have previously considered the IEEE 802.15.4 and the ZigBee protocols for medical data transport, none have considered a multi-hop topology or the amount of traffic typically generated by a single channel electrocardiogram (ECG) sensor used for monitoring purposes. For instance, Golmie et al., from the National Institute of Standards and Technology (USA), have used simulations to evaluate the suitability of the IEEE 802.15.4 standard to health care monitoring [75]. The most demanding sensor, a multiple-lead ECG, generates 1500 bytes every 250 ms (a much larger amount of data than that generated by the ECG sensor developed in this work). It is shown that using just three such devices results in an overload of the network capacity. In another work, Liang and Balasingham have also considered nonbeacon-enabled star network based on IEEE 802.15.4 protocol [130]. The network consists of ten, fifteen or twenty ECG sensors, each of them generating the double of the traffic generated by the sensors used in [75]. The highest value of delivery ratio (DR) does not reach 90% and is achieved for 10 sensors that transmit full-size packets.

This chapter describes the results of simulations and laboratory tests used to evaluate the performance of the ZigBee network used to transport medical data to the ZigBee-to-IP gateway. The first set of simulations evaluates the performance of the Carrier Sense Multiple Access (CSMA) mechanism used by the Medium Access Control (MAC) layer of the ZigBee protocol when handling the traffic generated by a crescent number of ECG sensors in star and tree topologies. The second simulation aimed on estimating the negative effects of hidden-nodes in the performance of ZigBee networks. Only ECG sensors were considered because they are, by a great deal, the ones that generate the most traffic. OMNet++ [160], a discrete event simulation environment

that can be employed as a network simulation platform, was used to model the IEEE 802.15.4-2003³⁹ MAC layer and evaluate the performance of the protocol applied to the system under development. In addition to simulations, experiments were carried out in the laboratory in order to confirm the results obtained using simulations. By examining packet capture files recorded during the experiments, it was possible to identify the occurrence of contention periods that result from different values of devices' clock drifts and lead to message losses. This issue is described after the presentation of the experimental tests results and is followed by the chapter summary.

5.1 Performance assessment at MAC level

The objectives of these simulations were threefold: a) to estimate the value of the main performance indicators as a function of the number of ECG sensors; b) to assess the impact of adding a new hop to a network; and c) to evaluate how data compression can improve the performance of the network.

5.1.1 Simulation configuration and assumptions

Simulations have considered the star and 2-hop tree topologies. Figure 70 (a) shows a star network consisting of three end devices (EDs), N1 – N3, directly associated to the personal area network (PAN) coordinator C. A 2-hop tree network is shown in Figure 70 (b) in which all EDs are associated to the router R. In this case, a message⁴⁰ generated by an ED cannot be sent directly to the PAN coordinator because an ED can only communicate through its parent. Therefore, every message generated by an ED is sent to the router, and then retransmitted to the PAN coordinator.

³⁹ As the original 2003 version of the IEEE 802.15.4 protocol is the basis of the ZigBee protocol, only this version is considered in this chapter.

⁴⁰ The terms message and packet are frequently used. A message represents information to be sent and a packet is an instance of a message that is transmitted in the channel. Depending on whether collisions occur, a node may need to transmit more than one packet to successfully send one message.



Figure 70 – Star and tree topologies

The EDs consist of ECG sensors which generate 114-byte messages at a fixed period of 250 ms or 500 ms. The first period is used by the developed ECG sensors programmed to transmit raw data, whereas the second period is used by ECG sensors programmed to compress data at a ratio of 2:1. Acknowledgement packets consist of 11 bytes sent without using the CSMA-CA mechanism (described in Section 2.2).

Simulations considered that EDs can either request or not the acknowledgement of transmitted messages. If acknowledgements are required, but the acknowledgement frame is not received, the ED assumes the message was not received by the destination and attempts to retransmit it. On the other hand, if an acknowledgement frame is not expected, the ED assumes the transmitted message was correctly received and does not perform any retry. The IEEE 802.15.4 standard defines that up to three retries can be done before the MAC layer declares that the message transmission has failed [91].

Before transmitting a message, the MAC layer delays for a random number of backoff periods and then requests that the PHY layer performs a CCA. If the channel is found busy, the MAC is required to backoff again before requesting another CCA. Simulations assumed that the CSMA-CA algorithm retries up to four times before declaring a channel access failure. Additionally, the value of the MAC attribute *macMinBE* is assumed to be equal to 3 for all devices (see Table 7, Section 2.2).

The wireless channel was considered free of fading and interferences. Each ED is capable of hearing the transmissions of all others EDs, that is, no hidden-nodes are present. EDs start to generate data at a random instant between the beginning of the simulation and the transmission period (250 ms or 500 ms). Simulation ends when the network coordinator receives 100,000 messages from any ED.

Table 19 summarizes the simulation parameters and their values. Table 20 lists the operation modes simulated, in which columns represent, respectively, the operation mode, the time interval between successive messages generated by EDs, the network topology, and if acknowledgements are required or not.

Table 19 – Parameters common to all simulations.

Parameter	Value
Data packet size (bits)	912
Acknowledgment frame size (bits)	88
The maximum number of retries allowed after a transmission failure (<i>aMaxFrameRetries</i> , [91])	3
The maximum number of backoffs the CSMA-CA algorithm will attempt before declaring a channel access failure (<i>macMaxCSMABackoffs</i> , [91])	4
The minimum value of the backoff exponent in the CSMA-CA algorithm (<i>macMinBE</i> , [91])	3
Messages to receive from any node before simulation stops	100,000

Table 20 – Modes considered in the first set of simulations.

Operation mode	T (ms)	Topology	Ack Tx?
NoACK_star_250ms	250	Star	N
AckTx_star_250ms	250	Star	Y
NoACK_tree_250ms	250	Tree	N
AckTx_tree_250ms	250	Tree	Y
NoACK_star_500ms	500	Star	N
AckTx_star_500ms	500	Star	Y
NoACK_tree_500ms	500	Tree	N
AckTx_tree_500ms	500	Tree	Y

5.1.2 Results

Packet collision and failed transmission attempt ratios

Despite the use of the CSMA-CA mechanism, messages may not be delivered due to collisions or to failure to access a busy channel. As the presence of hidden-nodes is not considered, collisions are only possible if two or more EDs simultaneously sense the channel idle and then transmit their packets. In this case, all packets are lost. In a star network, only EDs contend for the channel. However, in a tree network, both EDs and routers compete and, consequently, for the same number of EDs, more collisions are expected. The same occurs in terms of failed transmission attempts.

Figure 71 presents the packet collision ratio for networks that relay the traffic generated by EDs that transmit (a) raw ECG data and (b) compressed ECG data. In the first case, EDs generate a data packet every 250 ms; whereas, in the second case, the interval between consecutive packets generation is equal to 500 ms. In a star topology, the number of collisions corresponds to the sum of the packets from EDs to the coordinator that collide. For tree topologies, the number of collisions includes a) collisions between packets transmitted by EDs to the router, and b) collisions between a packet transmitted by an ED to the router and a packet transmitted by the router to the

coordinator. Collisions between data packets and acknowledgement frames are not included because these events are very rare, even in dense networks.

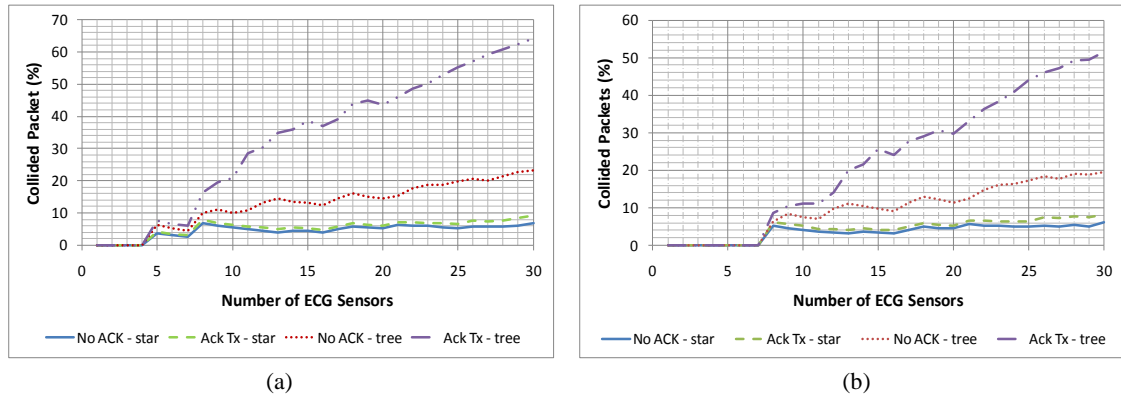


Figure 71 – Collided packet ratio curves for an increasing number of EDs that transmit (a) raw ECG data and (b) compressed ECG data.

In general, when the number of EDs increases, the percentage of packet collisions also increases as a result of the crescent number of messages. Also, in general, for the same number of EDs, more packets collide in case networks allow the use of retransmissions than in the opposite case. This is because the channel gets more congested in the first case than it does in the second case. Due to the limited number of EDs and hops considered, none of the networks collapsed. However, it can occur in case more EDs or hops are added.

Figure 72 presents the failed transmission attempt ratio curves for EDs that transmit (a) raw ECG data and (b) compressed ECG data. Similarly to the packet collision ratio, the failure ratio increases when the number of sensors increases. Additionally, failed transmission attempts are more frequent in deeper networks and in networks that allow the use of retransmissions.




Figure 72 – Failed transmission attempt ratio curves for an increasing number of EDs that transmit (a) raw ECG data and (b) compressed ECG data.

As expected, the curves presented in Figure 71 and Figure 72 show that, for the same topology and transmission mode, both the collision and failure ratios are larger for networks that relay the traffic generated by EDs that transmit raw ECG data than for networks comprised of EDs that transmit compressed ECG data. For instance, for networks that contain up to 7 EDs that transmit compressed ECG data, both ratios are smaller than 1%. In contrast, to achieve a similar result for networks that relay the traffic generated by EDs that transmit raw ECG data, it is necessary to reduce the number of EDs to 4.

Throughput

The normalized throughput represents the ratio between the amount of traffic successfully received by the network coordinator and the network data rate. Figure 73 (a) and Figure 73 (b) present the normalized throughput for star and tree networks that relay the traffic generated by EDs that transmit raw ECG data, whereas Figure 73 (c) and Figure 73 (d) present the normalized throughput for star and tree networks that relay the traffic generated by EDs that transmit compressed ECG data. The blue straight lines in all graphs represent the ideal throughput, which grows linearly as a function of the number of EDs.

For star networks that allow retransmissions, the normalized throughput curves present roughly no deviation from the ideal throughput (in fact, both curves are nearly superimposed). It happens because, in spite of the collided packets and failed transmission attempts, most of the generated messages are transmitted and reach the coordinator. Otherwise, for the same topology, if retransmissions are not allowed, the deviation increases as the number of EDs increases. For instance, for all topologies, the maximum deviation was observed for the maximum number of EDs (30). It is smaller than 0.1% and 0.3% for the `AckTx_star_500ms` and `AckTx_star_250ms` modes, whereas it is equal to 1.53% and 3.44% for the `NoAck_star_500ms` and `NoAck_star_250ms` modes.

Tree networks present a much worse performance than star networks. It happens because messages transmitted by each ED must be retransmitted by the router, resulting in a duplication of the network traffic load relatively to the star network and, consequently, in an increased number of collisions and failed transmission attempts.




Figure 73 – Normalized throughput curves for an increasing number of EDs considering: (a) star topology, raw ECG data; (b) tree topology, raw ECG data; (c) star topology, compressed ECG data; and (d) tree topology, compressed ECG data.

Delivery ratio

The DR represents the percentage of successfully delivered messages to the number of generated messages. The DR curves for networks that transport raw ECG data are shown in Figure 74 (a); whereas the DR curves for networks that transport compressed ECG data are shown in Figure 74 (b). It can be observed that if a small number of EDs are active, all networks present a high reliability, even if acknowledgment frames are not used. However, despite the efforts made by the MAC sublayer to deliver all messages, it is possible to observe a decrease in the DR for both star and tree networks when the number of EDs increases (notice, though, that for star networks that allow retransmissions, the DR does not drop significantly). This decrease reflects the deviations from the ideal throughput shown in the throughput curves, and is ultimately caused by collisions and failed transmission attempts.

Due to the decrease in the traffic load, for the same number of EDs, networks that relay the traffic generated by EDs that compress data achieve a larger DR than networks that contain EDs that transmit raw data.

In general, the use of acknowledgment frames increases the DR. However, when the number of EDs is high, as evidenced by tree networks curves, the traffic load increase caused by collisions (which trigger retransmission attempts) can cause a decrease in the DR relatively to the networks that do not allow retransmissions. It occurs, for instance, for tree networks that contain EDs that transmit raw ECG data when the number of EDs exceeds 13. Eventually, as shown in the DR and throughput graphs, networks can collapse if the number of EDs exceeds a critical value.

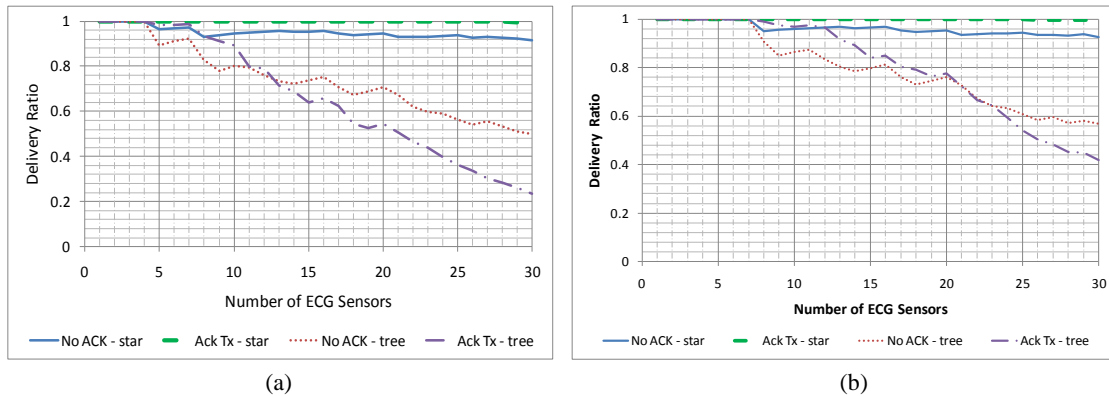


Figure 74 – DR curves for networks that relay the traffic generated by EDs that transmit (a) raw ECG data and (b) compressed ECG data.

Delay

The delay a message experiences comprises the access delay, the turnaround time, the transmission time, the propagation time (negligible) and the time spent during retries. Table 21 shows how to determine the maximum access delay a message can experience considering that a variable number of attempts is done until the channel is found idle. The MAC attributes *macMaxCSMAbackoffs* and *macMinBE* are set to their default values, 4 and 3, respectively [91]. For example, if the device must assess the channel twice to find it idle and, on both attempts, the MAC layer selects the largest possible number of backoff periods, the access delay the message experiences is equal to 7.296 ms. In the limit, a message can experience an access delay of 37.44 ms.

For the star network topology, the minimum time required to transmit an ECG data message is equal to 3.968 ms if acknowledgment frames are not required, or 4.512 ms otherwise. The required events, since the message is generated, and their minimum durations are listed in Table 22. For the tree topology, the minimum time necessary, considering the ED is associated to a router is, approximately, twice as long, as shown

in Table 23. Notice that, on all cases, the backoff periods selected by the devices are equal to zero.

Table 21 – Maximum access delay as a function of the number of transmission attempts done.

Number of channel access attempts done to find the channel idle	Maximum access delay
1 attempt:	2.368 ms
1 st try = 2.368 ms	
Max. random delay (BE = 3, NB = 0) = $7 * 0.32 = 2.24$ ms	
CCA = 0.128 ms	
2 attempts:	7.296 ms
1 st try = 2.368 ms (channel busy: backoff)	
2 nd try (1 st retry) = 4.928 ms	
Max. random delay (BE = 4, NB = 1) = $15 * 0.32 = 4.80$ ms	
CCA = 0.128 ms	
3 attempts:	17.344 ms
1 st try = 2.368 ms (channel busy: backoff)	
2 nd try (1 st retry) = 4.928 ms (channel busy: backoff)	
3 rd try (2 nd retry) = 10.048 ms	
Max. random delay (BE = 5, NB = 2) = $31 * 0.32 = 9.92$ ms	
CCA = 0.128 ms	
4 attempts:	27.392 ms
1 st try = 2.368 ms (channel busy: backoff)	
2 nd try (1 st retry) = 4.928 ms (channel busy: backoff)	
3 rd try (2 nd retry) = 10.048 ms (channel busy: backoff)	
4 th try (3 rd retry) = 10.048 ms	
Max. random delay (BE = 5, NB = 3) = $31 * 0.32 = 9.92$ ms	
CCA = 0.128 ms	
5 attempts:	37.440 ms
1 st try = 2.368 ms (channel busy: backoff)	
2 nd try (1 st retry) = 4.928 ms (channel busy: backoff)	
3 rd try (2 nd retry) = 10.048 ms (channel busy: backoff)	
4 th try (3 rd retry) = 10.048 ms (channel busy: backoff)	
5 th try (4 th retry) = 10.048 ms	
Max. random delay (BE = 5, NB = 4) = $31 * 0.32 = 9.92$ ms	
CCA = 0.128 ms	

Table 22 – Events and minimum periods of time involved in the transmission of an ECG data message for the star topology.

Event description	Action	Time required (ms)
ED selects a backoff period equal to 0	Backoff	0 (minimum)
ED senses the channel clear	CCA	0.128
ED's transceiver changes to TX mode	Turnaround RX-to-TX	0.192
ED transmits a data packet to the coordinator	TX	3.648 (114 bytes)
ED's transceiver changes to RX mode	Turnaround TX-to-RX	0.192
ED receives the acknowledgment	RX	0.352 (11 bytes)

Table 23 – Events and minimum periods of time involved in the transmission of an ECG data message for the 2-hop tree topology.

Event description	Action	Time required (ms)
ED selects a backoff period equal to 0	Backoff	0 (minimum)
ED senses the channel clear	CCA	0.128
ED's transceiver changes to TX mode	Turnaround RX-to-TX	0.192
ED transmits a data packet to the router	TX	3.648 (114 bytes)
ED's transceiver changes to RX mode	Turnaround TX-to-RX	0.192
ED receives the acknowledgment frame	RX	0.352 (11 bytes)
Router's transceiver changes to RX mode	Turnaround TX-to-RX	0.192
Router selects a backoff period equal to 0	Backoff	0 (minimum)
Router senses the channel clear	CCA	0.128
Router's transceiver changes to TX mode	Turnaround RX-to-TX	0.192
Router retransmits the data packet to the coordinator	TX	3.648 (114 bytes)
Router's transceiver changes to RX mode	Turnaround TX-to-RX	0.192
Router receives the acknowledgment frame	RX	0.352 (11 bytes)

Table 24 shows how to find the maximum delay an ECG data message can experience in a star network considering that the MAC layer fails three times to receive the acknowledgment frame and that, for each retry, the CSMA-CA mechanism makes up to five attempts to find the channel idle. As shown, in the limit, a message can experience a relatively large delay. For a 2-hop tree topology, the maximum delay is, approximately, twice the maximum delay experienced in a star network.

Table 24 – Maximum delay an ECG data message can experience in a star network.

Maximum delay (worst case) – single hop (aMaxFramesRetries = 3)	
1 st attempt, 1 st and 2 nd retries = $[3 * (37.440 + 0.192 + 3.648 + 0.864)] = 126.432$ ms	
5 attempts to access the channel (4 failed attempts have been done before the channel is found idle. See Table 21)	37.440 ms
Turnaround time (TX-to-RX)	0.192 ms
Time on air (114 bytes)	3.648 ms
Wait for the acknowledgment frame (ACK failure reported)	0.864 ms
3 rd retry = $37.44 + 0.192 + 3.456 = 41.088$ ms	
5 attempts to access the channel	37.440 ms
Turnaround time (TX-to-RX)	0.192 ms
Time on air (114 bytes)	3.648 ms
Maximum delay = 126.432 ms + 41.280 ms = 167.712 ms	

The delay also includes the queuing delay, that is, the time a message waits in the device's buffer. As the IEEE 802.15.4 specification leaves the buffer size definition up to the implementation, the transmit buffer size at the MAC layer differs depending on the manufacturer and the specific device. For instance, Jennic wireless microcontrollers have a transmit buffer at the MAC layer that can hold up to four packets, whereas ZigBee routers can queue up to eight packets [112]. In our specific case, as messages

are generated at intervals of 250 ms or 500 ms and the maximum delay per hop a message can experience is 167.712 ms, the buffer is never used by devices.

Figure 75 and Figure 76 show the mean and maximum delay curves for networks that relay the traffic generated by EDs that transmit raw and compressed ECG data, respectively. As shown, when the number of EDs is small, the delay messages experience for all modes of operation is approximately constant. However, when the number of EDs exceeds a certain limit, the delay increases with the number of nodes. As observed, in denser networks, the delay values are high for networks that allow retransmissions. As EDs get the opportunity to retransmit messages that collide, the channel gets busier. Consequently, EDs backoff more often and the number of collisions increases. Although it is possible that a message experience a maximum delay of 167.712 ms per hop, this is highly unlikely. The maximum delay obtained for a star network is 90 ms; whereas, for tree networks, this value is smaller than 190 ms.




Figure 75 – Mean and maximum delay for networks that only contain EDs that transmit raw ECG data.




Figure 76 – Mean and maximum delay for networks that only contain EDs that transmit compressed ECG data.

Energy consumption

The average value of the energy spent by one ED to deliver a message to its parent can be obtained as the ratio of the total energy spent by the ED to deliver all messages, including packets retransmissions where applicable, to the number of messages sent by the ED. Voltage and current consumption values specified for JN5139 ZigBee modules [110] were used since the developed sensors are based on these modules. A voltage supply of 3 V was considered. When the module's microcontroller is switched on, the module consumes 9.21 mA, as shown in Figure 77. During sampling, the ADC is switched on, and the current consumption increases to 9.79 mA. When the module switches the transceiver on during the backoff process, the current increases to 32 mA and, when the module starts transmitting, the current increases to 37 mA.




Figure 77 – JN5139 module current consumption per relevant activity.

Figure 78 presents the curves of the average energy consumed by an ECG sensor to transmit a message, considering only the transceiver consumption. For instance, for a 2-hop tree network that includes six EDs that transmits compressed ECG data and employs acknowledged data messages, the average energy consumed to transmit one message is 0.44 mJ. Since during any one second interval the ED transmits two messages, the consumption relative to the transceiver is 0.88 mJ. Additionally, the ADC is used to make 200 measurements, with each measurement taking 1 ms, which results in a consumption of 0.348 mJ. The consumption of the microcontroller is 27.63 mJ, so the total energy consumption is equal to 28.858 mJ. If the module is powered by a +3V, 1200 mAh battery (12960 J), a lifetime of 124 hours of continuous operation can be achieved. The calculations are detailed in Table 25. As observed, most of the energy is consumed by the microcontroller, which is kept on to process the samples of the ECG waveform.




Figure 78 – Energy consumption per message for networks that contain EDs that transmit (a) raw ECG data and (b) compressed ECG data.

Table 25 – Energy consumption calculation for an ED that transmits compressed ECG data.

Event description	Period of time	Energy consumption (mJ)
Transceiver consumption per message (determined by simulation)	2 messages in 1 s	$E = 2 * 0.44 \text{ mJ} = 0.88 \text{ mJ}$
ADC sampling (200 Hz)	$200 * 1 \text{ ms}$	$E = 3 \text{ V} * (9.79 - 9.21) \text{ mA} * (200 * 1 * 0.001) = 0.348 \text{ mJ}$
Microcontroller on	1 s	$E = 3 \text{ V} * 9.21 \text{ mA} * 1 \text{ s} = 27.63 \text{ mJ}$
Total energy consumption = 0.88 mJ + 0.348 mJ + 27.63 mJ = 28.858 mJ		

5.1.3 Simulation results outline

Table 26 summarizes the simulation results presented. Two important QoS markers are considered: a DR greater than 99.9% and a maximum delay smaller than 250 ms. The maximum number of ECG sensors each network can have and still comply with the established requirements is shown. For instance, if up to 23 ECG sensors that transmit compressed data are active, a star network can achieve a DR better than 99.9%. On the other hand, to achieve the same DR for a tree topology, it is necessary to reduce the number of ECG sensors to 7. As opposed to the DR, which is severely impacted the number of EDs, transmission mode and topology, the delay limit established is not reached on any mode.

Table 26 – Maximum number of ECG sensors to achieve established QoS parameters.

Operation mode	DR > 99.9%	Max. delay < 250 ms
NoACK_star_250ms	4	> 30
AckTx_star_250ms	20	> 30
NoACK_tree_250ms	4	> 30
AckTx_tree_250ms	4	> 30
NoACK_star_500ms	7	> 30
AckTx_star_500ms	23	> 30
NoACK_tree_500ms	7	> 30
AckTx_tree_500ms	7	> 30

It is possible to observe that the use of acknowledged transmissions contributes favorably to the performance of the networks, though an increase in the maximum delay is experienced and a small decrease in the lifetime of each ED is also expected.

5.2 Performance assessment with hidden-nodes

The performance of CSMA-based WSNs can be seriously degraded by collisions caused by hidden-nodes. In a CSMA-based network, a node can only transmit if it senses the channel idle. As explained on Section 2.2.3, the hidden-node problem occurs when the carrier sensing fails and a node starts transmitting when other node has already occupied the channel. If both transmissions are within the reach of a receiver, a collision occurs.

No specific mechanism to avoid the hidden-node problem is provided by the IEEE 802.15.4 protocol, which motivated some authors to consider specific scenarios and propose strategies to mitigate it. Three of the most prominent ones involve grouping nodes that have bidirectional connectivity between each other [86, 126, 179]. However, these strategies require the modification of the original protocol and consider beacon-enabled networks consisting of static nodes, which is not the scenario considered in this work.

This analysis aims at estimating the degradation of the network performance, namely the delivery ratio, in the high load scenario imposed by the transmission of ECG data, considering different percentages of hidden-nodes. The terms node and sensor, used throughout this section, are used interchangeably.

5.2.1 Simulation configuration and assumptions

The simulations involving hidden-nodes considered an increasing number of EDs which send data directly to the coordinator, in a star topology. Five scenarios were simulated, where each scenario considered a different percentage of hidden-nodes in the network, varying from no hidden-nodes up to 20% of hidden-nodes. The situation depicted in Figure 20 (b) was simulated since in the star topology all EDs transmit to the coordinator. In that case, if one ED starts to transmit when a node that is hidden from it is already transmitting, both packets are lost and no acknowledgements are sent back by the network coordinator.

The parameters used by this simulation are the same ones used to evaluate the performance of star and tree networks, and are presented on Table 19. As before, EDs start to generate data at a random instant between the beginning of the simulation and the transmission period (250 ms or 500 ms). The wireless channel was considered free of fading and interferences.

5.2.2 Results

Packet collision and failed transmission attempt ratios

Collisions can involve data packets and acknowledgment frames. A collision between a data packet and an acknowledgment frame occurs if an ED fails to detect a packet transmission and starts transmitting just after the packet is received by the network coordinator. These collisions are rare when compared to the collisions between packets because acknowledgment frames are relatively small packets that quickly follow a received packet.

Figure 79 shows the collided packet ratio curves for networks that contain EDs that transmit (a) raw ECG data and (b) compressed ECG data. When the number of EDs is small, the percentage of collided packets is small. However, when the number of EDs increases, the collided packet ratio increases for both modes. Also, the collided packet ratio increases as the percentage of hidden-nodes increases.




Figure 79 – Collided packet ratio curves for networks that contain different percentages of hidden nodes and a crescent number of EDs that transmit (a) raw ECG data and (b) compressed ECG data.

Figure 80 presents the failed transmission attempt ratio curves for networks that contain EDs that transmit (a) raw ECG data and (b) compressed ECG data. When the number of nodes increases, more collisions and, consequently, more retransmissions occur. As a result, the channel gets busy, and the failure ratio increases. Similarly, since

collisions are frequent for networks with high percentages of hidden-nodes, the failure ratio increases as the percentage of hidden-nodes increases.

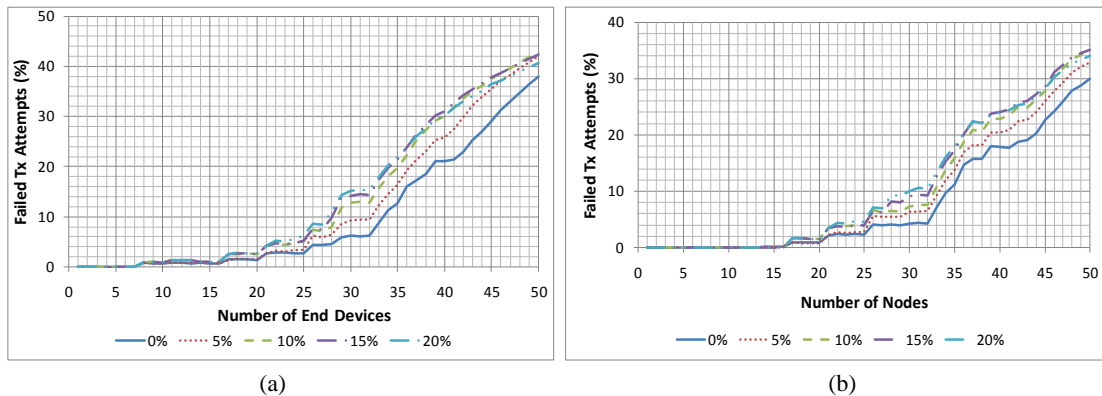


Figure 80 – Failed transmission attempt ratio curves for networks that contain different percentages of hidden nodes and a crescent number of EDs that transmit (a) raw ECG data and (b) compressed ECG data.

Throughput

Figure 81 shows the normalized throughput curves as a function of the number of EDs, considering a percentage of hidden-nodes ranging from no hidden-nodes up to 20% of hidden-nodes, for EDs that (a) transmit ECG raw and (b) compressed ECG data. The continuous straight lines shown in both graphs represent the ideal throughput. On all cases, it is possible to observe a deviation from the ideal throughput when the number of EDs increases. Additionally, it is observed that as the percentage of hidden-nodes increases, the deviation between the corresponding throughput curve and the ideal throughput curve increases. For instance, if raw data are transmitted, a deviation smaller than 1% from the ideal throughput can be achieved if the number of EDs does not exceed 34 and no hidden-nodes exist. However, for a percentage of hidden-nodes of 5%, the number of nodes must be reduced to 25 not to exceed this deviation. In this case, when the number of EDs reaches 40, the network starts to collapse due to excessive traffic. In the limit, if this percentage is increased to 20%, to achieve a deviation smaller than 1% of the ideal throughput, the number of nodes must not exceed 13.




Figure 81 – Throughput curves for networks that contain different percentages of hidden nodes and a crescent number of EDs that transmit (a) raw ECG data and (b) compressed ECG data.

Delivery ratio

The DR variation as a function of the number of nodes and the percentage of hidden-nodes is shown in Figure 82. Due to the random selection of hidden-node combinations, the curves exhibit some fluctuations. Nevertheless, they correctly evidence the highly negative effect of hidden-nodes in the network performance. In fact, if no hidden-nodes are present, high DR values are achieved with a relatively high number of nodes. However, even a small percentage of hidden-nodes can cause a considerable increase in the number of collisions and, consequently, seriously affect the performance of the network. For instance, consider the transmission of raw ECG data. If no hidden-nodes are present, more than 99.9% of the messages generated by up to 20 EDs are correctly delivered. However, if the percentage of hidden-nodes is increased to 5% or 10%, the number of EDs must be reduced to 10 or 4, respectively, to achieve the same DR.




Figure 82 – DR curves for networks that contain different percentages of hidden nodes and a crescent number of EDs that transmit (a) raw ECG data and (b) compressed ECG data.

5.2.3 Simulation results outline

Table 27 summarizes the simulation results presented. The first column refers to the percentage of hidden-nodes present in the network. The second and third columns present the maximum number of ECG sensors a network can contain and still present a DR greater than 99.9% for both operation modes. As observed, the hidden-node problem causes an increasing number of collisions which seriously affect the network throughput and, consequently, the DR. For instance, if no hidden-nodes are present, the network can correctly relay the traffic generated by a large number of ECG sensors. However, in the presence of hidden-nodes the network performance degrades considerably and the number of ECG sensors must be considerably reduced to maintain the DR.

Table 27 – Hidden-node analysis results summary for nonbeacon-enabled ZigBee/IEEE 802.15.4-based star networks.

Percentage of hidden-nodes (%)	DR > 99.9%	
	Raw data	Compressed data
0	20	23
5	10	13
10	4	11
15	4	11
20	4	11

5.3 Laboratory experiments

Experimental tests were performed in the laboratory environment to assess the performance of star and tree nonbeacon-enabled ZigBee-based WSNs consisting of EDs that transmit either raw or compressed ECG data. Additionally, capture files are analyzed to understand the causes of message failures and to find possible ways to improve the performance.

5.3.1 Experimental tests configuration

Experiments were accomplished using the test boards included in the JN5139-EK010 evaluation kit [109] and developed test boards, all based on the wireless module JN5139 [106], from Jennic. All modules run the ZigBee version 1.0 (2004). Figure 83 shows a developed test board and Figure 84 presents its schematics diagram. This design is based on the sensor boards supplied with the evaluation kit and include circuits

to program the microcontroller (J2 connector, PROG and RESET buttons), to monitor the state of the microcontroller's digital input output line DIO16 (LED D1), to sample an analog signal using the microcontroller's analog to digital controller input ADC1 (connector J3, pin 3), to power on or off the board (switch SW3), and to monitor its power state (connector J3, pins 1 - 2 and LED D2). The boards are powered by two AAA batteries (BATT1 and BATT2).




Figure 83 – Test boards: (a) top side (b) bottom side




Figure 84 – Schematic diagram of the developed test boards.

The topologies used to evaluate the performance of the networks are shown in Figure 85. The coordinator relied on the sequence number attached to the payload of each message transmitted by all EDs to determine the average DR. The average delay was estimated by determining the time interval between the instant a message is generated by a specific ED and the moment it is received by the network coordinator. The specific ED that generated messages used to measure the delay was designated Delay Test ED (DTED). One of the DIO pins of the DTED was wire connected to one of the DIO pins of the coordinator. This connection is labeled TRIGGER in Figure 85.

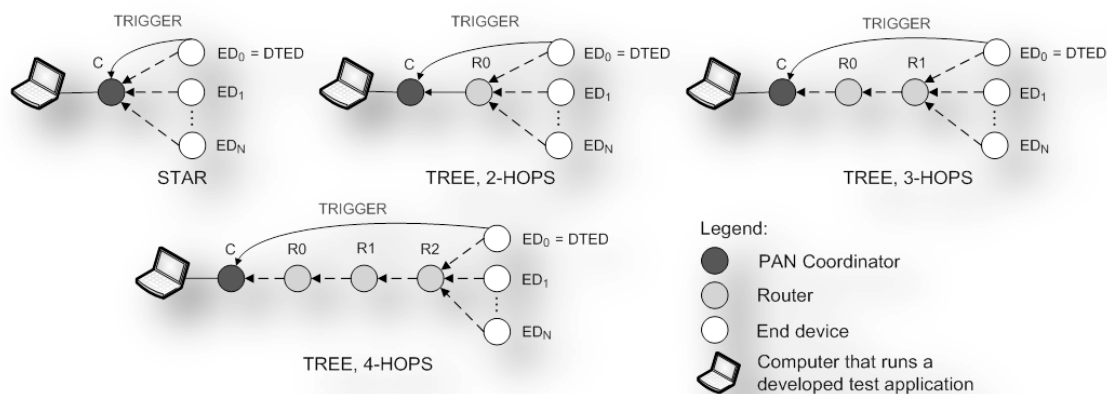


Figure 85 – Star and tree topologies used to measure the DR and delay.

The procedure used to measure the average delay is described as follows. Just before the DTED transmits a message, it informs the coordinator by changing the state of the DIO pin connected to the coordinator. The application running in the coordinator is interrupted and then reads and stores the counting value supplied by an internal 1 ms resolution timer. When the message is received, the coordinator reads the counting value again and computes the interval between the two events, which results in a good approximation of the end-to-end delay experienced by the message transmitted by the DTED.

Packets sent by EDs were addressed to the network coordinator and included a payload of 79 bytes, which consisted of 75 data bytes, 2-byte identification number and 2-byte sequence number added by the Application (APP) layer. All messages used hop-by-hop acknowledgements. Devices could make up to five attempts to access the channel and up to three retries were allowed. All experimental tests were carried out using channel 26, to reduce interference from WLANs. Test boards were positioned relatively close to each other to avoid the presence of hidden-nodes. It was necessary to

dynamically prevent the network coordinator and some routers (R0 in the three-hop topology and R0 and R1 in the four-hop topology) from accepting more than one child during network formation to create the desired tree network from devices that were at the radio range of each other. Additionally, to preserve the topology, it was necessary to prevent routers from establishing direct routes to the network coordinator.

Up to twelve EDs were used (including the DTED). The DR and delay were measured for each configuration (for instance, a 2-hop tree network, 6 EDs transmitting raw data) at least three times. For each run, at least 1,000 messages were sent by each ED to the network coordinator.

5.3.2 Results

Delivery ratio

Figure 86 shows the lowest DR values measured for all runs and considering different topologies and a crescent number of EDs that transmitted (a) raw ECG data or (b) compressed ECG data (for each topology and for a specific number of EDS, at least three independent test runs were performed). As shown, both star networks were able to successfully relay all traffic generated by EDs. However, the traffic increase in tree networks impacted their performance. The 4-hop tree networks presented the worst performance, followed by the 3-hop tree networks and so on. Additionally, due to the lower traffic load, an n-hop network that relayed compressed traffic behaved better than an analogous n-hop networks that relayed raw data.




Figure 86 – Lowest DR values measured for (a) raw ECG data transmission and (b) compressed ECG data transmission.

Delay

Figure 87 present the mean end-to-end delay curves for a variable number of end devices that transmit raw and compressed ECG data, respectively. Again, each curve represents a different topology. As shown, if the network traffic is not intense, the mean delay experienced by packets is almost constant despite the number of EDs. However, when the traffic increases, EDs are forced to backup more often and the delay increases. This situation was observed for 3-hop and 4-hop tree networks that relayed raw ECG traffic and for 4-hop tree networks that relayed compressed ECG data.




Figure 87 – Mean delay measured for (a) raw ECG data transmission and (b) ECG compressed ECG data transmission.

Figure 88 show the highest delay values observed for raw and compressed ECG data. In the first case, the highest delay observed for a 4-hop tree network was 360 ms. If data are transmitted at 250 ms intervals, the total delay experienced by the first data sample collected is equal to 610 ms (360 ms + 250 ms). In case of EDs that transmit compressed data, the highest delay observed was 181 ms and, as a result, the highest delay experienced by a sample is equal to 681 ms (181 ms + 500 ms). The IEEE Std. 11073-00101-2008 [2] specifies that for real-time waveform transmission, the communication delay shall not exceed 3 s (see Table 22). On both cases the delay values measured are under acceptable limits, however, it is necessary to add these values to the delay messages experience on next stages.




Figure 88 – Highest value of the delay measured for (a) raw ECG data transmission and (b) compressed ECG data transmission.

5.3.3 Laboratory tests outline and discussion

Some of the DR and delay curves presented are not monotonic, which means that DR and delay values do not strictly decrease or increase as the number of EDs increases. Additionally, for the same number of EDs, some measured values are lower for an n -hop topology than for an $(n-1)$ -hop topology. For example, for networks that relay raw data and consists of 6, 7 or 8 EDs, the lower DR measured for a 3-hop tree network was higher than the lower DR measured for 2-hop tree networks (see Figure 86). These unexpected results happened because in nonbeacon-enabled networks transmissions are not synchronized and, consequently, the test outcomes obtained are highly dependent on the initial conditions, specifically on the time intervals between successive transmission attempts made by EDs. Consequently, in the example just mentioned, it is possible and very likely that if additional tests have been done, lower DR values would have been measured for 3-hop tree networks.

The mentioned time intervals are initially established when EDs join the network and are, with small variations, maintained during a certain amount of time because all EDs generate messages at approximately regular time intervals (every 250 ms or 500 ms). During network operation, these intervals change due to clock drifts, which can aggravate or improve the overall network performance. It is necessary to significantly increase the duration of each experiment to better describe the performance of each network. Nevertheless, results presented are important references when considering the design of health monitoring systems based on the ZigBee protocol. For instance, it is not possible to guarantee an average DR greater than 99% when using a 2-hop tree network

consisting of 6 EDs that transmit raw ECG data. However, it is possible to use this topology if ECG sensors transmit compressed data.

Table 28 summarizes the experimental results. Star networks operating in the unslotted CSMA-CA mode can relay at least 99.9% of all the traffic generated by twelve ECG sensors on both modes. The increase of traffic load in multi-hop networks imposes a reduction in the number of ECG sensors, but these networks can still reliably relay the traffic of a significant number of sensors. Additionally, it was observed that messages do not experience large delays even in networks with several hops.

Table 28 – Experimental results summary.

Network topology	Max. number of EDs for DR \geq 99.9%		Max. end-to-end delay (ms)	
	Raw data	Compressed data	Raw data	Compressed data
Star	12	12	12	10
Tree, 2 hops	5	6	96	17
Tree, 3 hops	5*	5	151	51
Tree, 4 hops	3	5	360	181

(*) The result for the 3-hop tree topology exceeds the result for the 2-hop tree topology. As the first topology cannot outperform the second, the smaller number of EDs was considered for both topologies.

Table 29 compare the experimental and simulation DR results obtained for star and 2-hop tree networks that contain ECG sensors that transmit raw ECG data. The same analysis applies to networks that contain sensors that transmit compressed data. Delay results are not compared because both simulation and measurement results indicate they are acceptable.

Table 29 – Experimental and simulation results comparison.

Number of EDs	DR, star (%)		DR, 2-hop tree (%)	
	Experimental (Min. DR)	Simulation (ACK req.)	Experimental (Min. DR)	Simulation (ACK req.)
1	100	100	100	100
2	100	100	100	100
3	100	100	100	100
4	100	100	100	100
5	100	100	100	98.3
6	99.9	100	97.9	98.5
7	99.9	100	96.1	98.7
8	100	100	95.5	93.5
9	99.9	100	97.2	91.1
10	100	100	95.0	89.4
11	99.9	100	95.8	80.0
12	99.9	100	85.4	79.0

The DR differences between experimental and simulation results for star networks are small. However, for 2-hop tree networks some results disagree. It mainly occurs because simulations have considered that the time relationships between messages generated by EDs are kept constant. On the other hand, as experiments were repeated several times (three times minimum), for each run, devices assumed different time relationships. So, while simulations have not considered the occurrence of clock drifts, experiments were too brief to capture their effects. Future simulations applied to nonbeacon-enabled networks will consider clock drifts in order to account for message losses that may occur during contention periods. Additionally, the duration of future tests will be significantly extended to capture the effects of clock drifts.

5.3.4 Clock drift measurement

The clock drift is often expressed in parts per million (ppm) and gives the number of additional or missing oscillations a clock makes in the amount of time needed for one million oscillations at the nominal rate. Apart from manufacturing inaccuracies, the frequency of oscillators may derive because of aging effects and temperature variations, among other causes [118]. Since sensors run on relatively cheap oscillators, large clock drifts are expected. If no synchronization method is applied to a WSN, nodes keep no time relationship between their clocks.

In a nonbeacon-enabled WSN, the time differences between transmissions from different sensors can vary considerably in a relatively small amount of time due to clock drift. The following experience was run to observe the evaluation of the relative time differences between transmissions done by sensors. Four test boards based on JN5139 modules were programmed to generate one message at 500 ms intervals. These modules have a maximum clock drift of 40 ppm (that is, their clocks can deviate up to 40 μ s in 1 s) [106]. A nonbeacon-enabled star network was used and packets were captured using the SNA network analyzer [40]. Sensors were activated and, periodically, the transmission times of successive messages from all sensors were recorded. Collision avoidance was disabled during the first iteration of the CSMA-CA algorithm⁴¹ [91].

⁴¹ That is, the MAC layer *macMinBE* parameter was set to 0.

Representative results are presented in Table 30 where each line contains the transmission time of four consecutive messages transmitted by ED1 – ED4 and the transmission sequence. The transmission time is partly shown in the first column (hour, minutes and seconds, in hh:mm:ss format) and partly in the columns ms/EDX (milliseconds). For instance, in the first line, it is shown that ED1 transmitted a message at 09:00:00.594. It was followed by a message transmitted by ED2 at 09:00:00.646. The sequence of transmissions is shown in the last column. For instance, during the time recording performed at 9 am, the message transmitted by ED1 is followed by messages transmitted by ED2, ED3 and ED4, respectively. The last line of Table 30 shows the differences, in milliseconds, between the first and the last transmission recorded for each ED during the test. The recordings were done using the SNA network analyzer, from Daintree [40]. Consequently, all times recorded are relative to the portable computer running the SNA.

Table 30 – Transmission time measurements.

PC time (hh:mm:ss)	ms				Tx sequence
	ED1	ED2	ED3	ED4	
09:00:00	594	646	723	744	ED1 – ED2 – ED3 – ED4
09:30:00	581	628	706	716	ED1 – ED2 – ED3 – ED4
10:00:00	527	572	647	654	ED1 – ED2 – ED3 – ED4
10:30:00	492	534	609	598	ED1 – ED2 – ED4 – ED3
11:00:00	480	518	592	567	ED1 – ED2 – ED4 – ED3
11:30:00	464	501	572	536	ED1 – ED2 – ED4 – ED3
12:00:00	450	484	556	509	ED1 – ED2 – ED4 – ED3
12:30:00	436	467	536	479	ED1 – ED2 – ED4 – ED3
13:00:00	390	424	487	417	ED1 – ED4 – ED2 – ED3
13:30:00	353	381	447	368	ED1 – ED4 – ED2 – ED3
14:00:00	308	328	395	302	ED4 – ED1 – ED2 – ED3
Time diff. (ms)	286	318	328	442	

The test results presented in Table 30 show that, over time, all sensors slightly delay their transmissions due to clock drifts. Additionally, it can be observed that ED4 is the one with the greatest drift. As shown in Table 30, since the beginning of the test, the time periods between successive transmissions (and, consequently, between messages generation) made by ED3 and ED4 are becoming shorter. Figure 89 illustrates what happened between 10 am and 10:30 am, where each box represents a message transmitted by a sensor. At 10 am, ED3 transmitted a message 7 ms before ED4 made its transmission. Between 10 am and 10:30 am, the order of messages changed, that is, at 10:30 am, the transmission made by ED4 preceded the transmission made by ED3 by

11 ms. During this specific interval, the probability of collisions increases because both devices can simultaneously sense the channel idle and transmit packets that are lost at the receptor. As shown in Table 30, similar situations occur between 12:30 pm and 1 pm and between 1:30 pm and 2 pm.



Figure 89 – Relative transmission sequence change due to clock drift.

A rough estimation of the clock drift of each sensor can be done by considering the time differences registered during the test⁴². For instance, the clock of ED1 has an approximate drift of 15.9 ppm ($286 \text{ ms} / (5 \text{ h} * 60 \text{ m} * 60 \text{ s} * 1000 \text{ ms})$) relative to the clock of the personal computer used during the test.

5.4 Summary

This chapter presents the results of simulations and laboratory tests used to assess the performance of the ZigBee network that will transport data generated by medical sensors. Additionally, it addresses contention issues arising from different sensors' clock drifts. In view of the results obtained, the following conclusions apply:

- A nonbeacon-enabled star network can successfully relay the traffic generated by up to 12 ECG sensors on both modes with a $DR \geq 99.9\%$.
- A 2-hop tree network consisting of up to 6 ECG sensors that transmit compressed data can achieve an average $DR \geq 99.9\%$. For ECG sensors that transmit raw ECG data, the number of sensors should be reduced to 5 to achieve a DR equal to 99.9%.

⁴² An exact estimation of the clock drift should use a set up that eliminates uncertainties. This test did not aim at measuring the clock drift, but to observe its effect.

- The delay experienced by ECG data messages is within the acceptable limits for real-time waveform transmission according to the IEEE 11073-00101-2008 standard.
- As shown for star networks, the presence of hidden-nodes can severely affect the performance of IEEE 802.15.4 and ZigBee-based networks. For instance, if no hidden-nodes are present, a $DR \geq 99.9\%$ can be achieved using up to 23 ECG sensors that transmit compressed data. However, in the presence of 5% of hidden-nodes, to maintain the DR, it is necessary to reduce the number of ECG sensors to 13.
- Future simulations should be improved by considering the clock drift values typical of low cost devices used to develop sensors based on the IEEE 802.15.4 and the ZigBee protocols.

Chapter 6

HM4All evaluation based on field tests

This chapter presents the results of HM4All evaluation based on field tests executed in a non-acute in-patient area of *Hospital Privado de Guimarães*, a private Portuguese hospital. The floor where the tests were executed has three private rooms and nine semiprivate rooms (rooms that can be shared by two patients) in which low acuity patients and parturients stay for a period of time that, normally, does not exceed five days. The system was specified to continuously monitor up to six patients using electrocardiogram (ECG) and temperature sensors.

The following tasks were accomplished:

- A radio survey to determine the quietest available channels;
- A physical inspection to determine where fixed devices could be installed;
- Connectivity tests to define the radio coverage of each network or router;
- Short range and link quality tests to verify the expected range of devices and the quality of the wireless links; and
- Field tests to evaluate the system performance under different conditions.

The research team had no control on the environmental conditions during field tests, including the use of wireless devices. All tests were conducted under close supervision of the hospital staff.

6.1 Radio survey

All hospital areas have Wi-Fi connectivity based on the IEEE 802.11g protocol and there are no restrictions to the use of wireless devices. Bulky data files from picture archive and communication systems (PACS) and additional traffic related to the enterprise resource planning (ERP) system are transferred over the local area network

(LAN), whereas the wireless local area network (WLAN) is mainly used by health care providers, patients and visitors to access the Internet and legacy applications.

The WLAN operates on channels 1, 6 and 11, as observed in the screenshot shown in Figure 90. This radio activity recording was acquired using the Wi-Spy 2.4x [147], a portable spectrum analyzer, and the software Chanalyzer 3.3 for the Wi-Spy 2.4x [146]. The software presents three views, namely, the spectral, the topographic and the planar views.

The spectral view registers the activity on the 2.4 GHz spectrum over a time window where blue represents the minimum level of activity detected (-110.0 dBm) and red the maximum level (-38.0 dBm). As shown, at the time the screenshot was acquired, most of the activity was detected on channels 1 and 6. The topographic view is shown just below the spectral view. It presents channel signatures, which are drawn for each active channel based on the amplitudes of the signals detected during the survey period. The last one, the planar view, shows the maximum, the average and current amplitude of the signal detected as a function of the frequency. At the time the screenshot was acquired, channel 1 was the one with most traffic. The average activity observed, which is represented by small green spikes, was low. No other important source of interference was detected during the survey though some interference from Bluetooth devices is expected.




Figure 90 – 2.4 GHz spectrum activity

Table 31 shows the Wi-Fi activity reported by the spectrum analyzer during the survey. Three columns are shown: the Wi-Fi channel number, the channel grade and the average peak power detected. The channel grade is a measurement of how quiet the channel is. Numbers and letters are used to classify a channel. The highest values correspond to the quieter channels. An ‘A’ grade channel is quieter than an ‘A-’ grade channel, which, in turn, is quieter than a ‘B+’ grade channel and so on. As observed, channels 1 and 2 were the busiest, whereas channels 6 and 11 were relatively quiet.

Table 31 – Wi-Fi report

Channel	Grade	Avg. Peak (dBm)
1	89 (B+)	-40.37
2	90 (B+)	-36.46
3	93 (A-)	-35.2
4	95 (A)	-35.57
5	96 (A)	-35.14
6	95 (A)	-35.53
7	96 (A)	-35.39
8	96 (A)	-34.61
9	97 (A)	-34.71
10	97 (A)	-35.85
11	97 (A)	-38.99
12	97 (A)	-44.14
13	98 (A)	-50.42
14	99 (A)	-61.53

Table 32 shows the survey results for the ZigBee channels. As expected, ZigBee channels 11 to 14, which overlap Wi-Fi channel 1 (see Figure 15), are the busiest channels. All others were relatively quiet, including channels 15, 20, 25 and 26, which do not overlap with the Wi-Fi channels used in the hospital.

Table 32 – ZigBee report

Channel	Grade	Avg. Peak (dBm)
11	91 (A-)	-49.28
12	84 (B)	-37.00
13	85 (B)	-35.22
14	93 (A-)	-35.00
15	98 (A)	-34.56
16	96 (A)	-35.28
17	94 (A)	-36.56
18	94 (A)	-34.39
19	97 (A)	-37.72
20	98 (A)	-33.67
21	97 (A)	-34.11
22	95 (A)	-34.11
23	95 (A)	-40.39
24	98 (A)	-45.22
25	99 (A)	-56.94
26	99 (A)	-61.89

6.2 Physical inspection

The floor plan of the in-patient area in which the system was commissioned is shown in Figure 91. Rooms R211 to R213 are private rooms; whereas rooms R201 to R209 are semiprivate rooms. The floor also has a treatment room (TREAT. ROOM), a nursery, a consultation room (CONS. ROOM) and staff exclusive areas, which include a pharmacy (PH), a reception (REC), a nurses' room and a support area. Blue rectangles refer to Wi-Fi access points (APs).




Figure 91 – Floor plan of the in-patient area on the second floor.

As ZigBee coordinators and routers are mains powered, it was necessary to locate available mains plugs or places where new plugs could be installed. Patient rooms had no spare mains plugs and it was not possible to install new ones. However, all other rooms had spare plugs. The hall had no spare mains plugs, but it was possible to add new ones over each emergency light, as indicated by small yellow numbered squares. These plugs are positioned at height of, approximately, 2.2 m from the ground level. As shown in Figure 91, spots 1 and 4 are close to a Wi-Fi AP and, consequently, should be avoided⁴³ [104].

The monitoring station should be positioned where nurses are constantly present so that alarms can be heard. Possible places are the reception, the nurse's room, and the pharmacy.

⁴³ Jennic recommends a physical separation from a WLAN AP of 8 m for co-channel operation to achieve a PER of 1%. For a particularly saturated WLAN link, 9-10 m may be necessary.

6.3 Connectivity tests

Initially, it was necessary to program the sensors' transmit power based on the desired range. It was necessary that coordinators or routers positioned at one of the spots defined on physical survey could communicate with sensors. In this case, the range should be at least 13 m (the approximate distance between spot 1, on Figure 91, and room R211). The necessary transmit power was estimated using the path loss equation adapted to a fading channel [57]:

$$P_0 = P_r + F_m + (10 * n * \log_{10}(f)) + (10 * n * \log_{10}(d)) - (30 * n) + 32.44, \quad (16)$$

where:

P_0 is the transmit power, including the antenna gain, in dBm;

P_r is the receiver sensitivity, in dBm;

F_m is the fade margin, in dB;

n is the path loss exponent;

f is the signal frequency in MHz; and

d is the range in meters.

The fade margin of 6 dB to 10 dB is typical for indoor applications, whereas the path loss exponent typically ranges from 2.7 to 4.3 for an office building, without line of site [57].

For $P_r = -96.5$ dBm, $n = 3.5$ (average value for the path loss exponent for an office building), $f = 2480$ MHz (channel 26 central frequency), $d = 13$ m, and assuming $F_m = 10$ dB, then $P_0 = -1.3$ dBm. If sensors' transmit power is programmed to 0 dBm, an approximate range of 14.1 m is expected.

Several tests were performed to verify the wireless connectivity range between sensors based on the JN5139-M00 module [106], which is used on sensors. Two boards based on these modules were used: one was programmed to continuously send messages (transmitter device), whereas the other (receiver device) was programmed to blink a LED every time a message was correctly received. These visual clues were used to classify the wireless links. A link was considered satisfactory if messages sent from the transmitter device placed on area A were correctly received at any point of area B (the transmitter was fixed, whereas the receiver was mobile). Alternatively, the link was considered intermittent if messages could only be correctly received on some parts of area B. Finally, the link was considered unsatisfactory if messages could not be received

at any part of area B. As these tests could disturb patients, they were performed in an unoccupied in-patient area similar to the floor in which the prototype system was commissioned.

The results are shown in Figure 92 where lines and columns labels refer to patient room numbers and other areas accessible to patients, namely, the consultation room (C), the nursery (N) and the treatment room (T). Tests started with the transmitter device placed at room 201 and the receiver device being moved from one room to the next one. The results that correspond to this setting are shown in the first line where dark blue cells symbolize satisfactory links (201-202, 201-203, 201-204), medium blue cells represent intermittent links (201-C, 201-205 and 201-206) and light blue cells stand for unsatisfactory links. Finally, diagonal cells are filled in black.

TX/RX	201	202	203	C	204	205	206	N	T	207	208	209	211	212	213
201	A	Dark Blue	Dark Blue	Medium Blue	Dark Blue	Medium Blue	Medium Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue
202	Light Blue	Black	Dark Blue	Dark Blue	Dark Blue	Dark Blue	Dark Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue
203	Light Blue	Light Blue	Black	Dark Blue	Dark Blue	Dark Blue	Dark Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue
C	Light Blue	Light Blue	Light Blue	Black	Dark Blue	Dark Blue	Dark Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue
204	Light Blue	Light Blue	Light Blue	Light Blue	B	Dark Blue	Dark Blue	Dark Blue	Light Blue	Dark Blue	Dark Blue	Dark Blue	Light Blue	Light Blue	Light Blue
205	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Black	Dark Blue	Dark Blue	Light Blue	Dark Blue	Dark Blue	Dark Blue	Light Blue	Light Blue	Light Blue
206	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Black	Dark Blue	Dark Blue	Dark Blue	Dark Blue	Dark Blue	Light Blue	Light Blue	Light Blue
N	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Black	Dark Blue	Dark Blue	Dark Blue	Dark Blue	Light Blue	Light Blue	Light Blue
T	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Black	Dark Blue	Dark Blue	Dark Blue	Light Blue	Light Blue	Light Blue
207	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	C	Dark Blue	Dark Blue	Dark Blue	Light Blue	Light Blue
208	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Black	Dark Blue	Dark Blue	Dark Blue	Light Blue
209	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Black	Dark Blue	Dark Blue	Dark Blue
211	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	D	Dark Blue	Dark Blue
212	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Black	Dark Blue
213	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Black

Figure 92 – Link survey involving all rooms accessible to patients.

As shown, considering the link categories proposed, most links are symmetrical, that is, for a node pair (A, B) the link quality from A to B is similar to the link quality from B to A. The only exceptions were 202-206 and 205-T.

Figure 92 shows four sub-matrixes, A to D, which represent geographic clusters. Under the propagation conditions encountered during the tests, most of the devices belonging to these clusters can hear each other’s transmissions. These geographic clusters can be used to define the necessary wireless infrastructure. In case of star networks being considered, four networks operating in different channels should be

employed. If 2-hop tree networks are considered, two networks operating in two different channels should be used. In this case, it is not possible to avoid the presence of hidden-nodes.

6.4 Range and link quality tests

Figure 93 shows the results of tests used to verify the range and quality of links between a network coordinator or router and a sensor. The results of additional tests are included in Appendix F.

Two devices were used, a coordinator and an end device (ED). The first one was based on the JN5139-M02 high power module [110], which has a receiver sensitivity of -100 dBm. The ED was based on the JN5139-M00 module and was programmed for a transmit power output of 0 dBm. From Equation (16), an expected range of 17.8 m is obtained for a high fade margin of 10 dB. Alternatively, for a low fade margin of 6 dB, an expected range of 24 m results. The application software executed by the coordinator was developed to average the Link Quality Indicator (LQI) values attached to a certain number of sequentially received messages⁴⁴. In case a message was not correctly received, the software assigned it an LQI value equal to zero. The ED was programmed to sequentially send messages at regular time intervals.


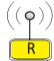

The icons shown in Table 33 were used to represent the network devices. Initially, the coordinator was positioned on one of the spots chosen during the physical inspection, as indicated in each setting shown in Figure 93. Then, the ED was carried to all rooms accessible to patients where average LQI values were measured. Each red number in Figure 93 represents an average LQI value measured, whereas its position in the picture match the spot at which the ED was positioned during the measurement. The areas marked in yellow represent sites where the average LQI value measured was equal to or less than 30 (according to Jennic, a LQI value equal or less than 30 is associated to a low quality link [107, 111]) or the reception was intermittent (INT). The areas marked in red correspond to sites where devices were unable to communicate.

⁴⁴ According to the IEEE 802.15.4-2003 specification, the LQI measures the received energy level and/or SNR for each received packet. In Jennic's implementation, the LQI value is proportional to the number of gain stages required to properly receive a message.



Figure 93 – Survey used to select adequate spots, as presented on (a) through (f), to position coordinators and routers.

Table 33 – Icons used to represent network devices.

Network device	Icon
PAN coordinator	
Router	
ED (sensor)	

Based on the test results obtained thus far, the following conclusions apply:

- The measured range involving a JN-5139M02 high-power module with a receiver sensitivity of -100 dBm and a regular JN5139-M00 module with a transmit power of 0 dBm is, approximately, 24 m if one or both devices are inside a room, or 40 m if both devices are placed at the corridor (line-of-sight).
- If star networks are used to relay the traffic generated by the devices included on geographic cluster A, a coordinator can be positioned at spot 2 (spot positions are shown in Figure 91) or at the reception. Spot 1 is an alternative. The pharmacy should be avoided.
- The traffic generated by devices included in the geographic clusters B, C and D can be relayed by star networks whose coordinators are positioned at spots 3, 5 and 6, respectively. The consultation room is not an adequate place to position a star coordinator that connects the devices belonging to geographic cluster D.
- Two 2-hop tree networks can be used to relay the traffic generated by sensors from geographic clusters (A, B) and (C, D). The first network can consist of a coordinator placed at spot 5 and a router placed at spot 6. Alternatively, the coordinator can be placed inside the consultation room and two routers at spots 5 and 6. The second network can consist of a coordinator placed at spot 1 or 2, whereas the router can be positioned at spot 3. Instead, the coordinator can be placed nearby the monitoring station, at the reception, whereas two routers can be positioned at spots 1 or 2 and 3. Collisions due to the presence of hidden-nodes are anticipated.

6.5 Experimental tests with no hidden-nodes

Long-term field tests performed in the hospital environment can provide a good estimative of the system performance because they can account for a) the effects of clock-drift in nonbeacon-enabled networks; b) the propagation effects caused by intermittent movement of people (including patients, health care providers, visitors and maintenance personnel), equipment and accessories (for instance, portable vital sign monitors, hospital chairs and meal carts); and c) the consequences of different levels of interference caused by other wireless networks that share the spectrum (mainly, the Wireless Local Area Network (WLAN)).

This section presents the results of six tests executed to estimate the performance of different networks in the absence of hidden-nodes and under different levels of WLAN interference.

6.5.1 Tests configuration

Three networks were evaluated: a ZigBee-based nonbeacon-enabled star network, a ZigBee-based nonbeacon-enabled 2-hop tree network and an IEEE 802.15.4-based beacon-enabled star network with guaranteed time slot (GTS) assigned to EDs.

The transmit power output used by EDs was increased to +2 dB (maximum transmit power for JN5139-M00 modules) in order to improve the quality of the links. Though it represents an increase in the transmission current of 14% (from 37 mA to 42 mA), the corresponding lifetime decrease is not expressive because, for ECG sensors, the transceiver consumption represents just 3.2% of the total energy consumption of the wireless module.

EDs were programmed to generate the traffic of ECG sensors that transmit compressed data. However, instead of just transmitting a sequence of dummy bytes, EDs include the following information into messages' payload: a) the cumulative number of messages lost because an acknowledge frame was not received after all possible retries; b) the cumulative number of failed transmission attempts (nonbeacon-enabled networks only); c) LQI values associated to received beacons (beacon-enabled networks only) and acknowledgment frames; and d) the cumulative number of synchronization losses (beacon-enabled networks only). In multi-hop networks, before

relaying an incoming message, routers add to its payload the LQI value attached to it. As well, for each received data message, the coordinator adds the LQI value attached to it to the meaningful bytes it extracts from the payload and sends this data, through the serial interface, to the test program. Raw ECG transmission is not considered because the waveform resulting from compressed data is adequate for non-acute patient monitoring⁴⁵.

Based on the received data, the test program continuously computes the (cumulative) average DR for each sensor. Results are computed every two-seconds and updated whenever a new message is received. Also, at 2-second intervals, the test program computes the proportion of received messages considering a window holding the status (received or dropped) of the last twenty messages expected for each sensor. The results can range from 100% to 0, in 5% intervals (100%, 95%, 90% and so on). This indicator, which was designated DR Window, is used to analyze the performance of the network along small time intervals. Additionally, the test program continuously records all performance indicators computed (the average DR, the DR Window and performance indicators computed by devices), displays, in real-time, part of the data, and allows the user to configure the serial interface. The user interface of the test program is shown in Figure 94. Three slightly different versions of the test program were developed, one corresponding to each evaluated network.

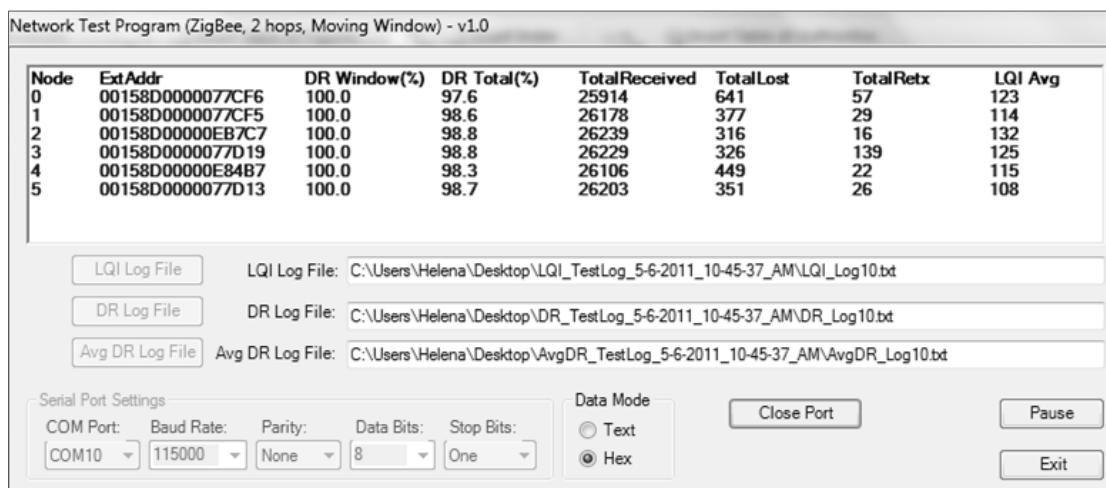


Figure 94 – Test routine user interface.

⁴⁵ This analysis was done by Otilio O. Rodrigues, MD, from Hospital Privado de Guimarães.

All tests used six test boards programmed as EDs. Hop-by-hop acknowledged transmission was used. Each ED was able to make up to three retransmissions. EDs operating in unslotted mode could backoff up to five times.

Since the hospital's WLAN is active on channels 1, 6 and 11, tests were performed on ZigBee channels 26 and 22. Channel 26 is free from WLAN interference, whereas channel 22 overlaps with Wi-Fi channel 11. By choosing these channels, it was possible to investigate the impact of the WLAN interference on the IEEE 802.15.4 and ZigBee networks.

The setting presented in Figure 95 (a) was used to perform tests using star networks. The coordinator was positioned on the hallway, near the entrance of the consultation room; whereas EDs were placed at rooms R201-R204. Figure 95 (b) presents the settings used for 2-hop tree topologies. The coordinator was placed on the consultation room whereas routers were placed on spots 5 and 6 (see Figure 91). No hidden-nodes were present on both settings. In addition to networking devices, the settings included the test application running on a PC connected to the coordinator; the SNA protocol analyzer running on a netbook; and the 2400E network adapter, which was placed outside the room, in the corridor, as shown in Figure 96.




Figure 95 – Settings used for tests using the (a) star topology with no hidden-nodes; and (b) 2-hop tree topology with no hidden-nodes.




Figure 96 – The test setting inside the consultation room (the 2400E, connected to the notebook, was placed outside the room, in the corridor).

The occupation of patient rooms varied. However, most of the time, rooms R201 and R204 were occupied by two patients, whereas rooms R202 and R203 were vacant. During the tests, EDs were positioned near the head of patients' beds, as shown in Figure 97. All tests were supervised by the hospital staff. Patients involved were aware of the objectives of the tests and agreed to have the test boards inside their rooms. No restriction was imposed regarding the use of any wireless equipment. Tests were run without interruptions.



Figure 97 – An ED on the head of a patient bed in room R202. The position of the device is assigned by the arrow.

6.5.2 Results

Table 34 lists the characteristics and the most relevant results. The first column lists the network configuration and the wireless channel used; whereas the following columns presents the test duration in hours, the time of the day each test started, and the resulting average DR. As shown, star networks that operated on channel 26 were able to deliver almost all generated messages, irrespective of the use of guaranteed time slots or the carrier sense multiple access (CSMA) mechanism. As anticipated by simulation and laboratory tests, the performance achieved by 2-hop networks is worse than the performance achieved by star networks. However, contrary to our expectations, the 2-hop tree network that operated on channel 22 achieved an average DR slightly higher than the one achieved by the 2-hop tree network that operated on channel 26. This is because the impact of WLAN interference on the DR is less than the effect of contention between devices, which is aggravated for multi-hop networks due to the relative increase in the traffic load.

Table 34 – Field test results. No hidden-nodes are present.

Protocol and other characteristics	Duration(h)	Starting time	Avg. DR (%)
IEEE 802.15.4, beacon-enabled, star network, channel 26 (GTS)	4.1	Morning	100
IEEE 802.15.4, beacon-enabled star network, channel 22 (GTS)	4.1	Morning	99.91
ZigBee, nonbeacon-enabled star network, channel 26	5.1	Afternoon	100
ZigBee, nonbeacon-enabled star network, channel 22	2.3	Afternoon	99.82
ZigBee, nonbeacon-enabled 2-hop, tree network, channel 26	16.7	Morning	98.56
ZigBee, nonbeacon-enabled 2-hop, tree network, channel 22	4.5	Afternoon	98.98

Table 35 details the DR results for star networks with no hidden-nodes. For each test and each ED, A – F, it is included: a) the number of generated messages; b) the average of the LQI values associated to received beacons (only for IEEE 802.15.4-based networks); c) the average of the LQI values associated to data messages; d) the number of messages that were not acknowledged; e) the number of failed transmission attempts; f) the number of lost messages; g) the number of duplicated messages; and h) the average DR. The quantities in a), b), d) and e) were continuously measured by EDs and the results were included in the payload of transmitted messages. The quantity in c) was

measured by the coordinator. The remaining results were continuously determined by the test routine based on data transferred by the coordinator.

Table 35 – Detailed results for field tests using star networks in the absence of hidden-nodes.

End Device	A	B	C	D	E	F
IEEE 802.15.4, beacon-enabled, star network, channel 26 (GTS)						
APP gen msg count	29,290	29,290	29,289	29,290	29,290	29,289
Beacon average LQI value	109	173	108	116	126	153
Data packet average LQI value	110	173	108	115	126	152
Missing ACK count	5	0	2	0	0	0
Failed transmission attempts	0	0	0	0	0	0
Lost msg count	5	0	2	0	0	0
Duplicated message count	0	0	0	0	0	0
Average DR (%)	99.98	100	99.99	100	100	100
IEEE 802.15.4, beacon-enabled star network, channel 22 (GTS)						
APP gen msg count	29,820	29,821	29,820	29,820	29,820	29,821
Beacon average LQI value	120	146	129	114	118	111
Data packet average LQI value	120	145	130	115	119	111
Missing ACK count	22	0	2	48	36	57
Failed transmission attempts	0	0	0	0	0	0
Lost msg count	22	0	2	48	36	57
Duplicated msg count	0	0	0	0	0	0
Average DR (%)	99.93	100	99.99	99.84	99.88	99.81
ZigBee, nonbeacon-enabled star network, channel 26						
APP gen msg count	36,484	36,484	36,478	36,403	36,480	36,479
Data packet average LQI value	150	155	122	138	139	85
Missing ACK count	0	0	1	0	0	0
Failed transmission attempts	0	0	0	0	0	0
Lost msg count	0	0	1	0	0	0
Duplicated msg count	0	0	0	0	0	0
Average DR (%)	100	100	100	100	100	100
ZigBee, nonbeacon-enabled star network, channel 22						
APP gen msg count	16,846	16,830	16,851	16,835	16,772	16,857
Data packet average LQI value	119	121	120	114	115	126
Missing ACK count	18	34	8	25	92	3
Failed transmission attempts	0	0	1	0	0	0
Lost msg count	18	34	9	25	92	3
Duplicated msg count	0	0	0	1	2	0
Average DR (%)	99.89	99.80	99.95	99.851	99.45	99.98

In the absence of hidden-nodes, all EDs associated to star networks achieved an excellent DR. This performance was, in part, guaranteed by the average good quality of links (only average values are presented, though all LQI values measured by devices have been recorded during tests).

As shown in Table 35, the number of channel access failures was small if compared to the number of messages that were not acknowledged. These results show that only

exceptionally EDs were unable to assess the channel idle. Consequently, most of the message losses occurred due to collisions or other effects the wireless channel is subject to, such as interference and momentary shadowing.

The number of duplicated messages is increased by the coordinator whenever it identifies a duplicate of a message previously received. As shown in Table 35, the number of duplicated messages and, consequently, the number of lost acknowledgment frames is very small.

Although networks subjected to WLAN interference have achieved a high average DR, the interference present on channel 22 forced EDs to retransmit several messages, as illustrated by the packet timelines shown in Figure 98. Both timelines were generated by the SNA protocol analyzer [40] and list packets sniffed from the IEEE 802.15.4 star networks tested using a 2400E network adapter [38].




Figure 98 – Packets captured for the IEEE 802.15.4-based networks with GTSS allocated to all EDs on (a) channel 26 and (b) channel 22.

Figure 98 (a) shows beacons and data packets sent by the EDs during the test on channel 26⁴⁶. Two red arrows point to beacon packets transmitted by the PAN coordinator at the beginning of two consecutive superframes. Following the first red arrow, a blue arrow points to a data packet transmitted by ED 0x0006. Five data packets

⁴⁶ Acknowledgment frames are shown too. However, since they are only presented on the first line, *Channel Summary*, they are hard to visualize.

transmitted by EDs 0x0005 – 0x0001 follow it. During this superframe and the followings, no retransmission is observed.

Figure 98 (b) shows packets sniffed during the test on channel 22. The red arrows point to beacon frames, whereas yellow arrows indicate packet retransmissions made by EDs. These retransmissions occurred because requested acknowledgment frames were not transmitted by the coordinator following data frames transmitted by the EDs; whereas acknowledgement frames were not sent because the coordinator was unable to correctly decode the data packets⁴⁷. Therefore, it is possible to conclude that the noise level increase on channel 22, in relation to channel 26, caused packet losses that, as a result, forced EDs to make retransmissions. Despite the allocation of GTSSs, if acknowledged messages have not been employed, all marked packets would have been lost. Consequently, in noisy environments, the use of acknowledge messages greatly contributes to increase the DR, even if a TDMA schema is used.

Table 36 presents the results for 2-hop networks on channels 26 and 22. For each ED, A – F, it includes a) the number of generated messages; b) the average value of the LQI attached to data messages transmitted by each ED (measured by routers); c) the average value of the LQI attached to data messages retransmitted by routers (measured by the coordinator); d) the number of messages not acknowledged to each ED; e) the number of messages not transmitted by the ED because the clear channel access (CCA) failed; f) the number of expected messages not received by the coordinator; g) the number of retransmitted messages received by the coordinator; and h) the average DR. The average DR and, consequently, the proportion of dropped messages, are calculated by the test routine based on information transferred by the coordinator.

⁴⁷ The network adapter could correctly decode the sniffed packets because its receive sensitivity is higher than the coordinator's. Moreover, its antenna has a higher gain.

Table 36– Detailed results for long tests using 2-hop networks in the absence of hidden-nodes.

End Device	A	B	C	D	E	F
ZigBee, nonbeacon-enabled 2-hop network, channel 26, no HNs						
APP gen msg count	120,096	119,081	115,801	118,170	120,098	119,444
ED to router link: avg. LQI value	144	127	154	140	140	122
Router to coordinator link: avg. LQI value	226	200	224	200	224	200
Missing ACK count	367	508	4,581	928	342	415
Failed transmission attempts	4	12	52	27	14	14
Lost msg count	378	1,393	4,672	2,303	375	1,030
Duplicated msg count	1	38	57	811	28	607
Average DR (%)	99.69	98.83	95.97	98.05	99.69	99.14
ZigBee, nonbeacon-enabled 2-hop network, channel 22, no HNs						
APP gen msg count	32,283	32,224	32,752	31,784	32,754	32,749
ED to router link: avg. LQI value	123	124	128	117	127	138
Router to coordinator link: avg. LQI value	191	189	210	190	212	210
Missing ACK count	493	601	5	1,004	2	2
Failed transmission attempts	16	10	0	27	0	0
Lost msg count	468	527	4	967	2	2
Duplicated msg count	104	104	0	165	0	0
Average DR (%)	98.57	98.37	100	97.06	99.99	99.99

As shown, both networks could deliver most messages generated by EDs. The worst individual DR result occurred for ED C, which lost 4.04% of all generated messages during the test using channel 26. In general, the quality of the links between devices was good.

In case of 2-hop networks, the test routine records only the total number of lost messages (*Lost msg count*) and the number of messages lost by EDs (*Missing ACK count* and *CCA failure count*). So, the number of messages lost by routers must be estimated. For instance, during the test on channel 26, ED A reported that 367 messages were not acknowledged. Additionally, it reported that 4 additional messages were lost because it could not access the channel. On the other hand, the coordinator reported that 378 messages were not received. In this case, except for a very small number of unacknowledged messages correctly received by the router, 7 messages generated by routers were lost.

Similarly to which occurred for star networks, most messages generated by EDs were lost due to collisions and interference, since the number of failed transmission attempts is relatively small. The number of duplicated messages was small too.

However, for EDs E and F, on channel 26, the proportion of duplicated messages to generated messages was equal to 0.69% and 0.51%, respectively.

LQI values provided by Jennic's implementation refer to the reliability of the channel, though it is not based on the bit error rate (BER) of the current packet, but on the number of gain stages necessary to correctly decode it. Although LQI values can be used to estimate the link quality, these values cannot be used to predict how a link will perform in the near future because short duration propagation effects, such as interference and fast fading, affect the link quality. For the purpose of illustrating this fact, refer to Figure 99. This data was recorded for the beacon-enabled star network based on the IEEE 802.15.4 protocol operating on channel 22. The horizontal axes of all graphs represent the sequence number of messages sent by an ED and correctly received by the coordinator. A message is generated every 500 ms. The curves shown in the first and second graphs represent the LQI values attached to received beacons and to successfully delivered messages, respectively⁴⁸. The third graph displays the cumulative number of lost messages. After successfully receiving 83 messages, 3 messages are lost. Despite the measured LQI values that remain constant, the number of lost messages continues to grow.




Figure 99 – LQI values and lost data packets count for an ED operating in the IEEE 802.15.4 star network on channel 22.

During all tests, the average DR was sequentially calculated at 2-second intervals. Also, a 10-second running window was used to continuously determine the DR achieved by the network and by individual EDs. In the first case, each DR value was

⁴⁸ Note that the Physical layer attaches an LQI values to each received packet. Consequently, there is no LQI information associated to a lost packet.

computed by dividing the total number of received packets by the number of expected packets since the beginning of the test. Consequently, if a few packets are lost in the beginning of the test, the DR values that follow these events are noticeably affected. However, as the test progresses, the loss of a few packets have little influence in subsequent DR values. On the other hand, the window-averaged results allow unbiased estimation of the network performance throughout the test.

Figure 100 (a) presents the cumulative distribution functions (CDFs) of the average DR for the networks that operated on channel 26, whereas Figure 100 (b) presents the CDFs of the window-averaged DR for the same networks. As shown, for all networks, the lower averaged DR values are higher than the lower window-averaged DR values. This inconsistency occurs because, in the first case, the effect of packet losses on the DR estimation varies with time, what is incorrect. For instance, while the CDF obtained from averaged DR values shows that the lower DR verified for the ZigBee-based 2-hop tree network was equal to 0.938, the CDF obtained from window-averaged DR values correctly shows that this value was equal to 0.783 (94 received messages out of 120 generated messages)⁴⁹.




Figure 100 – CDFs of DR for networks that operated on channel 26 in the absence of hidden-nodes: (a) obtained from the continuous values of the average DR and (b) obtained from the average of the DR over a 10-second sliding window.

⁴⁹ The range of possible DR values is a function of the window size. The window-averaged DR is obtained as the ratio of correctly received messages and expected messages. As a 10-second window was used, the number of expected messages is equal to 120 (one ED generates 2 messages per second; 6 EDs generate 120 messages every 10 seconds). For instance, if one message out of 120 messages is lost, the resulting window-averaged DR is equal to 0.992.

Figure 101 presents the CDFs of the window-averaged DR for the networks that operated on channel 22. Except for the 2-hop tree network, the CDFs indicate that the performance on channel 22 is worse than on channel 26. However, considering the particular hospital scenario in which tests were performed and the test conditions (good quality links and the use of acknowledged transmissions), even in presence of WLAN interference, the majority of the generated messages was delivered. The CDFs of the average DR are not presented for these networks because, as discussed, the average DR values just capture the overall performance of the network up to a certain point in time.




Figure 101 – CDFs of average DR over a 10-second sliding window for networks that operated on channel 22 in the absence of hidden-nodes.

Figure 102 presents the windowed DR computed for each ED (ED A to ED F) associated to the 2-hop tree network that operated on channel 26. The curves for the star networks are not presented because these networks only lost a small number of messages. Two graphs are used to present the data because the test was quite long. As shown, from time to time, two or more devices are unable to successfully deliver all generated messages. During the first minutes of test, no packet was lost. However, after, approximately, half an hour, the window-averaged DR values for ED C and ED D decreased for approximately half an hour before they started to increase again. A similar situation occurred for other pairs of EDs. As illustrated, despite the mechanisms employed by the Medium Access Control (MAC) layer, a contention-based network is subjected to collisions, which reduce its reliability.




Figure 102 – DR per 2-second calculated using a 20-message length window for the 2-hop tree network on channel 26 with no hidden-nodes.

6.5.3 Router deadlock

In multi-hop topologies, apart from EDs, routers also compete to access the wireless channel, which aggravates the contention problem. Figure 103 shows events observed experimentally concerning the exchange of messages generated by two end devices, ED1 and ED2. Both EDs are associated to a router, and this device is directly associated to the network coordinator. ED1 transmits a message just after ED2 receives an acknowledgement frame from the router. In this situation, the router can neither receive the message transmitted by ED1, because it has already initiated the CSMA-CA channel access mechanism, nor relay the received message, because it senses the channel busy.

ED1 extinguishes all possible retries and, consequently, its message is lost. Finally, as shown, the router relays the message transmitted by ED2. However, in other occasions, the router can drop the packet previously received if it reaches the maximum number of channel access.

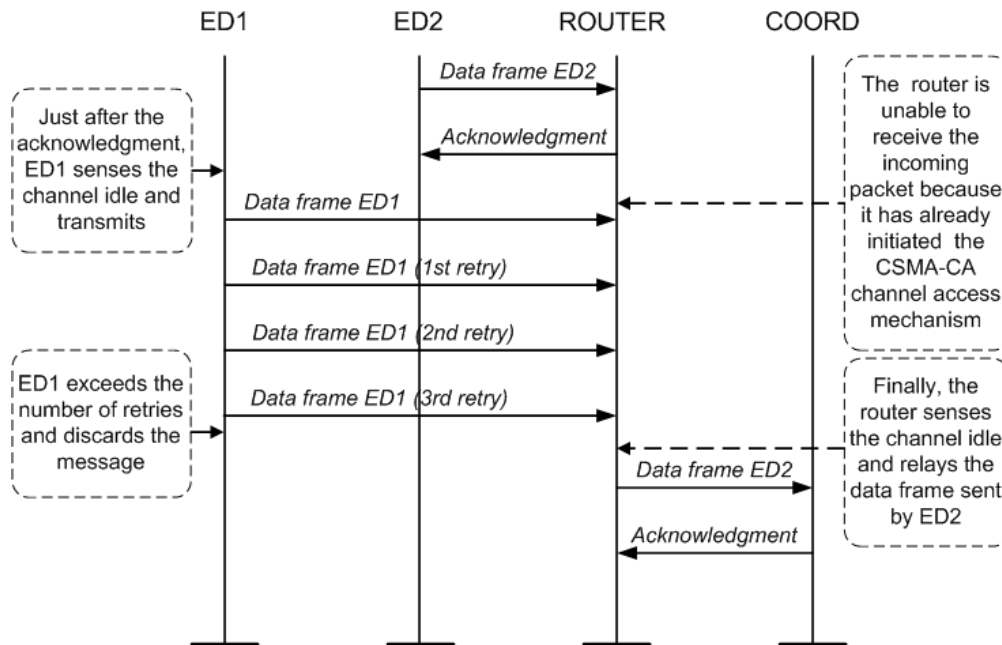


Figure 103 – ED1 accesses the channel just after a message from ED2 to the router is acknowledged. Unable to receive this message, the router just backs off while ED1 makes all possible retries.

The packet timeline shown in Figure 104 exemplifies the contention problem just described. Every 500 ms, each ED (network addresses 0x1430 – 0x1433) associated to a router (network address 0x0001) generated one message that was addressed to the network coordinator (network address 0x0000). The MAC parameter *macMinBE* of all devices was set to 3.




Figure 104 – Packet timeline captured during a test using a 2-hop ZigBee-based tree network.

The packet timeline shown in Figure 104 started with packet 1. Detailed information about these packets is also presented in the packet list shown in Figure 105, which also starts with packet 1. This packet was sent by ED 0x1431 and acknowledged by the router. The acknowledgment frame is followed by a failed transmission (packet 2) and two retries made by ED 0x1433. Then, the router finds the channel idle and relays packet 1. A few milliseconds after the acknowledgement frame sent by the coordinator, end device 0x1433 makes a successful retry and receives an acknowledgement frame transmitted by the router. In this case, despite having made three retries after a failed transmission, end device 0x1433 was able to deliver its message.

S...	Channel	Time	Time Delta	MAC Src	MAC Dest	NWK Src	NWK Dest	Protocol	Packet Type
41131	26	14:07:59.310	+00:00:00.242	0x1431	0x0001	0x1431	0x0001	ZigBee AF	0x0123: 0x12
41132	26	14:07:59.314	+00:00:00.004					IEEE 802.15.4	Acknowledgment
41133	26	14:07:59.318	+00:00:00.004	0x1433	0x0001	0x1433	0x0001	ZigBee AF	0x0123: 0x12
41134	26	14:07:59.324	+00:00:00.006	0x1433	0x0001	0x1433	0x0001	ZigBee AF	0x0123: 0x12
41135	26	14:07:59.330	+00:00:00.006	0x1433	0x0001	0x1433	0x0001	ZigBee AF	0x0123: 0x12
41136	26	14:07:59.335	+00:00:00.005	0x0001	0x0000	0x0001	0x0000	ZigBee AF	0x0123: 0x12
41137	26	14:07:59.339	+00:00:00.004					IEEE 802.15.4	Acknowledgment
41138	26	14:07:59.340	+00:00:00.002	0x1433	0x0001	0x1433	0x0001	ZigBee AF	0x0123: 0x12
41139	26	14:07:59.344	+00:00:00.004					IEEE 802.15.4	Acknowledgment
41140	26	14:07:59.350	+00:00:00.006	0x0001	0x0000	0x0001	0x0000	ZigBee AF	0x0123: 0x12
41141	26	14:07:59.354	+00:00:00.004					IEEE 802.15.4	Acknowledgment

Figure 105 – Packet list that includes the packets 1 and 2 shown in Figure 104.

The packet list shown in Figure 106 starts with packet 3, which was also sent by ED 0x1431 to the router. The acknowledgment frame sent by the router is followed by one failed transmission (packet 4) and three retries (the maximum number of retries allowed by the protocol) made by ED 0x1433. After these packets, the router successfully relays packet 3. In this case, ED 0x1433 failed to deliver its message.

S...	Channel	Time	Time Delta	MAC Src	MAC Dest	NWK Src	NWK Dest	Protocol	Packet Type
41166	26	14:08:00.310	+00:00:00.242	0x1431	0x0001	0x1431	0x0001	ZigBee AF	0x0123: 0x12
41167	26	14:08:00.314	+00:00:00.004					IEEE 802.15.4	Acknowledgment
41168	26	14:08:00.317	+00:00:00.004	0x1433	0x0001	0x1433	0x0001	ZigBee AF	0x0123: 0x12
41169	26	14:08:00.324	+00:00:00.006	0x1433	0x0001	0x1433	0x0001	ZigBee AF	0x0123: 0x12
41170	26	14:08:00.329	+00:00:00.005	0x1433	0x0001	0x1433	0x0001	ZigBee AF	0x0123: 0x12
41171	26	14:08:00.334	+00:00:00.005	0x1433	0x0001	0x1433	0x0001	ZigBee AF	0x0123: 0x12
41172	26	14:08:00.340	+00:00:00.006	0x0001	0x0000	0x0001	0x0000	ZigBee AF	0x0123: 0x12
41173	26	14:08:00.344	+00:00:00.004					IEEE 802.15.4	Acknowledgment

Figure 106 – Packet list that includes the packets 3 and 4 shown in Figure 104.

This situation could have been avoided if the router could process incoming packets during the backoff period. However, it is not possible for Jennic’s and other stack implementations. Such behavior is not an infringement of the IEEE 802.15.4 standard,

which specifies that, during the channel assessment part of the CSMA-CA mechanism, the device shall discard any received frames. However, it does not specify what procedure must be followed during backoff. In consequence, some manufacturers opted for ignoring incoming packets for the duration of backoff.

Alternatively, if the router had more priority to access the channel than the EDs, the contention problem described would occur less frequently. This can be achieved by setting different values of the MAC layer parameter *macMinBE* for each type of device. In order to confirm it, simulation was used to estimate the DR values considering a 2-hop tree network consisting of a variable number of EDs associated to a single router, as shown in Figure 107.




Figure 107 – 2-hop tree network.

Different values of the time interval between messages generated by end devices (0, 1 ms or 2 ms) were used in each simulation set. These small values were chosen to cause contention between devices. In addition, two different values of the router's parameter *macMinBE* (3 or 0) and a crescent number of end devices were used. All simulation runs considered that the parameter *macMinBE* was equal to 3 for all end devices. Each end device generated one ECG data message every 500 ms. Correctly received messages were acknowledged by the receiver device. Each run ended when the coordinator correctly received 100,000 messages.

Table 37 contains the results obtained where each *Delta time* column specifies the time difference between successive packets generated by all EDs. For instance, for *Delta time* = 0, all EDs generate a packet every $t = n * 500$ (ms) where n is the simulation cycle number. Alternatively, if *Delta time* = 1 ms and the number of EDs is equal to 2, ED₀ generates a packet every $t = n * 500$ (ms), whereas ED₁ generates a packet every $t = (n * 500) + 1$ (ms). Each *MinBE* column (*MinBE* = 0 and *MinBE* = 3) shows the value of the resulting DR. Every *Delta DR* column shows the relative improvement in the DR when the router *macMinBE* parameter is altered from 3 to 0.

Table 37 – DR results obtained from simulation.

Number of EDs	Delta time = 0		Delta DR	Delta time = 1 ms		Delta DR	Delta time = 2 ms		Delta DR
	macMinBE			macMinBE			macMinBE		
	3	0	3	0	3	0			
2	0.684	0.828	21%	0.745	0.902	21%	0.791	0.963	22%
3	0.544	0.710	31%	0.597	0.807	35%	0.572	0.797	39%
4	0.437	0.598	37%	0.465	0.691	49%	0.461	0.671	46%
5	0.352	0.503	43%	0.362	0.570	58%	0.376	0.589	57%
6	0.290	0.430	48%	0.295	0.492	67%	0.322	0.529	64%

As shown in Table 37, DR values increase significantly when the router's *macMinBE* parameter is set to 0. For instance, if two EDs generate messages at the same time and the router's parameter *macMinBE* is set to 3, the resulting DR value is equal to 0.684. However, if the router's *macMinBE* parameter is set to 0, the DR value increases to 0.828, an improvement of 21%. On the other hand, if the interval between messages generated by two devices is very small, this strategy alone is not enough to prevent packets from being lost.

6.5.4 Discussion

By observing the running windowed DR graphs shown in Figure 102, it is clear that test results are influenced by initial conditions, that is, by time relations between transmitters. Therefore, a reliable estimation of the expected performance of a network can only be obtained if sufficiently long tests are performed. Moreover, the global average value of the DR is not sufficient to estimate the performance of network. The performance of individual devices over time should be considered. Consequently, the measurement of window-averaged DR values decisively contributed to the estimation of the performance of the wireless networks evaluated.

6.6 Experimental tests with hidden-nodes

Up to this point, tests have not included the presence of hidden-nodes. As already discussed, the ZigBee protocol, which uses a CSMA mechanism, contains no provisions against hidden-node situations [118].

6.6.1 Tests configuration

The setting shown in Figure 108 was used to evaluate the perform of a nonbeacon-enabled ZigBee star networks in presence of hidden-nodes. The coordinator was placed just outside the consultation room, whereas room R201 and the treatment room each received three EDs. EDs placed on one room could not communicate or hear the others EDs' transmissions. Two tests were performed. The first one employed EDs that transmitted the same amount of traffic generated by ECG sensors; whereas the second one employed EDs that transmitted only heart rate traffic.

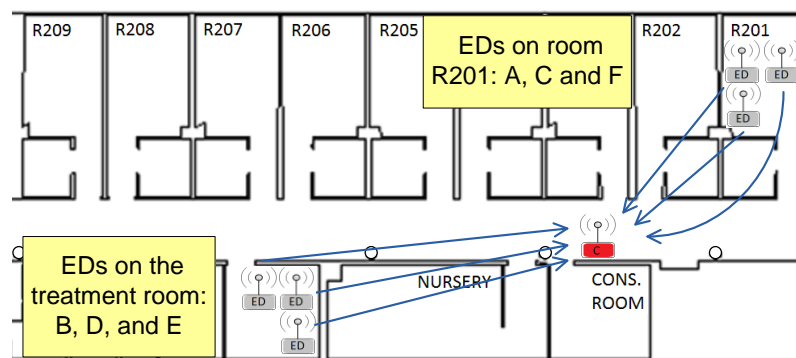


Figure 108 – Setting used to evaluate the performance of a ZigBee-based, nonbeacon-enabled star network with 50% of hidden-nodes.

6.6.2 Results for ECG traffic

During the first test, the devices were active for almost 8 h, but data recording was interrupted for approximately 50 minutes. As presented on Table 38, data were recorded for around 2.8 h and 4.2 h, respectively, during the first and second parts of the first test. In the first part of the test, the network performance was significantly worse than on previous tests. However, during the second part of the tests, the network achieved a high DR of 99.955%.

Table 38 – Average DR values for the test using a ZigBee-based, nonbeacon-enabled star network, on channel 26, 50% of hidden-nodes, ECG waveform transmission.

Part	Duration(h)	Starting time	Avg. DR (%)
1	2.8	Morning	83.960
2	4.2	Afternoon	99.955

Table 39 presents detailed test results. As shown, during the first part of the test, some devices lost a large proportion of the generated packets. For instance, ED A and

ED E lost 34.50% and 27.97% of the generated packets, respectively. As shown, in average, the quality of the links was adequate (the smaller average LQI value corresponds to the link between ED B and the coordinator and was measured during the second part of the test). As already observed during the previous tests, the number of channel access failures is much smaller than the number of not acknowledged messages.

Table 39 – Detailed results for the test using a ZigBee-based nonbeacon-enabled star network on channel 26 with 50% of hidden-nodes. ECG data transmission.

End Device	A	B	C	D	E	F
ZigBee, nonbeacon-enabled star network, channel 26, 50% hidden-nodes, ECG waveform transmission, part 1						
APP gen msg count	19,646	19,909	19,656	19,903	19,898	19,662
Data packet average LQI value	112	89	81	113	108	112
Missing ACK count	6,773	368	42	3,088	5,566	3,193
Failed transmission attempts	0	7	0	0	1	0
Lost msg count	6,775	375	42	3,087	5,566	3,193
Retx msg count	42	12	0	17	23	8
DR (%)	65.50	98.12	99.79	84.49	72.03	83.76
ZigBee, nonbeacon-enabled star network, channel 26, 50% hidden-nodes, ECG waveform transmission, part 2						
APP gen msg count	29,969	29,895	29,964	29,966	29,970	29,966
Data packet average LQI value	84	78	90	120	126	108
Missing ACK count	0	75	6	0	0	0
Failed transmission attempts	0	0	0	0	0	0
Lost msg count	0	75	6	0	0	0
Retx msg count	0	0	0	0	0	0
DR (%)	100	99.75	99.98	100	100	100

The CDFs of the window-averaged DR for the two parts of the test are presented in Figure 109. These distribution functions reflect the performance of the network, which during the first part of the test have achieved an average DR as low as 0.383%.




Figure 109 – CDFs of the windowed DR for the test using a ZigBee-based, nonbeacon-enabled star network, on channel 26, 50% of hidden-nodes. ECG traffic transmission (parts 1 and 2).

The windowed DR curves for individual EDs are presented in Figure 110. As shown, the curves for ED A and ED E have similar shapes, that is, when the window-averaged DR for ED A drops, the same effect is observed for the ED E. Also, when ED A stops losing packets, the same occurs to ED E. The same occurs to ED D and ED F and, just before the end of the first part of the test, to ED A and ED B. As shown in Figure 108, all these ED pairs are hidden-node pairs. For instance, ED A, which was placed at room R201, is hidden from ED E, which was placed at the treatment room.




Figure 110 – DR per 2-second calculated using a 20-message length window for a ZigBee, nonbeacon-enabled, star network on channel 26 with 50% hidden-nodes. ECG traffic transmission.

6.6.3 Results for heart rate traffic

A final test was employed six EDs that generated only heart rate data. The payload of data packets was reduced from 79 bytes to 8 bytes⁵⁰, whereas the period between messages was reduced from 500 ms to 3 s. The test lasted 10.2 h and an average DR of 99.90% was achieved. As shown on Table 40, which presents detailed results for individual EDs, the worst performance was computed for ED A, which achieved a DR equal to 99.71%. The average value of the LQI was relatively low for some EDs, which did not contribute to decrease the DR. No packet was lost due to failed transmission attempts and just few acknowledgement packets were lost.

Table 40 – Detailed results for the test using a ZigBee-based nonbeacon-enabled star network on channel 26 with 50% of hidden-nodes. Heart rate data transmission.

End Device	A	B	C	D	E	F
ZigBee, nonbeacon-enabled star network, channel 26, 50% hidden-nodes, HR tx						
APP gen msg count	12,219	12,219	12,219	12,219	12,219	12,219
Data packet average LQI value	66	77	53	120	43	78
Missing ACK count	40	1	21	12	28	12
Failed transmission attempts	0	0	0	0	0	0
Lost msg count	35	0	0	12	26	0
Retx msg count	12	0	0	0	5	3
Average DR (%)	99.71	100	100	99.90	99.79	100

As EDs generate data messages every 3 s, the length of the window used to continuously determine the window-averaged DR was increased to 12 s, which includes 120 expected messages. The CDF of the window-averaged DR is presented in Figure 112. As shown, the lowest DR value observed was equal to 0.892%, but most of the DR values measured were equal to 100%.




Figure 111 – CDFs of windowed DR for the test using a ZigBee-based, nonbeacon-enabled star network, on channel 26, 50% of hidden-nodes. Heart rate data transmission.

⁵⁰ It was not possible to further reduce the payload size because test data are transmitted by EDs.

Figure 112 presents graphs that show the values of the window-averaged DR computed for each ED. In two occasions, two pairs of devices, devices A and E, and devices A and D, which were hidden from each other, lost packets while contending to access the wireless channel. Comparing these graphs with the graphs shown in Figure 110, which represent the window-averaged DR measured for EDs that transmitted ECG data, it is possible to verify that, during contention periods, EDs that transmitted only heart rate data lost much fewer messages than EDs that transmitted ECG data. It occurs because much less traffic is generated by a heart rate sensor than by an ECG sensor (packets generated by heart rate sensors are smaller and less frequent than packets generated by ECG sensors), which decreases the probability of collisions between packets generated by heart rate sensors.




Figure 112 – DR per 2-second calculated using a 20-message length window for a ZigBee, nonbeacon-enabled, star network on channel 26 with 50% hidden-nodes. Heart rate traffic transmission.

6.6.4 Discussion

The results presented in this section evidence that CSMA-based networks cannot maintain a good reliability with hidden-nodes. However, if the traffic load is low and the length of transmitted packets is small, these networks are valid options to the TDMA-based GTS scheme supported by the IEEE 802.15.4 protocol, as shown by the results obtained using EDs that generated only heart rate traffic.

6.7 Comparing the communication performance of HM4ALL with a commercial vital signs monitoring system

The results shown in Table 41 were extracted from a study that evaluated the performance of FlexNet 802.11a wireless infrastructure. This study was supplied by Welch Allyn, which uses this infrastructure to remotely monitor patients using portable monitors and bedside monitors.

Hospitals B, D and H monitored patients using Micropaq wireless portable monitors (see Figure 10). The monitored areas range from 60 m² to 8,082 m² and included one or more departments. According to the authors of the study, the systems operated in environments free of interference from other wireless devices. The results were collected during a 7-day period. The average dropout was determined considering only communication loss intervals that do not exceed 120 seconds (longer periods were considered disconnections). As shown, the average dropout values are quite low, not exceeding 0.090%.

Table 41 – FlexNet (Welch Allyn) communication performance [210].

Hospital	B	D	H
Departments	Dialysis	Cardiac Care	Radiology, Echo, Telemetry
Total coverage area (m ²)	60	1,115	8,082
Number of Micropaq units	7	22	18
Seconds of dropout per 24 hours of patient monitoring	12	48	78
Average dropout time (%)	0.014	0.055	0.090

Unfortunately, it is not possible to directly compare the performance of the communication networks used by the prototype system and Welch Allyn's systems mainly because: a) the tests have followed different metrics (the prototype system tests considered the proportion of lost messages, whereas Welch Allyn's study considered the amount of time the system have been unable to communicate); b) the prototype system test did not consider patient mobility; c) the performance results for Welch Allyn's

systems were obtained in an environment free of WLAN interference⁵¹; and d) a variable number of monitoring units have been employed during Welch Allyn's tests, whereas the prototype system was tested using a fixed number of sensors. However, it is possible to infer that there are resemblances between the operation scenario in hospital B and the scenario used to evaluate the prototype system performance because roughly the same number of devices has been used and the mobility of patients in a dialysis unit is quite restricted. Therefore, it is possible to conclude that the performance observed for the IEEE 802.15.4 network on channel 26 was similar to the performance reported in Welch Allyn's study, which is an encouraging result. In the absence of hidden-nodes, the nonbeacon-enabled ZigBee-based network also performed as well as Welch Allyn's system.

6.8 Critical factors on the performance of ZigBee networks

This section presents issues that involve sending data in multi-hop ZigBee networks and the support to mobile sensors. Additionally, it discusses the use of end-to-end acknowledged messages to send high significance messages, such as alerts, and an alternative monitoring scenario which do not involve the transmission of ECG waveform data. Finally, it introduces a proprietary protocol recently proposed that is an adequate alternative to the IEEE 802.15.4 protocol for star networks comprised of a large number of devices that relay real-time traffic.

6.8.1 Mobility issues

The prototype system is based on spatially distributed networks. ZigBee coordinators and routers are static, whereas sensors, which join the network as EDs, are mobile. So, it is possible that a sensor moves away from its parent and communication is interrupted. Two scenarios are possible: a) inter-PAN mobility and b) intra-PAN mobility. In the first case, the sensor moves away from the coverage area provided by the network; whereas, in the second case, it moves away from its parent, but it is

⁵¹ FlexNet 802.11a infrastructure is based on the IEEE 802.11a protocol, which operates in the 5 GHz band. The WLANs installed on hospitals B, D and H operated in the 2.4 GHz band.

possible to find another parent among the devices that comprise the network. The second case is possible if the sensor operates on a multi-hop network.

In a star network where messages should be acknowledged, a sensor may conclude it should search for another network if the MAC layer reports the occurrence of one or more consecutive failed transmission attempts. Another strategy involves additionally monitoring LQI values attached to acknowledgment frames as an indicator of channel impairment. If the sensor concludes it has moved away from its parent, it scans all channels or a subset of channels and, if another network is found, it tries to associate. If successful, the whole process takes about two seconds⁵².

In a ZigBee multi-hop network, it is possible that the sensor moves away from its parent but enters an area covered by another potential parent operating in the same network. In this scenario, the sensor may decide to perform the orphaned device realignment procedure [91]. If successful, the device associates to another parent in less than one second⁵³. If it is unable to find another parent to associate with, it will be forced to search for another network.

During the reassociation process, the data flow is interrupted. The data flow from the sensor is resumed as soon as the device reassociates. However, the data flow in the other direction can only be resumed when a new route to the device is established. During the reassociation process, an ECG sensor would lose one or more packets. On the other hand, sensors that use a lower data transfer rate (such as the temperature sensor) would delay the transmission of data, but would not experience data loss.

As mentioned, in case a device is unable to communicate with its parent and must associate to another network, the data flow is interrupted for about two seconds. In the in-patient floor considered, this situation may occur if, for instance, a patient takes a

⁵² The time necessary to find another network and associate to it was measured using test boards based on the JN5139 wireless module. A timer was programmed to start counting just before the beginning of the scanning process and to stop counting just after the conclusion of the association process. According to these experiments, to perform these actions it is required between 1,912 and 1,919 ms.

⁵³ The period of time required to associate with another parent in the same network was measured using test boards based on the JN5139 wireless module. It was observed using the SNA protocol analyzer and was equal to 762 ms.

walk through the corridor or is seeing by a physician in the consultation room (see Figure 93). In several settings involving non-acute patients (these are the scenarios envisioned for health monitoring systems based on ZigBee and IEEE 802.15.4 protocols), this delay can be tolerated.

6.8.2 Use of end-to-end acknowledged messages

If, on one side, it is important to correctly deliver all generated data messages, it is essential to deliver data messages in case an unsafe condition is detected (for instance, fever or tachycardia).

In a star network in which the coordinator works as a sink, it is sufficient that a sensor employs acknowledgements to be sure that a message has correctly reached the destination. However, sensors that operate on multi-hop ZigBee-based networks may opt for end-to-end acknowledgments instead of hop-by-hop acknowledgments. In case end-to-end level acknowledgments are employed, in addition to hop-by-hop acknowledgments transmitted by the next-hop destination, acknowledgment frames are transmitted from the final destination, through the transmission path, to the device that originated the message.

The use of end-to-end acknowledgments increases the reliability of transactions above that available from the Network (NWK) layer alone because the application layer can request that lower layers make another attempt if a message is not confirmed by the final destination [220]. On the other hand, there is a significant cost increase in terms of network bandwidth, latency and power consumption. Therefore, the use of end-to-end acknowledgments should be restricted not to overload the network.

6.8.3 An option to IEEE 802.15.4: the eLPTR protocol

The results have revealed that the use of distributed star networks based on beacon-enabled mode of the IEEE 802.15.4 protocol is promising when used to concurrently monitor several patients using ECG sensors. However, the GTS scheme defined by this protocol can support a maximum of only seven devices. The recently proposed eLPTR (enhanced Low Power Real Time) protocol [2] cleverly addresses this problem.

The eLPTR protocol is a proprietary protocol that supports real-time traffic in star WSNs [2]. The protocol improves bandwidth utilization and increases the number of

supported devices through the division of the superframe into a much larger number of slots (500, in the current implementation) than the IEEE 802.15.4 (16 slots). This feature increases significantly the granularity of slot allocation in the CFP, avoiding the waste of bandwidth and, consequently, contributing to increase the throughput efficiency and the number of supported nodes.

Another interesting feature is the provision of a larger set of options for the superframe period, to closely match the packet generation interval imposed by the application. This feature is achieved by using 8 bits to encode the superframe period, allowing 256 options in comparison with the 15 options provided by the GTS scheme implemented by the IEEE 802.15.4 protocol [2].

6.9 Summary

The prototype system was commissioned in an in-patient area of a hospital covered by an IEEE 802.11b/g WLAN operating on channels 1, 6 and 11. The commissioning was followed by two sets of field tests, which employed six test boards programmed to associate as EDs. The first set of tests included no hidden-nodes. The test boards were programmed to generate the traffic of ECG sensors, tests was executed on channels 26 and 22, and employed three different network configurations: a) IEEE 802.15.4-based, star topology, beacon-enabled, GTSs assigned to EDs; b) ZigBee-based, nonbeacon-enabled, star topology; and c) ZigBee-based, non-beacon-enabled, 2-hop tree topology. All star networks achieved an average DR greater than 99.9%, except the nonbeacon-enabled ZigBee network that operated on the channel 22, which achieved an average DR of 99.8%. The 2-hop tree networks achieved an average DR of 98.6% for channel 26 and 99.0% for channel 22.

However, despite the promising average reliability, when the DR was computed for each ED that operated on a 2-hop network using a 10-second length running window, periods of contention between EDs were observed. During these periods, EDs could not maintain a high DR. These contention periods occurred because, from time to time, the instants of the transmission attempts made by two EDs tend to overlap as an effect of the different clock drifts of sensors.

Apart from collisions, during contention periods, devices that operate on multi-hop networks are subjected to router deadlock conditions. As the IEEE 802.15.4 protocol

does not specify if during backoff a device should interrupt or not the CSMA-CA mechanism to receive incoming packets, some vendors opt for neglecting these packets. Under this condition, which was termed router deadlock, multi-hop networks may experience contention issues that result in several messages being lost. If it is not possible to modify the MAC implementation, it was proposed reducing to zero the value of the MAC parameter *macMinBE* to minimize message losses. Simulations that corroborate with this assumption were presented.

Test results have also shown that the LQI values attached to received messages cannot be used to predict the quality of the link in the near future (for instance, after 500 ms from the reception of the last acknowledge frame) because the quality of the wireless link may change rapidly due to fast fading or interference from other devices that share the wireless channel. Finally, the analysis of the capture files recorded showed that, under WLAN interference, the number of retransmissions increased significantly. However, under the conditions experienced during the tests, the overall performance degradation of the networks was small.

The second set of tests was performed to measure the performance of networks under the presence of hidden-nodes. The test configuration consisted in a nonbeacon-enabled ZigBee star network comprised of six EDs where three of them were hidden from the other three. The first test used EDs programmed to generate the traffic of ECG sensors. During the first part of the test, the network achieved an average DR equal to 84.0%, what confirmed that the presence of hidden-nodes can severely compromise the performance of ZigBee-based networks. However, the second test used EDs programmed to send only heart rate data. By reducing the amount of traffic generated by EDs, the network achieved an average DR of 99.9%. This result shows that nonbeacon-enabled ZigBee networks are a valid option to the transport of vital signs in scenarios where wireless sensors generate low traffic.

The communication performance of HM4All was compared with the FlexNet 802.11a wireless infrastructure used by Welch Allyn to remotely monitor patients. For a restrict number of ECG sensors, IEEE 802.15.4 networks that allocate GTSs to sensors perform as well as Welch Allyn's infrastructure. In the absence of hidden-nodes or if only low amounts of traffic are generated by sensors, ZigBee-based nonbeacon-enabled networks also perform as well as Welch Allyn's infrastructure.

The support to mobile sensors and the use of end-to-end acknowledged messages is briefly discussed. After an end device recognizes it lost communication with its parent, it takes, approximately, 800 ms to reassociate to a new parent, in the same network. In case the end device must search for a new parent in another network, the process that involves nearby networks discovering and new parent association requires, approximately, 2 s. These data flow interruptions may be acceptable to some applications that involve remote monitoring of non-acute patients. The use of end-to-end acknowledged messages improves reliability, but increases the traffic load. However, the use of end-to-end acknowledged messages should be considered when transmitting high significance messages, such as alarms.

In the particular scenario in which the prototype system was commissioned, the best way to monitor ECG signals from up to six patients is to use four spatially distributed IEEE 802.15.4-based star networks that allocate one GTS to each sensor. On the other hand, if instead of transmitting ECG waveform samples and heart rate values, sensors transmitted only heart rate and other infrequent vital signs, much less traffic would be generated and a single multi-hop network would be able to relay all traffic generated by many sensors. This is a convenient solution to residential settings, nursing homes and several hospital environments, such as recovery, post-operative and other in-patient areas.

Chapter 7

Conclusions and future work

This chapter summarizes the objectives, contributions and conclusions of this thesis. It also proposes possible directions for future research.

7.1 Conclusions

The objective of this work was to study remote vital signs monitoring based on wireless sensor network (WSN) technologies, through the development, implementation and performance analysis of such a system.

The developed system, HM4All, which stands for *Health Monitoring for All*, is based on the ZigBee protocol. It consists of:

- Electrocardiogram (ECG) and axillary temperature sensors;
- ZigBee networking devices; and
- Three applications: ZigBee-to-IP gateway, Application Server and Monitoring Station.

Vital signs measured by wireless sensors are transported by ZigBee networking devices to a ZigBee-to-IP gateway, which processes and sends data to the Application Server. Remote clients execute the Web pages that comprise the Monitoring Station application to exhibit data and manage patients and sensors. Sensors, networking devices and the ZigBee-to-IP gateway application were developed under this work.

The ECG sensor measures electrocardiogram potentials, determines the heart rate, detects abnormal rhythms, namely tachycardia, bradycardia, background arrhythmia and

asystole⁵⁴, and sends data through a wireless channel. No cables between the sensor and the disposable electrodes are required, which causes the sensor to be unobtrusive and easy to clean. The ECG signal is sampled at a frequency of 200 Hz. However, the sensor can transmit 2:1 compressed ECG data, which does not affect the quality of the ECG signal for heart rate monitoring purposes. The sensor employs an established algorithm to accurately detect the QRS complexes and evaluates the heart rate by averaging the most recent ten consecutive R-R intervals. The results from tests that compared the heart rate measurements done by the sensor and by a vital signs bedside monitor revealed a maximum deviation of ± 2 beats per minute (bpm). The ECG sensor is powered by a CR2 lithium battery with power holding capacity of 850 mAh and a nominal voltage of 3.0 V. The average current consumption is equal to 12.3 mA, resulting in a battery life of, approximately, 70 hours.

The temperature sensor employs a thermistor medical probe to measure the axillary temperature. It was designed to measure temperature values in the range of 34 °C to 42 °C with a resolution of 0.1 °C and a precision of ± 0.2 °C. Tests compared temperature measurements done by the sensor and by a vital signs bedside monitor. The results revealed a maximum deviation of ± 0.3 °C, which is acceptable considering that the precision of the temperature measurements of the vital signs bedside monitor used is ± 0.1 °C. The temperature sensor is powered by a CR2540 coin battery with power holding capacity of 610 mAh and a nominal voltage of 3.0 V. The average current consumption is equal to 107 μ A, which corresponds to a battery lifetime of 241 days.

The ZigBee-to-IP gateway is a GUI-based application developed in C# language. It validates and processes data frames received from a ZigBee coordinator and sends processed data to the Application Server application through a remote connection. Additionally, it contains a user interface that exhibits sensor data, provides data recording functionalities and establishes and monitors remote connections.

Simulations and laboratory tests preceded HM4All commissioning in the hospital. Two simulation sets were executed. The first one considered an increasing number of ECG sensors operating in two different modes, raw or compressed data transmission,

⁵⁴ It should be noted that these automated detection procedures have been programmed in the sensor, but only tachycardia and bradycardia are featured in the visualization application in the form of alarms.

diverse network depths, and the use or not of acknowledgement frames. The second one considered different percentages of hidden-nodes for a star topology network consisting of a crescent number of ECG sensors. As expected, for dense networks, the best performance results were observed for sensors that transmitted compressed data and for networks that employed acknowledgement frames. Simulation results have also shown that it was necessary to restrict the network depth and the number of sensors to maintain a high delivery ratio (DR). Specifically, it was shown that to achieve a DR of 99.9% using ECG sensors that transmit compressed data and use acknowledged transmissions, it was necessary to restrict the number of ECG sensors to 23 in a star topology or to 7 in a 2-hop tree topology. In the presence of hidden-nodes, the number of nodes should be further reduced to 6 in a star topology. These results are due to the relatively small bandwidth offered by the ZigBee protocol and its susceptibility to collisions.

Laboratory tests followed simulations. These tests considered a crescent number of ECG sensors that transmitted raw or compressed data and different topologies. Up to 12 sensors were used and no hidden-nodes were considered. It was observed that a star network consisting of up to 12 nodes, on both modes, could achieve a DR of 99.9%. On the other hand, to maintain the same DR, a 2-hop tree topology should consist of up to 6 ECG sensors that transmit compressed data or 5 ECG sensors that transmit raw data. If more hops are added, fewer sensors should be employed to obtain a similar DR. The maximum and mean delay ECG messages experienced were also measured. For sensors that transmit raw and compressed data, the maximum delay values obtained were equal to 12 ms and 10 ms for star networks, and 10 ms and 17 ms for 2-hop tree networks. These delay values are within acceptable limits for real-time waveform transmission according to the IEEE 11073-00101-2008 standard.

Although sensors were programmed to transmit data at regular intervals, during network operation these intervals change due to clock drifts, which can aggravate or improve the overall network performance. Laboratory tests have evidenced significant DR drops that occur when two or more sensors try to access the wireless channel simultaneously or within a small time period. In order to correctly quantify and analyze the network performance, including transitory contention effects, long experimental tests were planned and executed in the hospital environment. New test routines were developed to measure quality of service indicators of individual sensors, including the DR calculated using a sliding window.

The prototype system was commissioned in an in-patient floor of Hospital Privado de Guimarães, an area served by a wireless local area network (WLAN) operating on channels 1, 6 and 11. The first set of field tests employed star and 2-hop tree topologies and hidden-nodes were avoided. Six test boards were programmed to transmit compressed ECG data and associate as end devices. As well as nonbeacon-enabled ZigBee networks, IEEE 802.15.4 beacon-enabled star network using guaranteed time slots (GTSs) assigned to EDs were evaluated. All star networks achieved an average DR greater than 99.9%, except the ZigBee network that operated on the channel 22, which achieved a DR of 99.8%. The 2-hop tree networks achieved an average DR of 98.6% for channel 26 and 99.0% for channel 22. However, despite the good average DR measured, when DR results were computed for individual sensors belonging to the 2-hop tree network that operated on channel 26, using a 10-second temporal window, it was observed that, occasionally, and for a variable time interval, two or more contending sensors were unable to correctly deliver their packets.

Apart from a small number of collisions caused by simultaneous channel assessments, the analysis of packet capture files revealed that the message losses were mainly caused by routers not being able to receive incoming packets during backoff. After receiving one packet it must relay, a router acknowledges the packet and then initiates the CSMA-CA mechanism. If another child end device transmits a new packet to the router, instead of interrupting the CSMA-CA to receive the incoming packet, the router backoffs because it senses the channel busy and, consequently, the packet is lost. If this sequence of events is repeated several times, the child end device may extinguish all possible transmission retries and the message is lost. On the other hand, the router may also drop the message it was trying to relay if it extinguishes the maximum number of transmission attempts. Unlike Jennic's CSMA-CA implementation, which ignores incoming packets for the duration of the backoff, other implementations, such as Texas Instruments', support it. However, such behavior is not an infringement of the IEEE 802.15.4 standard because it does not specify which procedure should be followed during backoff. In order to minimize message losses in multihop networks comprising routers that are unable to receive incoming packets during backoff, it was proposed changing the value of the MAC parameter *macMinBE*, from its default value of 3 to zero. Simulations that corroborate with this assumption revealed a good improvement in the resulting DR. For instance, if two ECG sensors always generate their messages at

the same time and their router parent has a *macMinBE* equal to 3, a DR of 68.4% is expected. On the other hand, if the *macMinBE* is set to zero, the DR will increase to 82.8%, an improvement of 21%.

The second set of field tests was executed using nonbeacon-enabled ZigBee-based star networks consisting of 50% of hidden-nodes. The first test employed six test boards programmed as end devices to generate the traffic of ECG sensors. Data was collected during two different periods. During the first part of the test, the network achieved a DR of 84.0%; whereas during the second part the DR was approximately equal to 100%. The low DR achieved during the first part of the test, which affected differently each end device, is explained by the increased number of collisions due to the presence of hidden-nodes. During the second part, due to clock drifts, nodes did not contend and, consequently, a high DR was achieved. This test demonstrated the negative effects of hidden-nodes in CSMA-based networks. The second test was executed using the same setting as before, but using test boards programmed to transmit only heart rate data. In this case, instead of transmitting an 114-byte packet every 500 ms, each test board transmitted a 44-byte packet every 3 seconds. The decrease in traffic caused a reduction in the number of collisions and a decrease in the contention periods, which resulted in a DR of 99.9%.

The system acceptance in the hospital environment was evaluated using questionnaires. All patients considered sensors lightweight and unobtrusive. All health care professionals considered the system useful. However, they have suggested including oxygen saturation and the blood pressure monitoring. Other suggestions included monitoring the respiration rate, providing patient reports and making it easier to access historical information.

Based on several simulations and experimental tests, this work concludes that ZigBee networks cannot reliably handle the transmission of a large number of ECG waveforms. In the hospital environment in which the prototype system was tested, we concluded that the best way to monitor ECG signals from up to six patients is to use four spatially distributed IEEE 802.15.4-based star networks that allocate one GTS to each sensor. However, if the traffic is considerably reduced, for instance, by just transmitting heart rate and temperature measurements, a single multi-hop network would be able to relay all traffic generated by a large number of sensors. This is a suitable solution to several hospital environments, such as recovery, palliative care and

emergency rooms. Another option would be transmitting only one ECG waveform when requested. On other even less demanding environments, which include assisted living facilities and residences, ZigBee can securely and reliably monitor health and wellness of chronically-ill and recovering patients.

In summary, we conclude that standard-based low-rate wireless personal area network (LR-WPAN) communication protocols can be used to remotely monitor vital signs if the traffic generated by sensors is restricted. The advantages of this technology include the size reduction of sensors, the reduction of the system cost, the increase of battery lifetime of sensors, and the possibility of mixing, in a single remote vital signs monitoring system, interoperable sensors from several manufacturers.

7.2 Future work

There are several possible directions of future work regarding the vital signs remote monitoring system developed, which include adding advanced features to the application software of the ECG sensor; developing an innovative physiological wearable sensor, such as a continuous blood pressure sensor; and validating emergent standard communication protocols.

The ECG sensor software could be further improved to include ECG waveform delineation to reduce the amount of data to be transmitted. Digital signal processing techniques to be applied include wavelet transforms [9, 139], and hidden Markov models [76]. Another possible future work is to develop an automatic heartbeat classifier that can be implemented by the ECG sensor's application software or as part of the ZigBee-to-IP application. In the first case, the classifier input could be either the ECG samples or resulting samples from the delineation process. In contrast, if implemented as part of the ZigBee-to-IP application, the classifier should accept, as input, the ECG samples that result from the delineation process. Several techniques could be employed to implement the classifier, including wavelet transforms [131, 132], support vector machines [166], and neural networks [97, 209]. Depending on the signal processing techniques chosen, the computational complexity involved may require the use of a dedicated microprocessor or a digital signal processor (DSP) to process sampled data. In this case, a major hardware design change should also be done. Both

ECG waveform delineation and heartbeat classification are major improvements that require long time to be implemented and evaluated against a standard database.

The development of a cuffless continuous blood pressure measurement is a very active research area, as discussed in Section 1.2. The majority of the systems proposed by academia have limitations that include motion artifacts; unsatisfactory accuracy; and the need of frequent calibrations. However, a recently released commercial product called BPro, a watch-like device that continuously determines blood pressure measurements from applanation tonography made at the radial artery, has proven that it is possible to determine systolic and diastolic pressure measurements with an unobtrusive device [186]. A possible research would involve applying micro-electrical-mechanical sensors (MEMS) and micro-manufacturing technologies, to reduce the device's size and increase its autonomy. Other approach involves improving upon other methods already under development by other groups [28, 30, 185, 194]. Despite the research line chosen, it would necessarily involve a multidisciplinary team, including mechanical, electrical and materials engineers.

Regarding alternative protocols, IEEE 802.11 networking has emerged as a viable solution for embedded applications as low-power chips and modules have been recently released [44, 68]. Low-power IEEE 802.11 devices have the advantages of native IP-network compatibility and well-known protocols and management tools. However, implementing IEEE 802.11 networking for embedded systems poses some challenges. Only IEEE 802.11b is presently available for low-power chips, which is a constraint in terms of data rate (the maximum data rate is 11 MHz with a downgrade to 1 MHz) and quality of service. Moreover, as opposed to laptops or mobile phones⁵⁵, embedded devices use Universal Asynchronous Receiver/Transmitter (UART) interface or Serial Peripheral Interface Bus (SPI) as the host connection, which can be an additional data rate limiting factor.

⁵⁵ For instance, Texas Instruments employs a Secure Digital Input Output (SDIO) interface to transfer data between the OMAP 5 family of Application processors and the WLAN function implemented by the WiLink family of mobile communication systems on a chip (SoC). The SDIO interface is roughly two times faster than a 4-bit SPI.

Emerging protocols, such as 6LoWPAN and ISA100.11a, can also be evaluated considering the requirements of patient monitoring under several scenarios. Like ZigBee, 6LoWPAN is based on the IEEE 802.15.4 Physical (PHY) and Medium Access Channel (MAC) layers and, consequently, offers a limited bandwidth if compared to IEEE 802.11. On the other hand, 6LoWPAN is IP-based, which allows devices to directly communicate with other IP-based devices without the need of complex gateways. ISA100.11a targets industrial applications. It uses the Physical (PHY) layer provided by the IEEE 802.15.4 protocol, but in addition to a CDMA-based MAC, it employs a Time Division Multiple Access (TDMA) MAC scheme based on the TSMP MAC protocol [172] to support low-latency applications and applications requiring specific bandwidth. It offers five different levels of quality-of-service and was designed to be extremely reliable.

Further investigation on the integration of wearable medical sensors, advanced signal processing techniques and low-power wireless communication protocols can lead to development of more unobtrusive, pervasive and dependable remote health monitoring systems.

Bibliography

1. Actibelt, 2010. http://www.actibelt.com/actibelt_flyer_march2010.pdf.
2. Afonso, J.A., Silva, H.D., Macedo, P. and Rocha, L.A. An Enhanced Reservation-Based MAC Protocol for IEEE 802.15.4 Networks. *Sensors*, 11. 3852-3873, 2011.
3. Age Concern. Personal Alarm Service, 2011. <http://www.aidcall.co.uk/personalcare/>.
4. Akyildiz, I.F., Weilian, S., Sankarasubramaniam, Y. and Cayirci, E. A survey on sensor networks. *Communications Magazine, IEEE*, 40 (8). 102-114, 2002.
5. AL-Khalidi, F.Q., Saatchi, R., Burke, D. and Elphick, H. Facial Tracking Method for Noncontact Respiration Rate Monitoring *Communication Systems Networks and Digital Signal Processing (CSNDSP), 2010 7th International Symposium on Newcastle upon Tyne 2010*.
6. Albaina, I.M., Visser, T., van der Mast, C.A.P.G. and Vastenburger, M.H. Flowie: A persuasive virtual coach to motivate elderly individuals to walk. In *Pervasive Computing Technologies for Healthcare, 2009. PervasiveHealth 2009. 3rd International Conference on*, (2009), 1-7.
7. Alert Me. AlertMe Home Monitoring, 2011. <http://www.alertme.com/products/home-monitoring>.
8. Alliance, Z. ZigBee RF4CE Overview, 2011. <http://www.zigbee.org/Specifications/ZigBeeRF4CE/Overview.aspx>.
9. Almeida, R., Martinez, J.P., Rocha, A.P. and Laguna, P. Multilead ECG Delineation Using Spatially Projected Leads From Wavelet Transform Loops. *Biomedical Engineering, IEEE Transactions on*, 56 (8). 1996-2005, 2009.

Bibliography

10. American Heart Association. Blood Pressure, 2011. <http://www.americanheart.org/presenter.jhtml?identifier=4473> (Jan 28th, 2011).
11. Amft, O. and Troster, G. On-Body Sensing Solutions for Automatic Dietary Monitoring. *Pervasive Computing Magazine, IEEE*, 8 (2)2009.
12. Anacom. Quadro Nacional de Atribuição de Frequências 2010/2011. Comunicações, I.-A.A.N.d. ed., 2010. <http://www.anacom.pt/download.jsp?contentId=1091654&fileId=1091653&channel=graphic>.
13. Anastasi, G., Borgia, E., Conti, M. and Gregori, E. IEEE 802.11b Ad Hoc Networks: Performance Measurements *Cluster Computing*, 8 (2-3). 135-145, 2005.
14. Andrushevich, A., Kistler, R., Bieri, M. and Klapproth, A. ZigBee/IEEE 802.15.4 Technologies in Ambient Assisted Living Applications *Pervasive Health 2010*, Munich, Germany, 2010.
15. Anglia2M2. Point to Point and Mesh Networking, 2011. http://www.anglia-m2m.com/Point_to_point.asp.
16. Baker, B.C. AN685 - Thermistors in Single Supply Temperature Sensing Circuits, Microchip, 1999. http://www.microchip.com/stellent/idcplg?IdcService=SS_GET_PAGE&nodeId=1824&appnote=en011704.
17. Baker, S.D. and Hoglund, D.H. Medical-Grade, Mission-Critical Wireless Networks [Designing an Enterprise Mobility Solution in the Healthcare Environment]. *Engineering in Medicine and Biology Magazine, IEEE*, 27 (2). 86-95, 2008.
18. Becker, D.E. Fundamentals of Electrocardiography Interpretation. *Anesthesia Progress* 53 (2). 53 - 64, 2005.

19. Benocchia, M., Tacconia, C., Farellaa, E., Beninia, L., Chiaria, L. and Vanzago, L. Accelerometer-based fall detection using optimized ZigBee data streaming *Microelectronics Journal*, 41 (11). 703-710, 2010.
20. Bianchi, V., Grossi, F., Matrella, G., Munari, I.D. and Ciampolini, P. A Wireless Sensor Platform for Assistive Technology Applications. In *Digital System Design Architectures, Methods and Tools, 2008. DSD '08. 11th EUROMICRO Conference on*, (2008), 809-816.
21. Bishop, J. Thermistor Temperature Transducer to ADC Application - Application Report SLOA052 Texas Instruments, 2000.
22. Bluetooth Special Interest Group. Bluetooth Basics, 2010. <http://www.bluetooth.com/English/Technology/Pages/Basics.aspx>.
23. Bourke, A.K., O'Donovan, K.J., Nelson, J. and OLaighin, G.M. Fall-detection through vertical velocity thresholding using a tri-axial accelerometer characterized using an optical motion-capture system *Engineering in Medicine and Biology Society, 2008. EMBS 2008. 30th Annual International Conference of the IEEE*, Vancouver, Canada, 2008.
24. Bourke, A.K., O'Brien, J.V. and Lyons, G.M. Evaluation of a threshold-based tri-axial accelerometer fall detection algorithm. *Gait and Posture* (26). 194–199, 2007.
25. Burke, T. Biomedical Engineering - ECG Assignment, 2007. <http://eleceng.dit.ie/tburke/biomed/assignment1.html>.
26. CardioNet. Welcome to CardioNet, 2011. <http://www.cardionet.com/index.htm>.
27. Casas, R., Marco, A., Plaza, I., Garrido, Y. and Falco, J. ZigBee-based alarm system for pervasive healthcare in rural areas. *Communications, IET*, 2 (2). 208-214, 2008.

Bibliography

28. Cattivelli, F.S. and Garudadri, H. Noninvasive Cuffless Estimation of Blood Pressure from Pulse Arrival Time and Heart Rate with Adaptive Calibration. In *Wearable and Implantable Body Sensor Networks, 2009. BSN 2009. Sixth International Workshop on*, (2009), 114-119.
29. Centralite. Centralite Residential Solutions, 2011. http://www.centralite.com/market_residential.php.
30. Chen, C., Kobayashi, T., Ichikawa, S., Takeuchi, Y. and Togawa, T. Continuous estimation of systolic blood pressure using the pulse arrival time and intermittent calibration. *Medical & Biological Engineering & Computing* 38. 569-574, 2000.
31. Chipara, O., Lu, C., Bailey, T.C. and Roman, G.-C. Reliable Patient Monitoring: A Clinical Study in a Step-down Medical Unit, Washington University in Saint Louis Saint Louis, MO, USA, 2009.
32. Comverge. Empowering Consumers to Control Energy Usage, 2011. <http://www.comverge.com/residential-consumer>.
33. Corventis. Avivo Mobile Patient Management (MPM) System, 2010. <http://www.corventis.com/AP/avivo.asp>.
34. Crossbow Technology. MicaZ. <http://www.xbow.com/Products/productdetails.aspx?sid=164>.
35. Culler, D. Secure, low power, IP-based connectivity with IEEE 802.15.4 wireless networks. *Industrial Embedded Systems* 2007.
36. Curtis, D., Shih, E., Waterman, J., Guttag, J., Bailey, J., Stair, T., Greenes, R.A. and Ohno-Machado, L. Physiological signal monitoring in the waiting areas of an emergency room *3rd International Conference on Body Area Networks - BodyNets '08*, ICST, Tempe, Arizona, USA, 2008.
37. da Frota Mattos, M.K., Biagioni, P.H. and Bassi, W. Electric field measurement on time domain generated by corona on insulators on distribution systems. In

-
- Electrical Insulation, 1996., Conference Record of the 1996 IEEE International Symposium on, (1996), 328-330 vol.321.*
38. Daintree Networks. 2400E Sensor Network Adapter 2010. <http://www.daintree.net/sna/adapter.php>.
 39. Daintree Networks. Getting Started with ZigBee and IEEE 802.15.4 2008. http://www.daintree.net/downloads/whitepapers/zigbee_primer.pdf.
 40. Daintree Networks. Sensor Network Analyzer (SNA), 2010. <http://www.daintree.net/sna/sna.php>.
 41. Daintree Networks. Understanding 802.15.4™ and ZigBee™ Networking, 2007.
 42. Daintree Networks. ZigBee Specification Comparison Matrix. <http://www.daintree.net/resources/spec-matrix.php>.
 43. Daintree Networks. ZigBee Specification Comparison Matrix. <http://www.daintree.net/resources/spec-matrix.php>.
 44. Dobkin, D.M. and Aboussouan, B. Low Power Wi-Fi (IEEE 802.11) for IP smart objects, GainSpan Corporation, Los Gatos, CA, 2009. http://www.gainspan.com/technology/white_papers.php.
 45. Dorothy, C., Eugene, S., Jason, W., John, G., Jacob, B., Thomas, S., Robert, A.G. and Lucila, O.-M. Physiological signal monitoring in the waiting areas of an emergency room *Proceedings of the ICST 3rd international conference on Body area networks*, ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), Tempe, Arizona, 2008.
 46. EDMI. Smart Metering Solutions, 2011. <http://www.edmi-meters.com/ProductsList.aspx?L0=1>.
 47. Elahi, A. and Gschwender, A. *ZigBee Wireless Sensor and Control Network*. Pearson Education, Inc., Boston, MA, 2009.

Bibliography

48. Elster. Smart Metering/Smart Grid, 2011.
http://www.elstersolutions.com/en/smart_metering_smart_grid.html.
49. Ember. ZigBee Security. <http://portal.ember.com/node/685>.
50. EMFIT. Non-Contact Vitals Monitor and IP Nurse-Call System, 2010.
http://www.emfit.com/en/care/products_care/vitals-monitor-and-nurse-call/.
51. Ericsson. Ericsson Mobile Health.
http://www.ericsson.com/solutions/enterprise/products/mhealth_solutions.shtml.
52. Estudillo-Valderrama, M.A., Roa, L.M., Reina-Tosina, J. and Naranjo-Hernandez, D. Design and Implementation of a Distributed Fall Detection System—Personal Server. *Information Technology in Biomedicine, IEEE Transactions on*, 13 (6). 874-881, 2009.
53. ETSI Directive 2009/114/EC of the European Parliament and of the Council of 16 September 2009 - Amending Council Directive 87/372/EEC on the frequency bands to be reserved for the coordinated introduction of public pan-European cellular digital land-based mobile communications in the Community. *Official Journal of the European Union* (16 September 2009)2009.
54. Exacon. Medical Temperature Sensors and Probes.
<http://www.exacon.com/Global/Global02/Global02FR2.htm#Neonatal>.
55. Falck, T., Baldus, H., Espina, J. and Klabunde, K. Plug 'n play simplicity for wireless medical body sensors. *Mobile Networks and Applications*, 12 (2-3)2007.
56. Faludi, R. *Building Wireless Sensor Networks*. O'Reilly, Sebastopol, CA, 2010.
57. Farahani, S. *ZigBee Wireless Networks and Transceivers*. Elsevier, Oxford, UK, 2008.

-
58. Fattah, H. and Leung, C. An overview of scheduling algorithms in wireless multimedia networks. *Wireless Communications, IEEE* 9(5). 8, 2002.
 59. Fernandez-Lopez, H. Interview with Warren Sandberg, MD, PhD (Associate Professor at Harvard Medical School, Massachusetts General Hospital Physician and Co-Project Leader of the CIMIT OR of the Future Project) - Monitoring Systems: Present Issues and Improvements, Boston, MA, 2009.
 60. Fernandez-Lopez, H., Afonso, J.A., Correia, J.H. and Simões, R. Experimental Evaluation of IEEE 802.15.4/ZigBee for Multi-Patient ECG Monitoring *eHealth2010*, Casablanca, Morocco, 2010, 8.
 61. Fernandez-Lopez, H., Afonso, J.A., Correia, J.H. and Simões, R. Extended Health Visibility in the Hospital Environment *BioDevices 2009 - 2nd International Conference on Biomedical Electronics and Devices*, Porto, Portugal, 2009, 422-425.
 62. Fernandez-Lopez, H., Afonso, J.A., Correia, J.H. and Simões, R. HM4All: A Vital Signs Monitoring System based in Spatially Distributed ZigBee Networks *Pervasive Computing Technologies for Healthcare (PervasiveHealth), 2010 4th Intl. Conf. on*, IEEE Munich, Germany, 2010, 1-4.
 63. Fernandez-Lopez, H., Macedo, P., Afonso, J.A., Correia, J.H. and Simões, R. Evaluation of the Impact of the Topology and Hidden Nodes in the Performance of a ZigBee Network. Hailes, S., Sicari, S. and Roussos, G. eds. *S-Cube 2009*, Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering, Pisa, Italy, 2009, 256-271.
 64. Fernández-López, H., Macedo, P., Afonso, J.A., Correia, J.H. and Simoes, R. Performance evaluation of a ZigBee-based medical sensor network. In *Pervasive Computing Technologies for Healthcare. PervasiveHealth 2009. 3rd International Conference on*, (2009), 1-4.

Bibliography

65. Fook, V.F.S., Leong, K.P., JianZhong, E.H., Jayachandran, M., Wai, A.A.P., Biswas, J., WeiSi, L. and Yap, P. Non-intrusive respiratory monitoring system using Fiber Bragg Grating sensor. In *e-health Networking, Applications and Services, 2008. HealthCom 2008. 10th International Conference on*, (2008), 160-164.
66. Fraile, J.A., Bajo, J., Corchado, J.M. and Abraham, A. Applying Wearable Solutions in Dependent Environments. *Information Technology in Biomedicine, IEEE Transactions on*, 14 (6). 1459-1467, 2010.
67. FTDI Chip. USB TTL Serial Cables, 2011. <http://www.ftdichip.com/Products/Cables/USBTTLSerial.htm>.
68. G2 Microsystems. Epsilon Family Wi-Fi Modules, Los Gatos, CA, 2011. <http://www.g2microsystems.com/products/modules.html>.
69. Gamboa, H., Silva, F. and Silva, H. Patient Tracking System - Continuous Monitoring and Location Solution for Ambient Assisted Living Facilities *Pervasive Health 2010*, Munich, Germany, 2010, 2.
70. GE Healthcare. ApexPro Enterprise-Wide Telemetry, 2010. http://www.gehealthcare.com/euen/patient_monitoring/docs/ApexPro_bro_M1171526_eng.pdf.
71. Ge Healthcare. ApexPro Enterprise-wide telemetry, 2009. http://www.gehealthcare.com/euen/patient_monitoring/docs/ApexPro_bro_M1171526_eng.pdf.
72. GE Healthcare. CARESCAPE V100 Vital Signs Monitor, 2010. http://www.gehealthcare.com/euen/patient_monitoring/docs/V100_bro_M1159221_eng.pdf.
73. Gislason, D. *ZigBee Wireless Network*. Elsevier, Oxford, UK, 2008.

-
74. Gnawali, O., Fonseca, R., Jamieson, K. and Levis, P. CTP: Robust and Efficient Collection through Control and Data Plane Integration. Stanford Information Networks Group Technical Report SING-08-02, 2008.
 75. Golmie, N., Cypher, D. and Rebala, O. Performance Analysis of Low Rate Wireless Technologies for Medical Applications. *Computer Communications*, 28. 1266-1275, 2005.
 76. Gomes, P.M.P.R. Sistema para aquisição, monitorização e análise da actividade cardíaca *Industrial Electronics Department*, University of Minho, Guimarães, 2010.
 77. Guerra, D., Gawlick, U. and Bizarro, P. An integrated data management approach to manage health care data *Proceedings of the Third ACM International Conference on Distributed Event-Based Systems*, ACM, Nashville, Tennessee, 2009.
 78. Hairong, Y., Hongwei, H., Youzhi, X. and Gidlund, M. Wireless sensor network based E-health system - implementation and experimental results. *Consumer Electronics, IEEE Transactions on*, 56 (4). 2288-2295, 2010.
 79. Hamilton, P.S. and Tompkins, W.J. Quantitative Investigation of QRS Detection Rules Using the MIT/BIH Arrhythmia Database. *IEEE Transactions on Biomedical Engineering*, BME-33 (12). 1157-1165, 1986.
 80. Hande, A., Polk, T., Walker, W. and Bhatia, D. Self-Powered Wireless Sensor Networks for Remote Patient Monitoring in Hospitals. *Sensors* 6(9). 1102-1117, 2006.
 81. Harris, M. and Habetha, J. The MyHeart project: A framework for personal health care applications. In *Computers in Cardiology, 2007*, (2007), 137-140.
 82. Hart Communication Foundation. Wireless Hart Technology, 2011. http://www.hartcomm.org/protocol/wihart/wireless_technology.html.

Bibliography

83. He, G., Wang, X. and Sun, G. Design of a Greenhouse Humiture Monitoring System Based on ZigBee Wireless Sensor Networks *Frontier of Computer Science and Technology (FCST), 2010 Fifth International Conference on*, Changchun, China, 2010.
84. HealthGrades. HealthGrades Hospital Quality in America Study, 2009. <http://www.healthgrades.com/press-releases/index.cfm?fuseaction=modNBG&modtype=b2b&modact=hospitalResearchRpts§ion=0> (Nov 20, 2009).
85. Hilton, P.A. *Fundamental Nursing Skills* Whurr Publishers, London and Philadelphia, 2004.
86. Hwang, L.-J., Sheu, S.-T., Shih, Y.-Y. and Cheng, Y.-C. Grouping strategy for solving hidden node problem in IEEE 802.15.4 LR-WPAN *Wireless Internet, 2005. First International Conference on* Budapest, Hungary, 2005, 26-32.
87. IEEE. Health Informatics - PoC Medical Device Communication - Part 00101: Guide - Guidelines for the Use of RF Wireless Technology *IEEE Std 11073-00101-2008*, 2008, 1-99.
88. IEEE. IEEE 11073 - Medical Device Informatics and Interoperability. <http://www.11073.org/>.
89. IEEE IEEE Health informatics--Personal health device communication Part 20601: Application profile--Optimized Exchange Protocol Amendment 1. *IEEE 11073-20601a-2010 (Amendment to IEEE Std 11073-20601-2008)*. 1-119, 2011.
90. IEEE. 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: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, 2007.

-
91. IEEE. IEEE Std 802.15.4-2003: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks IEEE ed., 2003.
 92. IEEE. IEEE Std 802.15.4™-2006 ISBN 0-7381-4996-9 SH95552, New York, USA, 2006.
 93. IEEE. Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (WPANs) - Amendment 1: Add Alternate PHYs, 2007.
 94. IEEE. Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (WPANs) - Amendment 2: Alternative Physical Layer Extension to support one or more of the Chinese 314-316 MHz, 430-434 MHz, and 779-787 MHz bands, 2009.
 95. IEEE. Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (WPANs) - Amendment 3: Alternative Physical Layer Extension to Support the Japanese 950 MHz Band, 2009.
 96. Ilyas, M. and Mahgoub, I. (eds.). *Handbook of Sensor Networks: Compact Wireless and Wired Sensing Systems*. CRC Press, 2004.
 97. Inan, O.T., Giovangrandi, L. and Kovacs, G.T.A. Robust Neural-Network-Based Classification of Premature Ventricular Contractions Using Wavelet Transform and Timing Interval Features. *Biomedical Engineering, IEEE Transactions on*, 53 (12). 2507-2515, 2006.
 98. Intelesens. Aingeal Patched-based Vital Signs, 2010.
<http://www.intelesens.com/pdf/aingealdatasheet.pdf>.
 99. ISO. ISO 13485:2003, 2011.
http://www.iso.org/iso/catalogue_detail?csnumber=36786.

Bibliography

100. ISO/IEC/IEEE. ISO/IEC/IEEE Health informatics--Personal health device communication--Part 20601: Application profile--Optimized exchange protocol *ISO/IEEE 11073-20601:2010(E)*, 2010, 1-208.
101. ISO/IEEE ISO/IEEE Health Informatics - Point-Of-Care Medical Device Communication - Part 10101: Nomenclature. *ISO/IEEE 11073-10101:2004(E)*. 0_1-492, 2004.
102. Jacobsen, L.A., Kent, M., Lee, M. and Mather, M. America's Aging Population, U. S. Census Bureau, Washington, DC, 2011. <http://www.prb.org/Publications/PopulationBulletins/2011/americas-aging-population.aspx>.
103. Jennic. Basic Operating System API Reference Manual - JN-RM-2021 - Rev. 3.1, 2007.
104. Jennic. Co-existence of IEEE 802.15.4 at 2.4 GHz - Application Note, 2008. http://www.jennic.com/files/support_files/JN-AN-1079%20Coexistence%20of%20IEEE%20802.15.4%20In%20The%202.4GHz%20Band-1v0.pdf.
105. Jennic. CodeBlocks IDE User Guide JN-UG-3028 Revision 1.8, 2008. http://www.jennic.com/files/support_files/JN-UG-3028-CodeBlocks-1v8.pdf.
106. Jennic. Data Sheet: JN5139-xxx-Myy JN-DS-JN5139-xxx-Myy 1v6 - IEEE802.15.4/ZigBee Module Family 2010. http://www.jennic.com/files/support_files/JN-DS-JN5139MO-1v6.pdf.
107. Jennic. How are the LQI and MSQ values for received data packets calculated?, 2008. <http://www.jennic.com/support/solutions/00004>.
108. Jennic. JenNet Protocol Stack, 2011. http://www.jennic.com/products/protocol_stacks/jennet.

-
109. Jennic. JN5139-EK010 ZigBee Evaluation Kit Getting Started, 2008. http://www.jennic.com/files/support_files/JN-UG-3030-JN5139-EK010-Getting-Started-1v2.pdf.
 110. Jennic. Product Brief – JN5139-xxx-Myy IEEE802.15.4/ZigBee Module Family, www.jennic.com, 2008. www.jennic.com.
 111. Jennic. Significance of LQI (forum post), 2008. <http://www.jennic.com/support/forums/thread.php?postID=0000003470>.
 112. Jennic. Technical Support: "Supress Route Discovery" 2010. <http://support.jennic.com/SelfService/Display.html?id=3554>.
 113. Jennic. ZigBee Stack User Guide JN-UG-3017 Rev. 1.5 2007.
 114. John G. Webster (Editor-in-Chief) *The Measurement, Instrumentation and Sensors Handbook*. CRC Press Boca Raton, FL, 1999.
 115. Kangas, M., Konttila, A., Lindgren, P., Winblad, I. and Jmsaa, T. Comparison of low-complexity fall detection algorithms for body attached accelerometers. *Gait and Posture*, 28. 285-291, 2008.
 116. Kangas, M., Konttila, A., Winblad, I. and Jamsa, T. Determination of simple thresholds for accelerometry-based parameters for fall detection. In *Engineering in Medicine and Biology Society, 2007. EMBS 2007. 29th Annual International Conference of the IEEE*, (2007), 1367-1370.
 117. Karantonis, D.M., Narayanan, M.R., Mathie, M., Lovell, N.H. and Celler, B.G. Implementation of a real-time human movement classifier using a triaxial accelerometer for ambulatory monitoring. *Information Technology in Biomedicine, IEEE Transactions on*, 10 (1). 156-167, 2006.
 118. Karl, H. and Willig, A. *Protocols and Architectures for Wireless Sensor Networks* Jhon Wiley & Sons Ltd, Chichester (West Sussex), UK, 2007.

119. Kiers, H.D., Hofstra, J.M. and Wetzels, J.F.M. Oscillometric blood pressure measurements: differences between measured and calculated mean arterial pressure. *The Netherland Journal of Medicine*, 6 6 (11). 6, 2008.
120. Kim, J.-S., Rahman, M.S., Kim, B.-Y., Jang, B.-J. and Kim, K.-D. DSP embeded hardware for non-contact bio-radar heart and respiration rate monitoring system *BioMedical Information Engineering, 2009. FBIE 2009. International Conference on Future 2009*, 4.
121. Klabunde, R.E. *Cardiovascular Physiology Concepts*. Lippincott Williams & Wilkins, Baltimore, MD, 2005.
122. Knight, H., Lee, J.-K. and Ma, H. Chair Alarm for patient fall prevention based on Gesture Recognition and Interactivity. In *Engineering in Medicine and Biology Society, 2008. EMBS 2008. 30th Annual International Conference of the IEEE*, (2008), 3698-3701.
123. Ko, J., Gao, T. and Terzis, A. Empirical Study of a Medical Sensor Application in an Urban Emergency Department *BodyNets '09 ICST*, Los Angeles, California USA, 2009, 8.
124. Kohler, B.U., Hennig, C. and Orglmeister, R. The principles of software QRS detection. *Engineering in Medicine and Biology Magazine, IEEE*, 21 (1). 42-57, 2002.
125. Konstantas, D. and Herzog, R. Continuous monitoring of vital constants for mobile users: the MobiHealth approach. In *Engineering in Medicine and Biology Society, 2003. Proceedings of the 25th Annual International Conference of the IEEE*, (2003), 3728-3731 Vol.3724.
126. Koubaa, A., Severino, R., Alves, M. and Tovar, E. Improving Quality-of-Service in Wireless Sensor Networks by Mitigating “Hidden-Node Collisions”. *Industrial Informatics, IEEE Transactions on* 5(3)2009.

-
127. Kwikset. Electronic Locks, 2011.
<http://www.kwikset.com/SmartSecurity/Electronic-Locks.aspx>.
128. Lan, M., Nahapetian, A., Vahdatpour, A., Au, L., Kaiser, W. and Sarrafzadeh, M. SmartFall: an automatic fall detection system based on subsequence matching for the SmartCane *BodyNets '09 Proceedings of the Fourth International Conference on Body Area Networks* Los Angeles, USA, 2009.
129. Li, Y., Chen, C.S. and Song, Y.Q. A Technical Review of Real-Time Protocols for Wireless Sensor Networks, INRIA - Institut National de Recherche en Informatique et en Automatique, 2007.
130. Liang, X. and Balasingham, I. Performance Analysis of the IEEE 802.15.4 based ECG Monitoring Network *Proceedings of the 7th IASTED International Conferences on Wireless and Optical Communications*, Montreal, Canada, 2007.
131. Llamedo, M. and Martinez, J.P. Heartbeat Classification Using Feature Selection Driven by Database Generalization Criteria. *Biomedical Engineering, IEEE Transactions on*, 58 (3). 616-625, 2011.
132. Llamedo, S.M. and Martinez, J.P. An ECG classification model based on multilead wavelet transform features. In *Computers in Cardiology, 2007*, (2007), 105-108.
133. Lo, B.P.L., Thiemjarus, S., King, R. and Yang, G.-Z. Body Sensor Network - A platform for pervasive healthcare monitoring. In *3rd International Conference on Pervasive Computing*, (London, UK, 2005).
134. Lorincz, J. and Begusic, D. Physical layer analysis of emerging IEEE 802.11n WLAN standard. In *Advanced Communication Technology, 2006. ICACT 2006. The 8th International Conference*, (2006), 6 pp.-194.
135. Lorincz, K., Malan, D.J., Fulford-Jones, T.R.F., Nawoj, A., Clavel, A., Shnayder, V., Mainland, G., Welsh, M. and Moulton, S. Sensor networks for

Bibliography

- emergency response: challenges and opportunities. *Pervasive Computing, IEEE*, 3 (4). 16-23, 2004.
136. Luštrek, M. and Kaluža, B. Fall Detection and Activity Recognition with Machine Learning. *Informatika*, 33. 205–212, 2009.
137. Lymberis, A. and Paradiso, R. Smart fabrics and interactive textile enabling wearable personal applications: R&D state of the art and future challenges. In *Engineering in Medicine and Biology Society, 2008. EMBS 2008. 30th Annual International Conference of the IEEE*, (2008), 5270-5273.
138. Mark, J.B. *Atlas of Cardiovascular Monitoring*. Churchill Livingstone, 1998.
139. Martinez, J.P., Almeida, R., Olmos, S., Rocha, A.P. and Laguna, P. A wavelet-based ECG delineator: evaluation on standard databases. *Biomedical Engineering, IEEE Transactions on*, 51 (4). 570-581, 2004.
140. Matos, A.C.L.P.d. Desenvolvimento de um sensor de electrocardiograma compatível com a tecnologia de redes sem fios ZigBee *Biomedical Engineering Department*, University of Minho, Braga, 2009.
141. Measurement Specialties. Zero-Power Resistance Characteristics. <http://www.meas-spec.com/searchresults.aspx?search=Zero-power%20resistance%20characteristic>.
142. MELZER, I., BENJUYA, N. and KAPLANSKI, J. Postural stability in the elderly: a comparison between fallers and non-fallers. *Age and Ageing*, 33 (6). 602-607, 2004.
143. Mentor Graphics. PADS tools, 2011.
144. Merriam-Webster. Merriam-Webster's Dictionary and Thesaurus 2007.
145. Meshnetics. Parking Lot Gets Smart with ZigBee, 2008. http://www.meshnetics.com/ZigBee_Parking_Automation_Case_Study.pdf.

-
146. MetaGeek. Chanalyzer 3, 2010. <http://www.metageek.net/products/chanalyzer-3>.
 147. MetaGeek. Wi-Spy 2.4x, 2010. <http://www.metageek.net/products/wi-spy-24x>.
 148. Min, S.D., Kim, J.K., Shin, H.S., Yun, Y.H., Lee, C.K. and Lee, M. Noncontact Respiration Rate Measurement System Using an Ultrasonic Proximity Sensor. *Sensors Journal, IEEE, 10* (11). 8, 2010.
 149. MindMe. Welcome to MindMe, 2011. <http://www.mindme.co.uk/>.
 150. MobiHealthNews. Wellcore unveils fall detection, activity monitor, 2011. <http://mobihealthnews.com/5923/wellcore-unveils-its-fall-prevention-and-activity-monitor/>.
 151. Montenegro, G., Kushalnagar, N., Hui, J. and Culler, D. Transmission of IPv6 Packets over IEEE 802.15.4 Networks *Internet Engineering Task Force RFC-4944*, IETF, 2007.
 152. Monton, E., Hernandez, J.F., Blasco, J.M., Herve, T., Micallef, J., Grech, I., Brincat, A. and Traver, V. Body area network for wireless patient monitoring. *Communications, IET, 2* (2). 215-222, 2008.
 153. MySQL AB. MySQL, 2011. <http://www.mysql.com/>.
 154. Najafi, B., Aminian, K., Loew, F., Blanc, Y. and Robert, P.A. Measurement of stand-sit and sit-stand transitions using a miniature gyroscope and its application in fall risk evaluation in the elderly. *Biomedical Engineering, IEEE Transactions on, 49* (8). 843-851, 2002.
 155. National Heart Lung and Blood Institute. Arrhythmia, 2009. http://www.nhlbi.nih.gov/health/dci/Diseases/arr/arr_whatIs.html.
 156. National Heart Lung and Blood Institute. High Blood Pressure, 2008. http://www.nhlbi.nih.gov/health/dci/Diseases/Hbp/HBP_WhatIs.html.

Bibliography

157. National Heart Lung and Blood Institute. Your Heart's Electrical System, 2008. http://www.nhlbi.nih.gov/health/dci/Diseases/hhw/hhw_electrical.html.
158. Nimmala, V.S.R., Penders, J., van Hyfte, D., Brands, M. and Gyselinckx, B. Semantic Wireless Body Area Networks. In *Engineering in Medicine and Biology Society, 2008. EMBS 2008. 30th Annual International Conference of the IEEE*, (2008), 3409-3412.
159. Nordic Semiconductors. Healthcare applications, 2011. <http://www.nordicsemi.com/eng/Applications/Healthcare>.
160. OMNet++. OMNet++ Discrete Event Simulation System, 2009. <http://www.omnet.org>.
161. Otto, C., Milenkovic, A., Sandres, O. and Jovanov, E. System Architecture of a Wireless Body Area Sensor Network for Ubiquitous Health Monitoring. *Journal of Mobile Multimedia* 1(4). 20, 2006.
162. Oum, J.H., Lee, S.E., Kim, D.-W. and Hong, S. Non-contact heartbeat and respiration detector using capacitive sensor with Colpitts oscillator. *Electronics Letters*, 44 (2). 2, 2008.
163. Pan, J. and Tompkins, W.J. A Real-Time QRS Detection Algorithm. *IEEE Transactions on Biomedical Engineering, BME-32* (3). 230-236, 1985.
164. Pantelopoulos, A. and Bourbakis, N.G. A Survey on Wearable Sensor-Based Systems for Health Monitoring and Prognosis. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 40 (1). 1-12, 2010.
165. Paradiso, R., Gemignani, A., Scilingo, E.P. and De Rossi, D. Knitted bioclothes for cardiopulmonary monitoring. In *Engineering in Medicine and Biology Society, 2003. Proceedings of the 25th Annual International Conference of the IEEE*, (2003), 3720-3723 Vol.3724.

-
166. Park, K.S., Cho, B.H., Lee, D.H., Song, S.H., Lee, J.S., Chee, Y.J., Kim, I.Y. and Kim, S.I. Hierarchical support vector machine based heartbeat classification using higher order statistics and hermite basis function. In *Computers in Cardiology, 2008*, (2008), 229-232.
167. Penders, J., van de Molengraft, J., Brown, L., Grundlehner, B., Gyselinckx, B. and Van Hoof, C. Potential and challenges of body area networks for personal health. In *Engineering in Medicine and Biology Society, 2009. EMBC 2009. Annual International Conference of the IEEE*, (2009), 6569-6572.
168. Pereira, D.C.R. Monitorização Remota de Pacientes *Information Systems Department University of Minho, Guimarães*, 2009.
169. Perry, A.G. and Potter, P.A. *Clinical Nursing Skills and Techniques*. Elsevier, St Louis, Missouri, 2010.
170. Philips Healthcare. Philips Intellivue Telemetry System Brochure, Philips Medical, 2010.
[http://incenter.medical.philips.com/doclib/enc/fetch/2000/4504/577242/577243/577247/582646/583147/IntelliVue_Telemetry_System_brochure_\(Non-US\).pdf%3fnodeid%3d585849%26vernum%3d5](http://incenter.medical.philips.com/doclib/enc/fetch/2000/4504/577242/577243/577247/582646/583147/IntelliVue_Telemetry_System_brochure_(Non-US).pdf%3fnodeid%3d585849%26vernum%3d5).
171. Philips Healthcare. Philips Lifeline, 2011.
<http://www.lifelinesys.com/content/home>.
172. Pister, K.S.J. and Doherty, L. TSMP: Time Synchronized Mesh Protocol *IASTED International Symposium on Distributed Sensor Networks IASTED*, Orlando, Florida, USA, 2008.
173. Qian, H., Mao, Y., Xiang, W. and Wang, Z. Recognition of human activities using SVM multi-class classifier. *Pattern Recognition Letters*, 31. 100-111, 2010.
174. Rangayyan, R.M. *Biomedical Signal Analysis - A Case Study Approach*. Wiley Interscience, Piscataway, NJ, 2002.

Bibliography

175. Rao, C.R. and Guha, S.K. *Principles of Medical Electronics and Biomedical Instrumentation*. Universities Press (India) Limited, Hyderabad, India, 2001.
176. Rautaharju, P.M., Surawicz, B. and Gettes, L.S. AHA/ACCF/HRS Recommendations for the Standardization and Interpretation of the Electrocardiogram Part IV: The ST Segment, T and U Waves, and the QT Interval. *Journal of the American College of Cardiology* 53 (11). 982-991, 2009.
177. Rienzo, M.D., Rizzo, F., Parati, G., Brambilla, G., Ferratini, M. and Castiglioni, P. MagIC System: a New Textile-Based Wearable Device for Biological Signal Monitoring. Applicability in Daily Life and Clinical Setting. In *Engineering in Medicine and Biology Society, 2005. IEEE-EMBS 2005. 27th Annual International Conference of the*, (2005), 7167-7169.
178. Ross, P.E. Managing care through the air [remote health monitoring]. *Spectrum, IEEE*, 41 (12). 26-31, 2004.
179. Ruzzelli, A.G., Jurdak, R., O'Hare, G.M.P. and Stok, P.V.D. Energy-Efficient Multi-hop Medical Sensor Networking *HealthNet'07*, San Juan, Puerto Rico, USA, 2007, 37-42.
180. Scalise, L., Ercoli, I. and Marchionni, P. Optical method for measurement of respiration rate *Medical Measurements and Applications Proceedings (MeMeA), 2010 IEEE International Workshop on* Ottawa, ON 2010.
181. Scheermesser, M., Kosow, H., Rashid, A. and Holtmann, C. User acceptance of pervasive computing in healthcare: Main findings of two case studies. In *Pervasive Computing Technologies for Healthcare, 2008. PervasiveHealth 2008. Second International Conference on*, (2008), 205-213.
182. Schiller, J.H. *Mobile Communications*. Pearson Education Limited, London, 2003.
183. Semmlow, J.L. *Biosignal and Medical Image Processing* Marcel Dekker, Inc., New York, 2004.

-
184. Shaltis, P., Reisner, A. and Asada, H. Calibration of the Photoplethysmogram to Arterial Blood Pressure: Capabilities and Limitations for Continuous Pressure Monitoring. In *Engineering in Medicine and Biology Society, 2005. IEEE-EMBS 2005. 27th Annual International Conference of the*, (2005), 3970-3973.
 185. Shaltis, P.A., Reisner, A.T. and Assada, H.H. Cuffless Blood Pressure Monitoring Using Hydrostatic Pressure Changes. *IEEE Transactions on Biomedical Engineering*, 55 (6). 1775-1777, 2008.
 186. Singer, E. A wristwatch that monitors blood pressure *Technology Review*, MIT, Massachusetts, Boston, 2011, 2. <http://www.technologyreview.com/biomedicine/37914/>.
 187. Solid Works Solid Works CAD. 2011.
 188. Soomro, A. and Cavalcanti, D. Opportunities and Challenges in Using WPAN and WLAN Technologies in Medical Environments *IEEE Communications Magazine* IEEE, 2006.
 189. Stallings, W. *Cryptography and Network Security: Principles and Practice* Prentice Hall, 2010.
 190. Steele, R., Secombe, C. and Brookes, W. Using Wireless Sensor Networks for Aged Care: The Patient's Perspective. In *Pervasive Health Conference and Workshops, 2006*, (2006), 1-10.
 191. Steinhart, J.S. and Hart, S.R. Calibration curves for thermistors. *Deep Sea Research*, 15. 497-503, 1968.
 192. Suriyachai, P., Roedig, U. and Scott, A. A Survey of MAC Protocols for Mission-Critical Applications in Wireless Sensor Networks. *Communications Surveys & Tutorials, IEEE PP* (99). 25, 2011.
 193. Tanenbaum, A.S. and Wetherall, D.J. *Computer Networks*. Prentice Hall, New York, 2010.

Bibliography

194. Teng, Q.Z.Y.S., Basran, D.A.D.S.-B.K.L.C., Bello-Haas, J.D. and Choi, V.Y. Pulse transit time-based blood pressure estimation using hilbert-huang transform *Engineering in Medicine and Biology Society, 2009. EMBC 2009. Annual International Conference of the IEEE*, Minneapolis, MN 2009, 1785 - 1788
195. Texas Instruments. Bluetooth Low Energy, 2011. <http://www.ti.com/ww/en/analog/bluetooth/index.htm?DCMP=BluetoothLowEnergy&HQS=NotApplicable+OT+bluetoothlowenergy>.
196. Texas Instruments. OPA336 <http://focus.ti.com/general/docs/lit/getliterature.tsp?genericPartNumber=opa336&fileType=pdf>, 2000.
197. Texas Instruments. REF200 - Dual current source/current sink, 1993. <http://focus.ti.com/general/docs/lit/getliterature.tsp?genericPartNumber=ref200&fileType=pdf>.
198. Texas Instruments. Wireless Connectivity - 1Q2011, 2011. <http://focus.ti.com/lit/sg/slab056/slab056.pdf>.
199. Thakor, N.V. Biopotentials and Electrophysiology Measurement In Webster, J.G. ed. *The Measurement, Instrumentation and Sensors Handbook*, CRC Press, Boca Raton, FL, 1999.
200. The Apache Software Foundation. Apache HTTP Server Project, 2011. <http://httpd.apache.org/>.
201. Tia, G., Greenspan, D., Welsh, M., Juang, R.R. and Alm, A. Vital Signs Monitoring and Patient Tracking Over a Wireless Network *Engineering in Medicine and Biology Society, 2005. IEEE-EMBS 2005. 27th Annual International Conference of the*, 2005, 102-105.
202. TinyOS Technology Exchange. Hardware Designs, 2009. <http://www.tinyos.net/scoop/special/hardware>.

-
203. Tompkins, W.J. (ed.), *Biomedical Digital Signal Processing - C Language Examples and Laboratory Experiments for the IBM PC*. Prentice Hall, Englewood Cliffs, New Jersey, 1993.
204. UC Berkeley Robotics and Intelligent Machines Lab Home Page. Motes And Radios, 2009.
<http://robotics.eecs.berkeley.edu/mast/wiki/pmwiki.php?n=Main.MotesAndRadios>.
205. Urrusti, J.L. and Tompkins, W.J. Performance evaluation of an ECG QRS complex detection algorithm *Proc. Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 1993, 800-801.
206. Varshney, U. and Sneha, S. Patient monitoring using ad hoc wireless networks: reliability and power management. *Communications Magazine, IEEE*, 44 (4). 49-55, 2006.
207. Vasu, V., Fox, N., Heneghan, C. and Sezer, S. Using the Lomb Periodogram for Non-Contact Estimation of Respiration Rates *32nd Annual International Conference of the IEEE EMBS*, Buenos Aires, 2010, 4.
208. Webster, J.G. *Medical Instrumentation Application and Design*. Wiley, Hoboken, NJ 2009.
209. Wei, J. and Seong Kong, G. Block-Based Neural Networks for Personalized ECG Signal Classification. *Neural Networks, IEEE Transactions on*, 18 (6). 1750-1761, 2007.
210. Welch Allyn. FlexNet for 802.11a life-critical wireless networks.
www.welchallyn.com/products/en-us/x-16-vo-96-1234190869014.htm
211. Welch Allyn. Micropaq Patient Wearable Monitor, 2005.


Bibliography

212. Welch Allyn. Normal Body Temperature Ranges, 2003. http://www.welchallyn.com/documents/Thermometry/Electronic%20Thermometry/SureTemp%20690/normalbodytempchart_20070323_suretempplus.pdf.
213. Welch Allyn. Welch Allyn 1500 Patient Monitor 2011. <http://www.welchallyn.com/products/en-us/x-19-pe-96-1290086676711.htm>.
214. WellDoc. Products and Services, 2010. <http://www.welldocinc.com/Products-and-Services.aspx>.
215. Wellkang tech Consulting. CE Marking, 2011. <http://www.ce-marking.org/index.html>.
216. Wheeler, A. Commercial Applications of Wireless Sensor Networks Using ZigBee. *Communications Magazine, IEEE*, 45 (4). 70-77, 2007.
217. Wong, M.Y.M., Pickwell-MacPherson, E. and Zhang, Y.T. Impedance cardiography for cuffless and non-invasive measurement of systolic blood pressure. In *Engineering in Medicine and Biology Society, 2009. EMBC 2009. Annual International Conference of the IEEE*, (2009), 800-802.
218. YSI Inc. Probes and Accessories Catalog, 2001.
219. ZigBee Alliance. The ZigBee Alliance and IPSO Alliance Collaborate in ZigBee IP and ZigBee Smart Energy 2.0, 2010. http://zigbee.org/imwp/idms/popups/pop_download.asp?contentID=17822.
220. ZigBee Alliance. ZigBee Alliance Document 053474r17, ZigBee Specification, v. 1.0 r17. Alliance, Z. ed., 2007.
221. ZigBee Alliance. ZigBee Alliance: Wireless Control that Simply Works, 2011. <http://www.zigbee.org>.
222. ZigBee Alliance. ZigBee Health Care Profile Specification, Revision 15, Version 1.0, San Ramon, CA, USA, 2010.


223. ZigBee Alliance. ZigBee Home Automation Public Application Profile - Revision 26, Version 1.1 *ZigBee Document 053520r26*, 2010.

224. ZigBee Alliance. ZigBee Specification - ZigBee Document 053474r06, Version 1.0, 2004.

Appendix A Schematic diagram of the electrocardiogram sensor board



Appendix B Schematic diagram of the temperature sensor board



Appendix C Temperature sensor battery lifetime calculations

The temperature sensor battery lifetime can be estimated using the current consumptions and the time periods listed in Table C1. For each temperature measurement, it is assumed that the temperature is sampled only once.

Table C1 – Current consumption per module or component.

Activity	Current consumption (μA)	ΔT (ms)
ADC on for temperature sampling	580	1
Front-end consumption (except wireless module)	2.4	1
ADC on for battery level sampling	580	1
LED on	3000	500
Radio on (TX or RX)	39000	2.27
CPU processing	9200	504.27
CPU sleeping with wakeup timer on	3.3	59495.73

As shown in Table C2, for each transmitted message, the transceiver is switched on for a minimum period of 2.72 ms.

Table C2 – Minimum period of time the transceiver is switched on for each transmitted message.

Action	Activity	ΔT (ms)
ED senses the channel clear	CCA	0.128
ED's transceiver changes to TX mode	Turnaround RX-to-TX	0.192
ED transmits a data packet to the coordinator (44 bytes)	TX	1.408
ED's transceiver changes to RX mode	Turnaround TX-to-RX	0.192
ED receives the acknowledgment (11 bytes)	RX	0.352
Transceiver ON time (ms):		2.272

All periods of time listed in Table C1 repeat every sixty seconds, as shown in Figure C1.

Appendix C. Temperature sensor battery lifetime calculations





Figure C1 – A temperature sensor cycle.

From Table C1, it is possible to get an average current consumption of 0.107 mA. For a 620 mAh battery, this consumption corresponds to a battery lifetime of, approximately, 241 days

Appendix D Schematic diagram of the coordinators and routers



Appendix E Questionnaires

ECG sensor evaluation - Patient questionnaire

Caro(a) Sr.(a),

Gostaríamos de saber a sua opinião sobre o sistema que temos desenvolvido na Universidade do Minho. Ele permite monitorizar sinais vitais, tais como o ECG e a temperatura corporal, através do uso de tecnologias de comunicação sem fios.

Este inquérito é anónimo; apenas solicitamos que nos indique a sua idade. As suas impressões e sugestões serão de enorme utilidade para a melhoria do sistema.

Desde já agradecemos a sua colaboração neste projecto académico.

A equipa de projecto.

Instituto de Polímeros e Compósitos
Universidade do Minho

Idade do(a) paciente: _____

Data de realização do inquérito: ___/___/___

Inquérito

	1 Discordo totalmente	2	3	4	5 Concordo totalmente
É portátil					
É pequeno					
É leve					
Não é obstrutivo					
É confortável					

O sensor caiu durante a sua utilização?

Sim Nao

Considera que o sensor de ECG pode contribuir de forma positiva para os cuidados prestados pelo hospital?

Sim Nao Não consigo avaliar

Caso tenha respondido sim à questão anterior, de que forma?

Teve algum receio quanto a utilização do sensor? Qual?

Criticas e sugestões de melhoria:

ECG sensor evaluation – Health care provider questionnaire

Inquérito Elaborado para os Enfermeiros que Utilizaram o Protótipo Final

	1 Discordo totalmente	2	3	4	5 Concordo totalmente
Sensor de ECG					
E fácil ligar e desligar o sensor					
E fácil trocar a bateria					
E fácil desinfetar o sensor					
E fácil colocar e retirar o sensor em pacientes					
E evidente a orientação da colocação do sensor					

Em algum momento a interface com o utilizador deixou de operar? Caso tenha ocorrido, descreva, sucintamente, o que sucedeu.

Algum dos sensores apresentou alguma avaria ou comportamento inesperado?

Críticas e sugestões de melhoria:

Temperature sensor evaluation - Patient questionnaire

Caro(a) Sr.(a),

Gostaríamos de saber a sua opinião sobre o sistema que temos desenvolvido na Universidade do Minho. Ele permite monitorizar sinais vitais, tais como o ECG e a temperatura corporal, através do uso de tecnologias de comunicação sem fios.

Este inquérito é anónimo; apenas solicitamos que nos indique a sua idade. As suas impressões e sugestões serão de enorme utilidade para a melhoria do sistema.

Desde já agradecemos a sua colaboração neste projecto académico.

A equipa de projecto.

Instituto de Polimeros e Compósitos
Universidade do Minho

Idade do(a) paciente: _____
Data de realização do inquérito: ____/____/____

Inquérito

	1 Discordo totalmente	2	3	4	5 Concordo totalmente
É portátil					
É pequeno					
É leve					
Não é obstrutivo					
É confortável					

O sensor caiu durante a sua utilização?

Sim Nao

Considera que o sensor de temperatura pode contribuir de forma positiva para os cuidados prestados pelo hospital?

Sim Nao Não consigo avaliar

Caso tenha respondido sim à questão anterior, de que forma?

Teve algum receio quanto à utilização do sensor? Qual?

Críticas e sugestões de melhoria:

Temperature sensor evaluation – Health care provider questionnaire

Inquérito Elaborado para os Enfermeiros que Utilizaram o Protótipo Final

	1 Discordo totalmente	2	3	4	5 Concordo totalmente
Sensor de Temperatura					
E fácil ligar e desligar o sensor					
E fácil trocar a bateria					
E fácil desinfetar o sensor					
E fácil colocar e retirar o sensor em pacientes					
E evidente a orientação da colocação do sensor					

Em algum momento a interface com o utilizador deixou de operar? Caso tenha ocorrido, descreva, sucintamente, o que sucedeu.

Algum dos sensores apresentou alguma avaria ou comportamento inesperado?

Críticas e sugestões de melhoria:

System evaluation – Health care provider questionnaire (page 1/2)



Universidade do Minho

Questionário sobre o funcionamento do sistema HM4AllInformações sobre o entrevistado:

Profissão: _____

Local de trabalho habitual (ex: cuidados intensivos em ambiente hospitalar):
_____Informações sobre o sistema e perguntas:

O sistema atual inclui sensores de electrocardiograma (ECG) e temperatura. O sistema permite observar o electrocardiograma, medir o ritmo cardíaco e a temperatura na região da axila. Além disso, permite a introdução de valores limites que quando ultrapassados geram um alarme. Também é possível monitorizar o nível de bateria dos sensores.




Figure 1. Componentes do sistema: (a) sensor de ECG; (b) sensor de temperatura; (c) coordenador/router; (d) janela de monitorização de paciente.

Avaliação da aplicabilidade do sistema:

Totalmente
inútil ou sem
interesseTotalmente
útil ou
interessant
e

	1	2	3	4	5
Quão útil é o sistema?					
Considera útil a extensão da monitorização de sinais vitais a pacientes fora dos cuidados intensivos e unidades intermediárias de cardiologia?					
Considera que um sistema capaz de apresentar apenas sinais vitais e gerar alarmes pode trazer benefícios a pacientes que não requerem cuidados urgentes ou cuja situação clínica ainda não está definida?					



Universidade do Minho

Avaliação do sensor de ECG:

	Discordo totalmente				Concordo totalmente
	1	2	3	4	5
E fácil ligar e desligar o sensor					
E fácil trocar a bateria					
E fácil desinfetar o sensor					
E fácil colocar e retirar o sensor em pacientes					
E evidente a orientação do sensor					

Avaliação do sensor de temperatura:

	Discordo totalmente				Concordo totalmente
	1	2	3	4	5
E fácil ligar e desligar o sensor					
E fácil trocar a bateria					
E fácil desinfetar o sensor					
E fácil colocar e retirar o sensor em pacientes					

Avaliação global:

1) Qual a função mais interessante do sistema? Por que?

2) Que função gostaria de ver incorporada ao sistema e como essa nova função traria benefícios ao seu trabalho?

3) Críticas e sugestões de melhoria:

Appendix F Range tests

The tests described in this annex were executed at Hospital Privado de Guimarães, on one of the impatient areas. They aimed on a) estimating the range of the wireless links; b) defining adequate spots to position routers; and c) quickly verifying the DR in several scenarios.

F.1 Tests configuration

Nonbeacon-enabled ZigBee networks operating in star and 2-hop tree topologies with a variable number of EDs were used. The end devices (EDs) were programmed to generate one 108-byte message every 500 ms, the same message size and period used by electrocardiogram (ECG) sensors that transmit compressed data. Up to six EDs were used because this is the maximum number of patients to be simultaneously monitored. Each message included a sequence number added by the Application layer. All messages used hop-by-hop acknowledgements. Sensors could make up to four attempts to access the channel and up to three retries were allowed. Each test involved the successful reception of at least 1,000 messages. All tests were conducted under the supervision of the hospital staff. Some rooms were occupied with patients who voluntarily granted the research team access to their rooms.

The boards used as coordinator and routers are based on JN5139-M02 high power modules, whereas EDs' boards are based on JN5139-M00 regular modules. The receiver sensitivity of high power modules and regular modules is -100 dBm and -96.5 dBm, respectively. The coordinator and routers were programmed to transmit at +10 dBm, whereas ED boards were programmed to transmit at 0 dBm.

F.2 Test results

Figure 113 shows the settings used on the first tests. The coordinator was positioned on the entrance. It was not possible to go into room R213, so two devices were positioned at room R211 and one at room R212. Blue curved vectors represents links used to send messages.




Figure 113 – DR test 1: coordinator at the entrance hall.

The tests results are shown in Table 42 where each line contains the test results for each setting, (a) and (b). The column *Average* presents the average delivery ratio (DR) considering the three EDs involved in the test. Column labels R211 (A), R211 (B) and R212 refer to the spots at which EDs were placed. These columns present the DR calculated from the messages generated by each device. All other tables in this section have similar structure.

The first test was performed using the setting shown in Figure 113 (a). During this test, it was observed that one of the devices on room R211 lost some messages when one of the nurses to handle the device, simulating its placement in a patient. It indicates that the quality of the link between the device and the coordinator was weak. In fact, observing the survey results shown in Figure 93 (a), it is possible to observe that the average link quality indication (LQI) associated to this link is 48, not a high value, despite being above the limit recommended by the manufacturer ($LQI > 30$). In consequence of that, one router was added at spot 1 (see Figure 93), as shown in Figure 113 (b). All devices associated to the router and, despite the additional hop, the DR increased slightly.

Table 42 – DR test 1 results.

Test	DR (%)			
	Average	R211(A)	R211(B)	R212
<i>a</i>	99.7	99.1	99.9	100
<i>b</i>	99.8	99.5	99.8	100

Then, the coordinator was moved to the reception, as shown in Figure 114. The tests results shown in Table 43 indicate that, if all devices associate to the router, the average DR measured is equivalent to the DR measured for the setting shown in Figure 113 (b). However, when one of the devices associated directly to the coordinator, as shown in Figure 114 (b), only 74.3% of the generated messages were correctly received, which demonstrated that the quality of the direct link between the device and the coordinator was much worse than the link between the device and the router (notice that the value of the LQI associated to this link is equal to 44, as shown in Figure 93 (b)).




Figure 114 – DR test 2: coordinator at the reception.

Table 43 – DR test 2 results.

Test	DR (%)			
	Average	R211(A)	R211(B)	R212
<i>a</i>	99.9	100	99.8	100
<i>b</i>	91.4	100	74.3	100

Parent selection is done using the results gathered during the network discovery procedure. This procedure returns to the APP layer a network list detailing the personal

area network (PAN) identifications within the device’s range. This procedure may be repeated a certain number of times defined by a Network (NWK) layer parameter. The purpose of repeating the network discovering procedure is to provide a more accurate neighbor list and associated link quality indications. The information gathered includes the PAN ID, the operational mode of the network, the identity of the ZigBee router or coordinator identified on the PAN, the depth of the ZigBee router on the PAN from the ZigBee Coordinator for the PAN, the capacity of the ZigBee router or coordinator and the routing cost [224].

In the configuration depicted on Figure 114 (b), one of the EDs chooses a poor direct link to the coordinator over a better link to the router. During the test, EDs were programmed to automatically join a specific network. Possibly, the ED failed receiving the beacon transmitted by the router during the discovery procedure. Otherwise, the beacon sent by the router could have been atypically attenuated. Some stack implementations automatically repeats the discovery procedure. It allows the joining device to obtain more accurate information about potential parents. Jennic’s stack implementation allows programmers to manually join a network and, consequently, to repeat the discovery process as many times as one wishes.

Two tests were performed after moving the coordinator to the reception and placing the router and the EDs as shown in Figure 115. All EDs have associated to the router and, for both tests, the resulting DR was good, as shown in Table 44.




Figure 115 – DR test 3: coordinator at the reception.

Table 44 – DR test 3 results.

Test	DR (%)			
	Average	R209	R208(A)	R208(B)
<i>a</i>	100	100	99.9	100
<i>b</i>	99.8	99.9	99.7	99.8

Two additional tests were performed after moving the coordinator to the pharmacy. The new setting and the associations between devices are shown in Figure 116.




Figure 116 – DR test 4: coordinator at the pharmacy.

As presented in Table 45, both results are worse than the ones obtained using the coordinator in the reception (shown in Table 44), which suggests that the quality of the link between the coordinator and the router deteriorated when the coordinator was moved to the pharmacy. It was subsequently confirmed, and may be explained by the existence of fully-loaded wall-mounted medication cabinets on the pharmacy, which can attenuate radiofrequency signals. The atypically bad results obtained on test (b) occurred because, frequently, one of two EDs was transmitting a message just after the other one has received an acknowledgement frame from the router.

Table 45 – DR test 4 results.

Test	DR (%)			
	Average	R209	R208(A)	R208(B)
<i>a</i>	97.6	98.1	98.0	96.6
<i>b</i>	88.1	84.0	98.3	82.1

A few more tests were executed on the right side of the in-patient area. The coordinator was kept on the consultation room while a variable number of end devices were distributed by the rooms. The first test was performed using the setting shown on Figure 117. As shown, a router was positioned on spot 6. Both EDs joined the router and the DR was equal to 100%.




Figure 117 – DR test 5: coordinator at the consultation room.

EDs were moved to room R206 and, after turned on, associated directly to the coordinator, as shown in Figure 118 (a). Due to the low link quality, several messages were lost, as shown in Table 46. A router was added to spot 5, as shown in Figure 118 (b). The test was repeated and, that time, all generated messages were correctly received.




Figure 118 – DR test 6: coordinator at the consultation room.

Table 46 - DR test 6 results.

Test	DR (%)		
	Average	R206(A)	R206(B)
<i>a</i>	95.6	94.2	97.0
<i>b</i>	100	100	100

Finally, two additional tests with six EDs were performed using the settings shown in Figure 119. The results are presented in Table 47 where EDs are designated ED1 to ED6 despite the spot at which they were placed.




Figure 119 – DR test 7: coordinator at the consultation room.

Table 47 – DR test 7 results.

Test	DR (%)						
	Average	ED1	ED2	ED3	ED4	ED5	ED6
<i>a</i>	100	100	100	100	100	99.9	100
<i>b</i>	94.7	99.7	99.8	99.8	99.8	86.1	82.9

The average DR measured for the first test was approximately 100%. However, during the second test, several messages from two devices located on room R206 were lost (ED5 and ED6). The captured files analysis reveals that a message from one ED quickly follows a message just transmitted by the other one, which indicates that the period between messages generation is very small, which blocked the router.

F.3 Conclusions

Considering the specific environment in which tests were executed and the tests results, good DR results can be achieved if the PAN coordinator and routers are positioned in the spots defined in Section 6.2. However, the results presented in Table 47 for test (b), indicate that, despite the CSMA-CA mechanism used by the MAC layer,

when two EDs generate messages within a small interval, a reasonable percentage of messages can be lost by both devices.