**Ricardo Dias
Fernandes**

Projecto de um Sensor sem Fios Passivo

**Ricardo Dias
Fernandes**

Projecto de um Sensor sem Fios Passivo

Design of a Battery-free Wireless Sensor Node

Dissertação apresentada à Universidade de Aveiro para cumprimento dos requisitos necessários à obtenção do grau de Mestre em Engenharia Electrónica e Telecomunicações, realizada sob a orientação científica do Doutor Nuno Miguel Gonçalves Borges de Carvalho, Professor Associado com Agregação do Departamento de Electrónica, Telecomunicações e Informática da Universidade de Aveiro, e sob a co-orientação científica do Doutor João Nuno Pimentel da Silva Matos, Professor Associado do Departamento de Electrónica, Telecomunicações e Informática da Universidade de Aveiro.

O júri

presidente

Professor Doutor José Carlos Esteves Duarte Pedro

Professor Catedrático do Departamento de Electrónica, Telecomunicações e Informática da Universidade de Aveiro

vogais

Professor Doutor José Alberto Peixoto Machado da Silva

Professor Associado do Departamento de Engenharia Electrotécnica e de Computadores da Faculdade de Engenharia da Universidade do Porto

Professor Doutor Nuno Miguel Gonçalves Borges de Carvalho

Professor Associado com Agregação do Departamento de Electrónica, Telecomunicações e Informática da Universidade de Aveiro

Professor Doutor João Nuno Pimentel da Silva Matos

Professor Associado do Departamento de Electrónica, Telecomunicações e Informática da Universidade de Aveiro

Agradecimentos

É verdade que o meu nome é o único escrito na capa deste documento, mas isso não significa que eu tenha sido eu o único autor, e nem tão pouco que a lista de pessoas que contribuíram de uma forma ou de outra tenha sido propriamente pequena, e por isso,
ao meu orientador, prof. Nuno Borges Carvalho,
ao meu co-orientador, prof. João Nuno Matos,
aos meus colegas de laboratório, de quem destaco o Luís Nunes, o Diogo Dias, o Daniel Lourenço, o Diogo Cunha e a Vanessa Silva,
aos meus colegas de residência, em particular ao Luís Rodrigues e ao Luís Pereira,
aos meus padrinhos,
aos meus avós, que infelizmente já partiram,
à minha não avó mas quase, Carminda Nogueira,
em especial, aos meus pais,
a todas as outras pessoas que influenciaram a minha vida,
o meu mais sincero obrigado.

Palavras-chave

captação de energia RF; captação de energia electromagnética; redes de sensores sem fios; RFID; modulação por retro-espalhamento de ondas electromagnéticas; redes RFID com sensores; redes de sensores passivos sem fios.

Resumo

Redes de sensores sem fios são hoje em dia utilizadas numa grande variedade de aplicações, o que justifica o facto de que os sensores que as constituem sejam igualmente diversificados. Ainda assim, quase todos eles dependem de baterias, as quais ficam sem carga normalmente muito antes do fim de vida dos restantes componentes. Para além disso, o tamanho das baterias é neste momento um impeditivo à redução do tamanho dos sensores. Uma forma de contornar estes problemas consiste em retirar as baterias dos dispositivos e em alternativa captar a energia das ondas electromagnéticas radiadas por uma fonte colocada próxima destes. Neste documento descrevem-se em detalhe um sensor sem fios projectado para captar energia de ondas rádio a 866.6MHz e a respectiva antena. Além de ser passivo, o sistema proposto é também programável, uma vez que o sensor integra um microcontrolador de uso geral, e inclui um conector de 50Ω e uma interface para depuração e expansão, composta por um total de 26 pinos. Em termos de performance prática, o sistema proposto é capaz de executar tarefas relacionadas com comunicação e processamento até um máximo de 4.1 metros de distância de uma antena transmissora a operar dentro das limitações impostas pelas entidades reguladoras locais, no que diz respeito a potência.

Keywords

RF energy harvesting; electromagnetic energy harvesting; wireless sensor networks; RFID; backscatter modulation; RFID sensor networks; wireless passive sensor networks.

Abstract

Wireless sensor networks are currently of primary importance in a multitude of applications, and therefore, it comes as no surprise that there are many types of sensor nodes as well. Yet, almost all of them operate on batteries that normally deplete long before the predicted life span of basically all the other hardware components. Not only that, the large size of the batteries is indeed actually preventing sensor nodes from becoming smaller. One way of overcoming the drawbacks related to batteries is to remove them and harvest all the necessary energy from electromagnetic waves being radiated by a nearby source. In this document, a wireless sensor node designed to harvest energy from radio waves at 866.6MHz and its antenna are proposed and described in detail. In addition to being passive, the proposed system is also programmable, given that the sensor node includes a general-purpose microcontroller, and features a 50Ω port, and an interface for debugging and expansion, comprised of a total of 26 pins. Lastly, with regards to practical performance, the proposed system is able to carry out communication and processing tasks at up to a distance of 4.1 meters away from a transmitter antenna radiating within the limits imposed by local regulatory entities, with respect to power.

Contents

Contents	i
List of Figures	iii
List of Tables	v
1 Introduction and objectives	1
1.1 Introduction	1
1.2 Objectives	3
1.3 An overview of the relevant concepts and technologies	4
1.3.1 Radio Frequency IDentification (RFID)	4
1.3.2 Wireless Sensor Networks (WSNs)	9
1.3.3 Antennas and radio wave propagation	14
1.4 Unusual applications involving WSNs and RFID	17
2 State of the art	19
2.1 Introduction	19
2.2 Wireless Passive Sensor Networks (WPSNs)	20
2.3 RFID Sensor Networks (RSNs)	20
2.3.1 Intel Wireless Identification and Sensing Platform (WISP)	21
2.4 Antennas for electromagnetic energy harvesting	22
3 System architecture	25
3.1 Introduction	25
3.2 Description of the tools used	28
3.2.1 Easily Applicable Graphical Layout Editor (EAGLE)	29
3.2.2 Advanced Design System (ADS)	29
3.2.3 High Frequency Structure Simulator (HFSS)	30
3.2.4 Code Composer Studio (CCS)	30
3.3 System overview	31
3.3.1 Energy harvesting and communication mechanisms	32
3.3.2 Supply voltage supervisor	34
3.3.3 Non-volatile memory	35
3.3.4 Other peripheral components	36

4	System optimization	39
4.1	Introduction	39
4.2	Radio front-end performance and design factors	39
4.2.1	Matching network and voltage multiplier (envelope detector)	40
4.2.2	Proposed antenna	43
4.3	Software algorithms and power management	44
4.3.1	Deep sleep mode	44
4.3.2	Proposed program for communication and control	44
5	Experimental results	47
5.1	Introduction	47
5.2	Harvested voltage, range and power consumption	47
5.2.1	Harvested voltage as a function of input power	47
5.2.2	Frequency and power sweeps	48
5.2.3	Operating distance in practice	49
5.3	Antenna measurements and efficiency considerations	52
5.3.1	Reflection coefficient	52
5.3.2	Gain and radiation diagram	53
6	Conclusions and future research	55
6.1	Introduction	55
6.2	Main conclusions	55
6.2.1	Further work	56
6.2.2	Concluding remarks on electromagnetic energy harvesting	57
A	Source code	59
A.1	Deep sleep mode	59
A.2	Proposed program for communication and control	60
	Bibliography	65

List of Figures

1.1	Star and mesh network topologies, represented on the left and right sides of the figure, respectively (in both network configurations, the empty circles represent sink nodes whereas the filled circles correspond to ordinary sensor nodes).	13
2.1	The result of bending an arm of a dipole antenna several times, following perhaps the most common of many techniques.	22
3.1	PCB layout of the developed sensor node shown on a 2:1 scale (darker-colored lines refer to the top copper layer whereas lighter areas and lines refer to the bottom layer, the circles represent vias and the remaining drawings refer to the component outlines).	27
3.2	Proposed dipole-based PCB antenna designed for use with the sensor node, shown on a 1:1 scale (the rectangular area represents the substrate and the other two structures correspond to the arms of the dipole, with the upper arm being on the top side of the substrate and the lower arm being on the bottom side).	28
3.3	Block diagram of the proposed sensor node (in which the tips of the arrows indicate in which directions power or data, or both, are allowed to flow).	31
3.4	Proposed circuitry for energy harvesting and communications (individual hardware parts are highlighted with filled areas and accompanied by relevant information).	33
3.5	Proposed supply voltage supervisory circuit (once again, individual hardware parts are highlighted with filled areas and accompanied by relevant information).	35
3.6	Non-volatile memory circuitry (as with previous figures, individual hardware parts are highlighted with filled areas and accompanied by relevant information).	36
3.7	Other hardware components integrated into the sensor node, which, from left to right correspond to a crystal oscillator and a to a single resistor.	36
3.8	Two pin interfaces where a number of signals from within the developed sensor node are routed to, having a combined total of 26 pins.	37
4.1	The model developed for simulating the RF front-end of the sensor node, in which microstrip line lengths are shown (in millimeters) and parasitic resistances are either added in the form of resistors, or in the case of the inductor, included directly into the component to which they refer to. Different line widths were utilized in the model, corresponding to 0.7, 0.8, 0.9 and 1.77mm, as well as circular vias with diameters of 1.5mm.	42

5.1	Harvested voltage seen at the storage capacitor within the sensor node in function of increasing and decreasing power levels injected at its input (plotted with filled and empty circles respectively), in comparison to the predicted results (dotted line with borderless circles), and always considering a sinusoid at 866.6MHz. Solid lines were plotted to highlight points located at -11.44 and -8.64dBm.	48
5.2	Reflection coefficient measured at the input of the sensor node (represented by filled circles) in function of frequency, plotted against the predicted magnitude and phase curves (dotted lines with borderless circles) and assuming a test power of -8.64dBm. Solid lines were utilized to mark points at 866.6MHz.	50
5.3	Reflection coefficient measured at the input of the sensor node (represented by filled circles) in function of power, plotted against the predicted magnitude and phase curves (dotted lines with borderless circles) and assuming a test frequency of 866.6MHz. Solid lines were used to mark points at -8.64dBm.	51
5.4	Reflection coefficients of the first and second antennas (plotted with empty and filled circles respectively) in function of frequency, compared to the curve (dotted line with borderless circles) expected for the first antenna, magnitude only. The highlighted data value corresponds to a frequency of 866.6MHz.	52
5.5	Normalized radiation patterns of the second antenna, measured at the planes that contain the direction of maximum radiation plus either the electric or magnetic field vectors (upper and lower plots respectively, both indicated by filled circles), plotted against expected curves (dotted lines with borderless circles), considering a frequency of 866.6MHz.	54

List of Tables

1.1	European frequency and power regulations for both UHF and microwave RFID (in case a certain frequency has more than one maximum output power level associated with it, the higher power limit applies).	8
4.1	A comparison between the RO4003C and FR4, in terms of dielectric constant and loss tangent characteristics at UHF.	43

Chapter 1

Introduction and objectives

At a fundamental level, the main purpose of this document is to describe the design and practical implementation of a battery-free wireless sensor node.

In this chapter, the first section introduces electromagnetic energy harvesting as a method of providing a remote power source for wireless sensor nodes, and also explains the main advantages of this approach, compared to disposable batteries. In the second section, the objectives of this work are presented. In third section, Radio Frequency IDentification (RFID) and Wireless Sensor Networks (WSNs) are introduced and discussed in general terms, and a few basic notions related to antennas and to radio wave propagation in free space are discussed from the point of view of electromagnetic energy harvesting. The fourth section describes a couple of peculiar applications in which RFIDs and WSNs play decisive roles.

The second chapter further extends the literature reviews, by discussing the state of the art in terms of how the aforementioned technologies are evolving (converging, actually), in an effort to enable wireless sensing and computation totally free of batteries, and consequently, totally free of lifetime constraints. Additionally, the same chapter also discusses the significant importance of antennas in the process of gathering energy from radio waves.

The third chapter marks the transition between research work and the actual implementation of a passive wireless sensor node. In this chapter, the major software applications involved in the design process are introduced, and a detailed hardware overview of the proposed sensor node and antenna is given. The fourth chapter continues to describe the system, but focuses essentially on subjects related to high frequency behavior and software algorithms. The fifth chapter describes a series of tests that were performed on the prototype, and finally, the sixth chapter concludes this document, by summarizing the key features of the system and indicating possible directions for further research.

1.1 Introduction

Wireless sensor nodes currently include processors designed for low-power operation that support various operating modes and with the ability to change from a sleep state to a fully active mode quickly and seamlessly. In addition to that, most sensor nodes have various power switches that can cut the power to peripheral hardware modules, if the processor finds those to be unnecessary

at a given moment of time. Of course, this delicate management of energy is only achievable if a suitable set of software routines are loaded onto the processor. These routines essentially instruct the processing unit to reduce its power consumption as much as possible while still maintaining a minimum wireless awareness, so that incoming transmissions can be received and decoded, and are usually referred to in literature as energy-aware software routines. When combined together with energy-efficient communication devices, these advancements in hardware and software design effectively reduce the power consumption requirements of sensor nodes, not only allowing them to operate autonomously for extended periods that can stretch up to a few years, but also creating a significant opportunity for the use of energy harvesting systems, because the primary challenge in WSN deployment still is the limited network lifetime due to finite capacity power sources, that is, batteries (in [1], many design issues that currently affect WSNs are identified, not only in terms of hardware but also in terms of networking protocols).

Energy harvesting enables sensors to operate from theoretically unlimited alternative forms of energy, such as thermal, biological, mechanical, photovoltaic (perhaps the most discussed these days), chemical or electromagnetic energy, thus avoiding batteries, and as a consequence, many of the drawbacks related to them. Because of the advantages it provides, the concept of collecting energy from the environment is now a highly desirable and increasingly important capability in sensor networks.

The conventional sensor network communication model assumes the deployment of low-cost and low-power, multifunctional wireless sensor nodes operating on the limited power capacity of disposable batteries. Although this model actually works decently in small WSNs composed of a few wireless nodes, it does not scale well. It is easy to understand that the potential maintenance cost of replacing batteries for a few hundreds or thousands of nodes quickly becomes a continuous and cost-prohibitive undertaking. Considering for instance no more than a couple hundred sensor nodes, each one capable of unattended operation for roughly one year before its energy source collapses, the maintenance costs immediately begin to escalate because of the need to replace batteries every few days, on average. Furthermore, it is important to realize that the mentioned lifetime of one year already assumes a low duty-cycle operation (something that is typical in these networks), as a fully active sensor node would deplete its battery in a matter of weeks, or perhaps days. Battery replacements occurring at such a fast pace can easily result in an unsustainable cost for many applications.

Maintaining battery-operated sensor nodes in hard to service locations is yet another significant challenge. In this not so unrealistic scenario, maintenance operations become much more difficult to achieve and hence more expensive, thus adding a substantial cost to the already high cost of replacing batteries. This can even limit where sensor nodes are located, potentially reducing the effectiveness of the overall deployment. Additionally and very importantly, for WSNs to achieve a true ubiquitous deployment, the size of the nodes must decrease drastically. Small batteries with suitable form factors effectively help reducing the size of a sensor node, but generally contain less energy than other (more traditional) larger batteries and thus have shorter life spans.

In short, the need shared by most sensor nodes for long lifetimes and small form factors does not match up well with the power density of available battery technology. Although exponential improvements have indeed occurred in other hardware components, such improvements have not occurred in battery technology, and no major changes are anticipated for now (the authors of [2]

consider that energy harvesting stands as a promising new approach precisely due to the fact that batteries are not progressing at the same rate that other hardware components are).

Electromagnetic energy harvesting, also commonly referred to as Radio Frequency (RF) energy harvesting, enables wireless sensors to be located in inaccessible or hazardous areas, or locations where battery replacement is highly impractical. With the ability to power multiple devices from a single power source, this technology allows WSNs to scale by supplying maintenance-free power for hundreds or thousands of wireless sensor nodes. Moreover, traditional energy harvesting sources such as solar, piezo or thermal, share a common limitation of being reliant on ambient sources that are generally beyond control. A wireless power solution based on RF energy harvesting overcomes this lack of control because power can be replenished when desired, either in a continuous fashion or using other schemes based on scheduled or on-demand recharge cycles. This mode of operation allows for a WSN to have a true zero standby power consumption.

In this document, a passive sensor node designed to operate exclusively on power harvested from Ultra High Frequency (UHF) radio waves is presented and described in detail. Nearly passive backscatter communication is used in this sensor node. It is important to notice that RF energy harvesting and backscatter radio links are essential aspects in passive RFID tags, as these devices transmit information using reflected electromagnetic waves and rely entirely on the readers as their remote power source.

1.2 Objectives

The main objective of this work was to design a wireless sensor node with energy harvesting and communication capabilities based on those found in passive RFID tags. The wireless sensor node to be developed was envisioned to be battery-independent, self-sustainable, fully programmable and tuned for maximum range, while still maintaining a small size, preferably no larger than other more conventional (battery-operated) sensor nodes. Even though the sensor was always idealized to be inexpensive, this was considered a secondary, less important objective.

Understanding how RFID readers and tags operate proved to be important in order to select suitable energy harvesting and communication mechanisms for the sensor node. Other than that, the sensor followed standard design guidelines common to most sensor nodes, and therefore, well documented and field-proven. As the harvesting performance was verified to be quite dependent of how well designed the high frequency frontend of the sensor node was, the construction of a simulation model encompassing all hardware components subject to higher frequencies was also added to the list of objectives. Tasks related to the actual practical implementation of the sensor node were then carried out and a prototype was built. At this phase, the main objectives were to implement software routines for communication and energy management and to build an antenna system. Once all objectives were considered accomplished, a series of tests were performed on the final system, so that theoretical assumptions could be compared to experimental data.

1.3 An overview of the relevant concepts and technologies

Only recently remote feeding with electromagnetic waves has started to be considered as a viable alternative to the batteries that still equip most sensor nodes, and the same applies to backscatter modulators versus conventional active transmitters, regarding communication. Even though being relatively new to sensor networks, these two technologies have already proven their worth (and in fact still prove) in RFID, in passive tags.

In the next subsection, a brief historical overview of RFID is presented first, followed by some text explaining what a RFID system is and how it works. Moving on to more advanced subjects, the two different RFID coupling mechanisms are introduced and the radio regulations that apply to this widely used technology are shown. The content related to RFID ends with an important part devoted to the backscatter modulation technique. In the subsection afterwards, some historical references related to WSNs are presented, and after that, the typical hardware architecture of a sensor node is shown. The following text explains how sensor nodes normally organize themselves to form networks and also identifies some techniques that can be (and usually are) used to reduce the power consumption in these networks. In the third and last subsection, subjects related to antennas and radio wave propagation are discussed, taking into account the specificities of WSNs and RFIDs, and especially considering RF energy harvesting.

1.3.1 Radio Frequency IDentification (RFID)

During the 1930s, the primitive biplanes of fabric and wood that had populated the skies above the battlefields of World War I had turned into all-metal monoplanes capable of carrying thousands of kilograms of explosives and traveling at hundreds of kilometers per hour. The microwave radar (also under rapid development at that time) was able to detect airplanes beyond visual range, but was not able to identify them. In fact, an enemy plane was only identified as such when in line of sight, often too late to prepare an appropriate response. It was exactly this inability that led to the incoming Japanese aircraft to be identified as a United States bomber flight, thus ensuring a surprise effect at Pearl Harbor in 1941.

The Luftwaffe solved this problem (initially) by using an ingeniously simple maneuver. During engagements with German pilots at the beginning of the war, the British noted that squadrons of fighters would suddenly (and simultaneously) execute a roll for no apparent reason. That curious behavior was eventually correlated with the interception of radio signals from the ground, and it then became apparent that the Luftwaffe pilots would roll when they received indication that they were being illuminated by their radar, in order to change the backscattered signal reflected from their airplanes. The consequent modulation of the blips on the radar screen allowed the German radar operators to identify them as friendly targets. This is, allegedly, the first known example of a backscatter radio link (according to [3] and [4] this is still an unverified report).

The first work exploring RFID was perhaps the landmark paper entitled “Communication by Means of Reflected Power” [5], published by Harry Stockman in October 1948. Stockman predicted that “...considerable research and development work has to be done before the remaining basic problems in reflected-power communication are solved, and before the field of useful applications is explored.” Indeed, commercial activities exploiting RFID only began to appear years later, during

the 1960s, with the first application being the Electronic Article Surveillance (EAS). The EAS was a simple 1-bit tag (since only the presence or the absence of the tag could be detected), meant to counter the theft of merchandise.

The rapid evolution of microelectronic technology during the 1970s triggered the development of a variety of practical applications for RFID, such as animal tracking, vehicle tracking, factory automation, access control and toll roads, and led to a mass deployment of RFID technology in the decade that followed.

The increase in the commercial use of RFID also prompted a need for standards, which led to many standardization activities in the 1990s. Most of these activities were conducted by the International Standards Organization (ISO) and the International Electrotechnical Commission (IEC). A relevant milestone came in 1996 with the standardization of RFID as a data carrier by the Article Number Association (ANA) and European Article Numbering (EAN) groups. In 1999, EAN International and the Uniform Code Council (UCC) of the United States, now both known as GS1, adopted a Ultra High Frequency (UHF) frequency band for RFID and established the Auto-ID Center at the Massachusetts Institute of Technology (MIT).

The Auto-ID Center then proposed a new Electronic Product Code (EPC) to identify products just like the standard bar code does. The available standards today for UHF RFID are the EPC Global and the ISO 18000 (in [3], some early military radio identification systems are described, in [6] a brief history of RFID is presented with a slightly higher level of detail compared to the text in this document, and if a more complete historical overview is needed, an appropriate reference would be perhaps [7]).

Infrastructure and terminology

The typical RFID system consists of one or more readers (either stationary or mobile) and many tags which by definition are attached to the objects they are meant to identify. In this system, a reader communicates with the tags located within its wireless range in order to collect data about their corresponding objects, usually in the form of an alphanumeric identification code (which is basically the serial number of the tag).

The reader, often also referred to as interrogator, is a fairly complex device in both hardware and software, and therefore it comes as no surprise that this device is usually the most expensive component in a RFID system, by a large margin. In order to be able to extract information from within a reasonable population of tags quickly and reliably, a typical reader normally provides a few antenna ports (rather than just one) and implements anti-collision algorithms. Ethernet and other communication ports are also frequently included to enable the reader to connect to higher level network infrastructures, or to provide access for configuration, for instance. Contrarily to the readers, the tags are much simpler, much smaller and much cheaper. In order to keep them that way (and therefore suitable for mass production) nearly all complexity is transferred to the side of the reader, hence making the resulting system extremely asymmetrical. Depending upon their operating principle, RFID tags can be classified as passive, semi-passive (or semi-active) or active ([3] and [8] describe not only readers but also all types of tags in detail, and contain photos of a few commercially available tags). It is useful to note that the acronym RFID, when expressed in general terms, refers to a system assumed not to contain active tags.

Passive tags are the least complex and hence the cheapest. They derive all their energy from an electromagnetic field supplied by a close-proximity antenna attached to an interrogator. The lack of an onboard power source allows operation with virtually no need for maintenance, but also seriously limits the range at which passive tags can be powered and read. Even so, and despite having up to several centimeters in length (considering UHF tags), passive tags can be paper-thin. Passive tags transmit data back to readers using a technique, called backscatter, that grants them the ability to send information without having to resort to power-expensive active transmitters. They usually consist of a single Integrated Circuit (IC) and an antenna, with the size of the former normally being insignificant when compared to the size of the latter (again, for UHF tags). The IC, because of being designed for a particular application rather than being general-purpose, is known as an Application-Specific Integrated Circuit (ASIC).

Semi-passive tags are reasonably similar to passive tags. Semi-passive tags include an onboard battery, but the latter is only utilized to supply power to their internal circuitry and perhaps to some peripheral sensors, if any. Even though their power source does not interfere with communication in any way, these tags generally achieve longer ranges, if compared to their passive counterparts, due to an increase in receive sensitivity. Nevertheless, with the longer range comes a limited life span that is closely tied to the battery.

Active tags are in fact full-fledged radios. They are equipped with a battery, a receiver, an active radio transmitter, and often include additional circuitry for control and sensing. In addition to an increase in functionality, as with the semi-passive tags, the battery affords these tags a greater wireless coverage, due to the introduction of filtering and amplification, and also due to a significant transmit power. Active tags are frequently capable of transmitting and demodulating sophisticated phase-based modulations, which can be more efficient users of available spectrum and provide superior noise robustness, in comparison to the amplitude modulation that passive and semi-passive tags are generally limited to. Because of having a more significant power consumption than semi-passive tags, active tags frequently deplete their batteries faster.

Near-field and far-field

RFID systems that operate at higher frequencies (above the 30MHz mark, according to [9]) are called far-field systems. In these systems, the readers use the real power contained in free space propagating electromagnetic plane waves to establish wireless links to tags, for power transfer and communication purposes. In contrast, near-field systems operate at lower frequencies and employ a different coupling mechanism, called inductive (or magnetic) coupling, that takes advantage the reactive energy located near the transmitting antenna.

Whether or not a tag is in the near or far field depends on how close it is to the power source and on the operating frequency (or wavelength, λ). As explained in [10], the near-field begins at the transmitting antenna and extends up to the surface of a conceptual sphere, known as radian sphere, whose radius is given in equation 1.1. Once outside the sphere, the electromagnetic waves separate from the antenna and start to propagate into free space as plane waves.

$$r = \frac{\lambda}{2\pi} \quad (1.1)$$

Equation 1.1 is only valid if the maximum dimension of the radiating structure (which in this case corresponds to D) is smaller than λ . Instead, if D is larger than λ , the boundary between the fields can be estimated using equation 1.2.

$$r = \frac{2D^2}{\lambda} \quad (1.2)$$

Nevertheless, the radius r is not precisely defined in either case, due to the fact that changes in electromagnetic fields occur gradually, meaning that the transitional state between the fields is not exactly represented by a thin spherical surface, but by a hollow sphere of relevant thickness instead. It should be noted that anything within the near-field region will couple with the antenna and distort its radiation pattern, thus making antenna gain a meaningless parameter under these conditions. Moreover, in the transitional region, as the antenna pattern is taking shape but is not yet fully formed, gain measurements vary with distance.

Radio regulations, EIRP and ERP

Since RFID readers generate and radiate electromagnetic waves, they are legally classified as radio systems. Naturally, the function of other radio systems must under no circumstances be disrupted or impaired by the operation of RFID systems, and due to that, the range of suitable operating frequencies available for RFID is limited. In fact, it is usually only possible to utilize frequency ranges that have been reserved specifically for industrial, scientific or medical applications, known worldwide as the ISM (Industrial Scientific Medical) frequency ranges.

Most of the far-field RFID activity is concentrated in a pair of fairly narrow ISM bands, first from 860 to 960MHz and then from 2.4 to 2.45GHz. Even though both of the frequency ranges are within the UHF band, which formally ends at 3GHz, only systems operating at the former are normally referred to as UHF systems. Readers and tags operating at the other ISM frequency band often receive the designation of microwave devices. In the 860 to 960MHz region, the worldwide regulatory environment is very complex, since RFIDs at these frequencies compete directly with cellular telephony and other popular and important applications, and as a consequence, different countries have made different choices about what can operate where. The 2.4 to 2.45GHz band, on the other hand, is available for unlicensed operation in nearly every major jurisdiction, but it is also crowded with other devices so that interference is a major issue.

In terms of power regulations the situation is basically identical, with the maximum power the readers are allowed to radiate being heavily dependent on the country. Table 1.1 shows the current radio regulations for UHF and microwave RFID systems in Europe (in [3], there is a figure that shows all frequency bands where RFID systems, near-field or far-field, may operate and organizes them according to how often are they utilized, [6] shows the most common bands and associates each one to relevant standards, typical applications, typical ranges, and so on, and [8] shows the differences in spectrum allocation and usage, considering Europe versus United States).

The same table also introduces the Equivalent Radiated Power (ERP), which is the power that would have to be supplied to an ideal half-wave dipole antenna in order to get the same electrical

field strength that the antenna under test produces, at a same distance. The Equivalent Isotropic Radiated Power (EIRP) is another reference power figure normally found in radio regulations and technical literature. The latter assumes an isotropic radiator instead of an ideal half-wave dipole as a reference antenna, but is otherwise identical to ERP ([9] explains these two reference power figures in greater detail). The ERP is currently the most frequently quoted power figure in RFID reader and tag specifications in Europe, whereas EIRP tends to be the preferred reference unit in the United States.

Table 1.1: European frequency and power regulations for both UHF and microwave RFID (in case a certain frequency has more than one maximum output power level associated with it, the higher power limit applies).

Frequency band [MHz]	Power [Werp]
865.0 .. 868.0	0.1
865.6 .. 868.0	0.5
865.6 .. 867.6	2
2.446 .. 2.454 GHz	0.5 or 4 (indoors) Weirp

Considering $G_{\lambda/2}$ to be the gain of an ideal half-wave dipole antenna (cited in many literature texts and equal to approximately 2.15dBi), the EIRP and ERP power figures that correspond to an antenna under test with a gain G and being supplied with a power equal to P_{in} can be calculated using equation 1.3.

$$P_{eirp} = P_{erp} G_{\lambda/2} = P_{in} G \quad (1.3)$$

Backscattering modulation technique

Backscatter modulators are currently very important in RFID, since they enable tags to send data to readers at little energy expense, by taking advantage of the reflection coefficient variation at the interface between the antenna of a tag and the tag itself.

Normally the antenna is designed to be impedance matched to its respective analog front-end so that no reflection occurs at the interface. When the tag decides to transmit data to the reader, it deliberately modifies the load seen by its antenna, thus forcing the latter to re-radiate a portion of the incoming power back to the reader. By alternating between these two states according to a data stream, an amplitude modulated response can be generated. In the extreme case, the load of the antenna switches back and forth from a perfectly matched load to a short circuit or to an open circuit. This makes the backscattered signal more pronounced and hence more easily detected by the reader, but also prevents the tag from receiving energy while in the reflecting state, because all the received power is reflected back.

Since it can be changed in both amplitude and phase, the reflection coefficient also allows the use of phase-based modulation techniques, which have the advantage allowing the tag to receive the same power regardless of being in communication or not, since the module of the coefficient

never changes during operation. However, selecting that modulus to be very small in an effort to increase the transfer of power to the tag results in a poorer communication performance, with any minor impedance deviation (introduced during manufacturing, for instance) leading to a dramatic variation in phase. Increasing the modulus in order to make the tag less sensitive to impedance deviations is not a viable option either, since the critical quantity currently limiting the range of passive RFID systems is the power transferred from the readers to the tags ([3] further describes backscatter in terms of modulation and coding schemes, while [11] addresses backscatter from a more theoretical point of view).

1.3.2 Wireless Sensor Networks (WSNs)

Similarly to many other technologies (RFID for instance), the development of sensor networks has been largely driven by military requirements. The SOund SURveillance System (SOSUS) was one of the earlier applications involving sensor networks and consisted of a network of acoustic sensors placed on strategic locations in the deep basins of the Atlantic and Pacific oceans to detect and track the quiet Soviet submarines during the Cold War ([12] traces the history of research in sensor networks over the last few decades). It should be noted however that in this case even though the sensing itself was wireless, the sensors were actually connected to one another (and to land base stations) using cables.

The first obvious sensor networks wireless application appeared during the Vietnam war, with the Air Delivered Seismic Intrusion Detector (ADSID). Each ADSID sensor node was more than 1 meter long and weighted approximately 17 kilograms. Equipped with a sensitive seismometer, these ADSID sensors were planted along the Ho Chi Minh Trail to detect vibrations from moving personnel and vehicles. The sensed data were transmitted from each node directly to an airplane, over a channel with a unique frequency. Even though the ADSID nodes were large, and the high energy cost of direct communication limited the lifetime of the sensors to only a few weeks, they successfully demonstrated the concept of wirelessly networked sensors (other early applications can be found in [12], [13] and [14]).

Modern research on wireless sensor networks started around 1980 and was mainly sponsored by the Defense Advanced Research Projects Agency (DARPA), the research and development office for the United States Department of Defense. At that time, wirelessly networked sensors became a crucial component of a new network-centric warfare, by enabling weapon systems to collaborate with each other over a communication network rather than operating independently, resulting in faster response times and improved detection and tracking capabilities.

Over the years, sensor nodes have been getting increasingly smaller, cheaper and more efficient and are now crucial in a multitude of applications, both military and civilian, mostly due to their outstanding adaptability (in [14], a few sensor nodes representative of the actual status of WSN technology are described in terms of hardware characteristics). And in fact, the era of networked sensing and communication is highly anticipated in the near future. The proof is that in September 1999, WSNs were identified by the "Business Week" as one of the most impactful technologies for the 21st century. Additionally, in January 2003, the "Technology Review" (published by the MIT) stated that WSNs are one of the top ten emerging technologies. It was also estimated that

WSNs generated less than 115 million euros in sales in 2004, but would reach more than 5 billion euros by 2010[13].

Overview

Building a sensor network first of all requires the constituting nodes to be developed and readily available. The wireless sensors, in turn, are usually highly dependent on the application for which they are going to be used. It is the latter that normally establishes the requirements with regards mostly to size, cost, energy consumption and communication. While there is certainly not a single standard available, nor would such a standard necessarily be able to support all application types, a certain common trend is observable in the literature when looking at typical hardware platforms for wireless sensor nodes.

The typical wireless sensor node is composed of a controller, a transceiver, a power supply, and one or more sensors and actuators, all integrated on a single or multiple Printed Circuit Boards (PCBs), and packaged in a few cubic centimeters (the following descriptions were largely based in those found in [15], because of their high level of detail, and complemented by the information contained in [16]).

Controller is responsible for processing all the relevant data and is usually capable of executing arbitrary code. It definitely plays a major role in managing data collection from the sensors, in performing power management functions, in interfacing the sensor data to the physical radio layer and also in managing the radio network protocol. In most cases, the controller is chosen to be a microcontroller, but it can also be a microprocessor, an ASIC, or even a Field Programmable Gate Array (FPGA). The choice obviously entails several trade-offs with regard to flexibility, performance, energy efficiency, and cost. Some of the key characteristics why microcontrollers are particularly suited to embedded systems are their flexibility in connecting with other devices (like sensors), their instruction set amenable to time-critical signal processing, and their typically low power consumption. In addition, they are freely programmable and hence very flexible, and often have flash memory built in. Moreover, most of the microcontrollers currently available have the possibility to reduce their power consumption drastically by going into sleep states where only parts of the controller are active. Digital Signal Processors (DSPs) provide another choice for a controller, these are specifically geared, with respect to their architecture and instruction set, for processing large amounts of vectorial data, as is typically the case in signal processing applications. In WSNs however, the requirements on wireless communication are usually quite modest and the signal processing tasks related to the actual sensing of data are not overly complicated either, which makes DSPs somehow overpowered for the task and hence rarely used. ASICs can be a superior solution, but only if the duties of the sensor nodes are not supposed to change over their lifetime. The number of nodes also has to be large enough to warrant the investment in ASIC development. However, using energy-efficient ASICs to perform fixed low-level processing tasks and leaving the remaining high-level, flexible and rarely invoked tasks to a microcontroller stands out as an attractive design and research option. FPGAs offer a tremendous flexibility as they can be fully reconfigured on-site, but also constitute the most energy-expensive solution, and are not compatible with traditional programming

methodologies. Also, it is important to realize that it is not practical to reprogram an FPGA at the same frequency a microcontroller can change programs.

Transceiver allows sensor nodes to wirelessly exchange data between themselves and consists of a receiver and a transmitter housed in the same package. Even though this is the standard configuration currently in use, it may prove useful to physically separate the two functions if the performance requirements for uplink and downlink communication are determined to be significantly different. Assuming the transmission medium to be the air, the most commonly utilized communication schemes make use of radio frequencies, optical communication or ultrasound technologies. Of these choices, RF-based communication is widely recognized as being the most relevant, since it fits the requirements for most WSN applications, as it provides relatively long range and high data rates, acceptable error rates at reasonable energy expenditure, and does not require line of sight between sender and receiver. However, this is the only option that requires antennas, which (depending on the operating frequency) may increase the size of a sensor node substantially. In general, the radio transceivers used in sensor nodes are based on typical RF frontend configurations using one or more intermediate frequencies and consisting of the expected elements, such as Power Amplifiers (PAs), Low Noise Amplifiers (LNAs), mixers and oscillators. Such mixture usually yields a decent radio coverage, but at the cost of power consumption. In fact, the transceiver is nearly always the single most power-expensive component in a sensor node. These conventional transceivers, despite being obviously important for radio communications, will not be further addressed in this document. Independently of their power consumption, virtually all transceivers share a common issue related to the minimum power spent while maintaining a minimal listening functionality. During the time a transceiver is waiting for a transmission to come in it must be powered on so that the wireless channel can be observed, thus spending energy without any immediate benefit. While it seems unavoidable to provide a receiver with power during the actual reception of a packet, it would be desirable not to have to invest energy while the node is only waiting for a packet to come in. This represents a critical trade-off between energy efficiency and wakeup latency. In an attempt to alleviate this issue, current wireless sensor nodes often include a specialized receiver circuit, called a wakeup receiver, designed to notify other components of an incoming packet without needing a significant amount of power to do so (more information on wakeup receivers in [17]).

Power supply provides the energy necessary for a proper operation. The power source of a node often consists of a single battery, either non-rechargeable or rechargeable. A rechargeable battery, as the name suggests, can be recharged once depleted, but in general has a smaller capacity when compared to a non-rechargeable alternative having roughly the same size. In both cases, the rated battery capacity specified by a manufacturer is only valid as long as maximum discharge currents are not exceeded. Capacity drops or even a premature battery failure quickly become more likely events if such limits are ignored. This issue becomes very important knowing that a typical sensor node consumes quite different levels of power over time (depending on the mode of operation), and may draw a significant amount of current from the battery once in a while. For most manufacturing technologies, the larger the battery, the more power that can be delivered instantaneously. Another critical topic

related to batteries is their ability not to lose charge when no power is drawn. How rapid this self-discharge phenomenon occurs highly depends on the type of battery and plays an important role in the selection of a battery-based energy source for a sensor node. In case of a rechargeable battery, an efficient recharging is equally important, especially considering the low and intermittently available recharge power that energy harvesters are usually only able to produce (in case the wireless sensor is supposed to make use of energy harvesting techniques). Conventional battery technology essentially does not recharge in these extreme conditions, and even when it does, the capacity gradually decreases over charging cycles due to the so-called memory effect.

Sensors and actuators are highly application dependent. Sensors are usually available in many different forms with many individual peculiarities precisely due to this fact. Currently, there are sensors designed to sense or measure a wide range of physical quantities. On the other hand, despite being as diverse as sensors, actuators are frequently only utilized to open or close a switch or a relay or to set a certain value in some way. The consequences that arise from these actions (such as unlocking a door or starting a motor, for instance) usually fall under the responsibility of other higher-powered systems. Actuators are sometimes paired with controlling sensors in order to create a feedback mechanism to inform the sensor node in case some malfunction occurs. Currently, there are many trade-offs associated with sensors and actuators, which include accuracy, sampling speed, dependability, cost, size and form factor and, very importantly, energy consumption.

Additional components may include secondary memories for extended storage space, location finding systems, Analog to Digital Converters (ADCs), push-button switches for reset and other functions, and sometimes a small backup battery designed to retain clock and date information while the primary battery is being replaced.

Radio regulations

Most (non-military) WSNs operate in the license-free ISM frequency bands that were previously introduced, and therefore, the power limitations that apply to RFID also apply to WSNs. However, while the readers often radiate all the power they are allowed to broadcast in an effort to increase range, most sensor nodes do not come close to exceeding the imposed limits, otherwise they would deplete their batteries too fast. At least nowadays, power regulations do not pose a limiting factor for (conventional) WSNs.

Network topologies

Once deployed, wireless nodes quickly organize themselves to form a network (so that tasks that demand a cooperative effort can be achieved), normally using star or mesh configurations, both depicted in figure 1.1. Other more sophisticated arrangements, such as hybrids between the two configurations mentioned, are not uncommon ([15] includes a comprehensive analysis of networks and network-related concepts in sensor networks, yet, in the framework of this work, the content in [18] was considered to be sufficient).

Figure 1.1: Star and mesh network topologies, represented on the left and right sides of the figure, respectively (in both network configurations, the empty circles represent sink nodes whereas the filled circles correspond to ordinary sensor nodes).

In a star network there is a unique sensor node, usually called a sink, whose task is to manage the other sensor nodes that are part of the system. A sink sends messages to other nodes, reads their responses and usually either logs the relevant events or reports them directly to higher level network infrastructures. Each sensor node only connects with the sink and is not even aware of other nodes operating within the network. In other words, ordinary nodes are not permitted to send messages to each other. This constraint requires all these nodes to be within radio range of the sink at all times, and therefore makes the network less robust due to its dependency on a single node to manage the entire system. However, the star network is also characterized by its simplicity and its ability to keep the power consumption of the nodes to a minimum precisely because of this constraint. Another important advantage of this topology is that it also allows for low latency communications between ordinary sensor nodes and the sink. Therefore, it comes as no surprise that this was (and still is) the network topology of choice for RFID (not only because of the aforementioned advantages, but also because this particular configuration facilitates the transfer of power from readers to tags and allows for the tags to be simple). In practice, the sink does not necessarily need to be a mobile sensor node and in fact, it sometimes is a stationary base station connected to the power grid, that is, the complete opposite.

A mesh network still assumes the same two classes of sensor nodes but eases the restrictions on communication by allowing each node to communicate to any other node within the network and also by allowing any node to act as a sink (as needed). Once a node wants to send a message to another node that is out of radio range, it can use one or more intermediate nodes to forward the message to its intended destination. If an individual node along the communication path fails, other nodes can still route the message using another path. In addition to adding redundancy to the network, this multi-hop communication scheme also allows mesh networks to outpace star networks in terms of radio coverage, as the range of the former is not necessarily limited by the range in between single nodes. In fact, the higher the number of nodes, the greater the potential range of the network. Nevertheless, nodes that implement multihop communications often have a higher power consumption compared to those who do not have such capability. Additionally, as the number of communication hops to a destination increases, the time to deliver the message also increases, especially if low-power operation is a requirement.

A hybrid between the star and mesh networks (a star-mesh hybrid as it is often called) yields a configuration that can be both robust and versatile, while keeping the power consumption of the sensor nodes to a minimum. In this arrangement, the network is composed of a set of higher power nodes with multihop capability which connect to one another following the rules of a mesh network. Attached to each one of these nodes in a star fashion are other (less capable) nodes who are not enabled with the ability to forward messages. Other more specific network arrangements are sometimes found in practice, depending on the application requirements.

Power saving techniques

Regardless of the topology selected, the average power consumption of a sensor node should be always as reduced as possible. In an effort to achieve that goal, virtually all wireless sensor nodes implement some kind of technique to reduce their energy consumption.

A commonly used power saving technique consists of placing the sensor node in a low power consumption sleep mode whenever possible, waking it up only when an interrupt is raised from a peripheral such as a wakeup receiver or from an internal timer, for instance. Assuming the sleep mode power consumption to be much lower than the active power consumption (which is what happens in practice), this duty-cycling approach results in a remarkably smaller average power consumption. As would be expected, the average power consumption decreases even more as the duty-cycle ratio decreases. As this tendency continues, the sleep power consumption starts to become more critical than the active power consumption in determining the average power consumption. Hence, and due to the fact that most sensor nodes operate at low duty-cycles, the sleep power consumption is often referred to as an important parameter in the design of a sensor node whereas the active power is sometimes seen as not being so critical. Yet, the low duty-cycle approach is not without practical consequences, because the lower the duty-cycle is, the greater the communication latency will be.

A more sophisticated possibility than using discrete operational states is to adapt the speed with which the sensor node operates according to the task it has to perform. One rather obvious solution (that may not be so obvious to implement in practice) is to switch the controller to full operation mode, compute the task at highest speed, and go back to sleep mode as quickly as possible. An alternative approach would be to compute the task only at the speed that would be required to finish it just before its deadline. The rationale is the fact that a controller running at a lower speed, that is, lower clock rates, consumes less power than operating at full speed. Even though the task takes longer to accomplish, the balance in terms of energy expenditure generally turns out to be positive (a more detailed explanation of this can be found in [15]). In addition, some controllers allow their supply voltage to be reduced in a significant manner at lower clock rates while still guaranteeing correct operation.

1.3.3 Antennas and radio wave propagation

At the most basic level, an antenna is a specialized device that converts the guided waves coming from a feeder cable, waveguide or transmission line into radiating waves travelling in free space, and vice versa (hence, a transducer). Antennas are already a key part of almost every wireless system because they allow wireless communication (in fact, that is their sole function in almost

every application), and they quickly become even more critical if the devices to which they are attached to are supposed to completely rely on them, not only for communication, but also for energy harvesting purposes, which is the case of passive RFID tags.

Antenna parameters

Both sensor nodes and tags normally utilize omnidirectional antennas because of their radiation pattern, which is reasonably similar to that of an isotropic radiator, but realizable. Such design choice is supported by the fact that sensor nodes and tags are generally assumed to be mobile (or at least randomly deployed). Even though a more directional antenna would enable tags to operate at a greater distance from the readers and sensor nodes to be placed further apart from one another, most applications simply cannot guarantee a precise alignment between tags and readers or between wireless sensors at all times. Additionally, even considering omnidirectional antennas, there is still the problem of polarization mismatching. In order to circumvent this issue, most practical RFID systems utilize circularly polarized antennas for readers and linearly polarized antennas for tags, in an effort to provide a suitable power transfer from readers to tags, regardless of how the latter are oriented in space, while simultaneously keeping tag antennas simple and cheap. It is true that when a circularly polarized wave reaches a linearly polarized antenna half of its power is simply discarded, but at least that loss does not increase as the receiving antenna is tilted at any angle within the plane perpendicular to the axis of propagation ([3] continues to explain antenna polarization from the point of view of RFID).

In addition to an isotropic-like radiation pattern, tag antennas are usually required to be as small as possible and sometimes even physically flexible, while maintaining a robust and durable structure and a low price (the principal design considerations for tag antennas are summarized in [19]). Even though the size of an antenna ultimately depends on the selected operating frequency, there are techniques that can effectively reduce the size of an antenna without modifying this important system parameter, but those will only be addressed later. For the time being, it is only important to note that these techniques inevitably lead to a loss of efficiency that increases as the ratio of antenna dimension to wavelength decreases. This trade-off between size and efficiency is very important in the design of passive RFID tags, because a higher efficiency improves range in a significant manner but a smaller size, on the other hand, frequently leads to a lower manufacturing cost, which is something of uttermost importance in a mass production environment. According to [20], there are three efficiencies associated with an antenna that once combined together yield the total (or overall) antenna efficiency given in equation 1.4.

$$e_0 = e_r e_c e_d \quad (1.4)$$

The total efficiency (which is represented by e_0) is defined as the ratio of the radiated power to the total input power supplied to the antenna. The first partial efficiency (e_r) is used to take into account losses due to reflections arising from an impedance mismatch between the antenna and whatever is attached to it. The second and third partial efficiencies that appear in the equation are related to ohmic losses in conductors (e_c) and in dielectrics (e_d). Because of these ohmic losses, part of the power that is delivered to the antenna is simply dissipated as heat (instead of being radiated). As can be concluded from the expression, an antenna designed to harvest power from

RF waves obviously benefits from a reduced reflection coefficient, but it is important to note that conduction and dielectric losses are equally relevant.

Operation in different frequency bands

Instead of applying (often non-trivial) size compression techniques in an attempt to obtain smaller antennas, a simpler option would be to increase the operating frequency, thereby reducing the size of the antenna without compromising efficiency. Unfortunately, such approach generally does not yield the results that would be expected due to the fact that different operating frequencies have different propagation characteristics and are subject to different power regulations.

Comparing the UHF and microwave frequency ranges introduced earlier, the latter is found to provide a larger bandwidth, in addition to a very reasonable propagation through non-conductive materials, which include wood and wood-based products, natural and synthetic garments, and plastics. Also, good reflections off metal surfaces allow for a fairly good propagation in cluttered environments. Inexpensive flexible antennas that are able to withstand considerable bending are commonly achievable at this higher frequency. In contrast, the microwave range exhibits a reduced propagation distance in comparison to UHF, if the same transmit power is considered. Not only that, microwave active devices are often more expensive than their UHF counterparts and more susceptible to electronic noise. Moreover, multipath and fading effects may need to be taken into account at these higher frequencies. Due to the added cost and complexity of a microwave solution, the UHF band is currently the most popular choice for RFID, and represents a suitable choice (in principle) for other electromagnetic energy harvesting systems as well. Nonetheless, it should be noticed that UHF electromagnetic waves are partially absorbed by liquids and behave quite unpredictably near metal (the information contained in this paragraph was extracted from [3], [9] and [10], and it is important to mention that in these references other frequency ranges besides UHF and microwave are compared in terms of radio wave propagation).

Free space path loss

The power that is radiated from a transmitter antenna spreads in all directions as it moves through free space and away from the source (how it actually spreads depends on the radiation pattern of the antenna), meaning that the power absorbed by a receiving antenna diminishes as the two antennas are placed further apart from each other. This loss in signal strength is called the free space path loss, and does not depend on the gains of the antennas nor on any loss arising due to hardware imperfections. Considering both antennas to be 100% efficient and perfectly aligned in terms of directional gain and polarization and assuming no multipath effects, the power the receiving antenna captures can be estimated using the Friis transmission equation, which is displayed in equation 1.5.

$$\frac{P_r}{P_t} = G_t G_r \left(\frac{\lambda}{4\pi R} \right)^2 \quad (1.5)$$

The equation relates the available power at the terminals of the receiving antenna (given by P_r) with the transmitted power P_t and takes into consideration the gain of the transmitting and receiving antennas (represented by G_t and G_r , respectively), the wavelength λ corresponding to

the operating frequency, and the distance that separates the antennas (that is, R). While the ideal conditions required by the Friis formula are almost never achieved in practical systems due to obstructions, reflections from objects and most importantly reflections from the ground, this formula still provides at least a reasonable first approximation. Although there are other versions of this equation that take into account several non-idealities that the basic version does not (they can be found in [20]), those will not be considered.

1.4 Unusual applications involving WSNs and RFID

RFID technology is playing an essential role in improving safety and productivity into one of the dirtiest environments, the mining industry. Maintaining the needed safety precautions throughout a complex network of tunnels in a mine, usually far below the surface, became easier and more efficient with RFID. The geographical zones in a Norwegian mine are charted out in a centralized control room, and employees, suppliers, visitors, and vehicles who enter the mine are registered and tagged with RFID tags. This grants security personnel the ability to track movement of people and assets through the mine from the control room. In addition, the tags are also used for access control. In the case of an emergency, finding those who are injured or trapped can be extremely difficult because of poor visibility and ventilation conditions. RFID allows a much more effective search in a time-critical situation, allowing mining operations to be much safer.

Investigating the reduction of the salmon population in the Columbia river and ensuring the proper tire management in NASCAR races are other examples of RFID unique applications ([8] describes these applications and describes some more).

Sensor networks also have several unique and interesting applications of their own, ranging from accurate sniper localization and smart self-healing mine fields using acoustic sensors to bird observation on Great Duck Island, Maine (in [21], other examples can be found, such as the one that follows, for instance).

In particular, wireless sensing technologies are being used to monitor glacier environments at Briksdalsbreen, Norway, in order to better understand glacier dynamics in response to climate changes. Wireless sensor nodes equipped with pressure and temperature sensors and a tilt sensor for measuring orientation are deployed in drill holes at different depths in the ice and also beneath the glacier. The nodes communicate with a base station mounted on top of the glacier, which measures supra-glacial displacements using differential GPS and transmits the collected data via GSM. Once the sensor nodes deplete their batteries, their service life ends, since they are not recoverable after deployment. Each node consists of three octagonal PCBs (for analog, digital and radio sub-systems) and a polyester egg-shape capsule.

Chapter 2

State of the art

In this chapter, the first section shows that the concept of combining RFID and WSNs is feasible and gaining momentum in terms of research. The second and third sections introduce the results of integrating RFID technology into WSNs and sensors into passive RFID tags, that is, the Wireless Passive Sensor Networks (WPSNs) and the RFID Sensor Networks (RSNs), respectively. The fourth section concludes this chapter by discussing antennas designed specifically with energy harvesting applications in mind.

2.1 Introduction

There is no doubt that RFIDs and WSNs are emerging as two of the most ubiquitous computing technologies in history due to their important advantages and their high applicability. While the industry has witnessed rapid growth in developing and applying RFID technology, and the network research community has devoted great efforts to improve WSNs, these two communities would benefit greatly by learning from each other. In that sense, it is indeed encouraging to notice that they are already beginning to converge. The integration of RFID and WSNs is now an imminent step that will lead to a higher level of synergy and more technological advances, and there are already some ongoing research projects aiming in that direction, in which tags are integrated with sensors, sensors are integrated with tags, readers are outfitted with interfaces for communicating with sensors, and so forth ([22] describes in detail the various forms of integration that may be used, and in [23] a system for dispensing medication to elderly people using both RFID and WSN devices is presented).

It is true that these technologies have almost completely converged considering active RFID tag solutions, in which each tag frequently carries one or more sensors, in addition to its unique identification number. Additionally, while still capable of communicating to readers using RFID protocols, active tags often provide support for some (multi-hop) communication standards well known to sensor networks. However, in this work only a mixture between passive tags and sensor nodes is considered, so that their individual advantages can be brought together to create a single device that combines the battery-free operation and low power consumption of the former with the ease of reprogrammability and modular design of the latter.

2.2 Wireless Passive Sensor Networks (WPSNs)

Unlike traditional sensor nodes, wireless passive sensor nodes are powered remotely by an external RF power source and remain operable as long as power is delivered in. Once grouped together, these nodes form WSPNs. By definition, passive sensors have no batteries and utilize modulated backscatter to send data to sink devices, which can be the same devices that supply power to them (resembling what RFID readers do) or not. In addition, there may be multiple sources of power in a same network, arranged in a clustered architecture for instance, in case a single source proves not to be enough (advantages and disadvantages of different powering topologies are discussed in [24]). If no RF power sources are active at a given moment of time, the energy stored in the passive sensor nodes eventually depletes and they cease to function until a sufficient amount of energy is once again accumulated as a result of energy harvesting.

One quite interesting feature of passive sensor nodes is that their duty-cycle automatically adapts itself according to the power available within the network. For instance, if little power reaches a passive node, the latter takes longer to recharge its transient energy reservoirs, and therefore begins to spend more time sleeping (that is, recharging) and less time in active mode, resulting in a lower duty-cycle (but of course, there is a threshold below which the node simply does not recharge, regardless of the time window). Once more power becomes available, the node starts to recharge faster, and because of that the duty-cycle goes up. In general, passive sensor nodes assume capacitors as energy storage devices, but other options exist (a very comprehensive review of potential power sources for wireless sensor nodes is given in [25]). Super capacitors, for instance, are portable sized capacitors with a very high value that are commonly used for buffering energy, without the aging or rate capacity effects found on batteries.

In terms of hardware, a wireless passive sensor node deviates from a typical sensor node on the power unit and the transceiver as explained, but otherwise there are no major differences. Even though passive nodes lack the capacity to communicate with one another through a cooperatively formed ad-hoc network, they are definitely beyond simple passive tags (more information regarding WPSNs can be found in [26] and [27]).

2.3 RFID Sensor Networks (RSNs)

Current passive tags are fixed function devices that use a minimalist and non-programmable state machine to report a hard-coded identification number when energized by a reader. While these tags find uses in many applications (mainly where low price outweighs functionality), a new class of sensor-enhanced passive tags is currently emerging and in fact, "...in the near future, tags will likely be used as environmental sensors on an unprecedented scale", according to [28]. Passive sensor tags, as their name suggests, are passive tags that contain onboard sensors. These tags are able to report back to RFID readers not only their (static) identification numbers but also the dynamic data collected by their sensors. Once queried by a standard RFID reader, sensor tags sample their sensors and encode the resulting data into the packet they would normally send. The reader is usually configured to just forward the packets it receives to an application developed for the particular purpose of extracting the actual sensed data, as it would be unable to decode the new (and non-standard) data on its own. Since passive sensor tags are designed to operate

in conjunction with conventional RFID readers, RSNs can be seen as an extension of the existing RFID technology, in which passive tags gain additional sensing capabilities.

While the only characteristic that distinguishes passive sensor nodes from passive sensor tags is perhaps the wireless technology they have evolved from, it seems that passive sensor nodes are generally envisioned as replacements for conventional sensor nodes in applications in which there is little problem in sacrificing some range if that effectively translates into no more battery replacements, whereas passive sensor tags tend to aim at the original vision of smart-dust that was articulated in the late 1990s, where walls, clothes, products, and personal items are all equipped with wirelessly networked sensors. For this vision of dense sensor deployment to become a reality, truly unobtrusive sensing devices are of primary importance, and that is why size and lifetime are considered to be two extremely important issues in RSN design ([29] describes some applications suited to the space between existing WSNs and RFID and highlights the most important research challenges in realizing RSNs).

2.3.1 Intel Wireless Identification and Sensing Platform (WISP)

One of the most notable examples of a RFID sensor tag is the Intel Wireless Identification and Sensing Platform (WISP). The WISP is a software-defined hardware platform for computation and sensing, designed to be compatible with basically any standards compliant RFID reader. The most distinctive features of this device are a wireless power supply, backscatter communication and a fully programmable 16-bit flash microcontroller. WISP devices operate in the UHF RFID frequency band.

The first WISP, called the α -WISP, consisted simply of an antenna, two RFID IC chips and two mercury switches. The switches were mounted in a geometrically anti-parallel configuration in order to experience an acceleration identical in magnitude but opposite in sign, so that they could select which chip is to be connected to the antenna and which is not. Once queried by a reader, the α -WISP returns one identification number while under positive acceleration and other while under negative acceleration, thus behaving as a 1-bit accelerometer. In this case, the switches function both as sensing and modulating elements.

The α -WISP later evolved into a more complex triaxial accelerometer with one bit of dynamic range per axis. This new device, called the π -WISP, included three mercury switches (instead of one) for sensing and one analog high frequency switch for selecting the appropriate tag chip according to the output of a MSP430F1121 low-power microcontroller, thus clearly separating sensing and modulating functions. The system also included an RF energy harvester composed of a matching network and 4 cascaded voltage doublers, so that the microcontroller would be able to operate from the power supplied by the RFID reader, instead of using a battery (more detailed information about earlier WISP platforms can be found in [30]).

Newer WISP devices are able to communicate arbitrary, multi-bit data in response to a single RFID poll event and include a super capacitor having a capacity in the order of 0.1F to maintain sensing and logging capabilities up to 24 consecutive hours when out of range of a reader (in [31] the most recent version of the platform is discussed, but [32], [33], [34] and [35] all contain valuable information related to various hardware revisions and refinements).

2.4 Antennas for electromagnetic energy harvesting

In terms of commercial availability, nearly all antennas designed for use with UHF tag chips consist of a thin layer of some conductive material on top of a substrate that can be either rigid or flexible, depending on the application. Whatever the case may be, the truth is that an enormous variety of antenna shapes exist on the market today. According to the manufacturers, different antenna shapes are designed to meet specific criteria requirements (optimized operation near liquids, for instance). What is very interesting to notice is that many of those complex shapes are actually modified half-wave dipoles.

First of all, the size of a UHF half-wave dipole is still too big for many applications. In order to overcome this problem, the idea is simply to squeeze it. By bending the arms of the dipole, its linear extent can be reduced. A dipole that is shortened in this manner is (unsurprisingly) called a bent dipole. As more bends are added, the dipole starts to look like a river meandering across flat terrain, as in figure 2.1, and becomes known as a meandered dipole. However, the more bending is done to fit the antenna into a small space, the lower the radiation efficiency becomes. This makes sense, since the direction of current flow is inverted in some segments of the dipole because of the bends. When looking at the dipole from a long distance (in other words, far-field), these opposed currents cancel themselves and therefore contribute with no radiation. In an unmodified dipole no current cancellations occur, simply because there are no opposed currents.

Figure 2.1: The result of bending an arm of a dipole antenna several times, following perhaps the most common of many techniques.

Second, even once the dipole antenna is squeezed down and an acceptable trade-off between size and efficiency is determined, there is still a problem in terms of impedance matching. As the impedance of the IC is nearly always unchangeable, the responsibility of matching falls entirely on the antenna. It is precisely due to this fact that commercially available tag antennas normally include some kind of matching structures, frequently realized as lengths of conductive line. The matching is usually done at the minimum threshold power level enough to turn on the chip, in an effort to maximize range. In certain cases, matching may become easier if relatively large areas of conductive material (circles, for instance) are attached to the ends of the dipole, in a process called capacitive tip-loading (this technique is explained in [3]).

Third, the width of the arms of the dipole effectively establishes a trade-off between quality factor and bandwidth. Using thinner lines yields a more resonant antenna, potentially allowing a very effective impedance matching, but only if the antenna is designed and built with enough precision, so that it indeed resonates at the frequency it was designed to. Because of the narrow bandwidth, any frequency deviation (even if little) can cause a significant increase in the reflection coefficient, possibly rendering the antenna useless at the desired frequency. Also, the thinner the

lines are, the more precise the manufacturing technique must be. In contrast, larger lines lead to a wider bandwidth antenna that is much less sensitive to frequency deviations, but at the cost of a lower quality factor (this subject is further explained in [3]). Finally, it should be noted that this reasoning applies not only to printed antennas, but to other antenna types as well.

A particular tag antenna

A quite good example of a RFID tag antenna designed with high efficiency (and hence high read range) in mind is presented in [36]. The antenna has an overall efficiency of 95% at a frequency of 915MHz (which corresponds to an ISM range within the United States), which is more than what leading commercial designs are able to offer (according to the reference). The antenna was based on the half-wave dipole and was fabricated on a very thin (and flexible) layer of Liquid Crystal Polymer (LCP) substrate. It consists of a pair of slightly bent dipole arms and two stubs designed for matching the antenna to its corresponding IC. One of the stubs controls mostly the resistive matching while the other controls the reactive matching. This omnidirectional antenna yields a practical read range of about 9.4 meters, which is higher than the 7.9 meters that a commercially available and slightly larger tag antenna obtains, under equal conditions. In terms of radiation diagram, this particular tag antenna shows patterns identical to those that would be expected of the traditional half-wave dipole. Also, according to this source, most commercially available tag antennas suffer from low efficiencies, in the order of 50 to 60%.

A particular application of RF energy harvesting

An application involving electromagnetic energy harvesting, in which RF power is harvested from a TV tower located at a distance of 4.1Km is reported in [37]. In this experiment, and as described, the harvesting system consisted of a broadband log periodic antenna having a gain of 5dBi (a commercial model designed for TV applications), a matching network and four cascaded voltage multipliers. Once the matching network was tuned to the 674 to 680MHz frequency band in which 960KW ERP were being broadcast by the tower, $60\mu\text{W}$ were able to be harvested in a continuous fashion. While not much, this value is not far from the $220\mu\text{W}$ that the Friis formula yields, and was enough to power a commercially available temperature and humidity meter equipped with an LCD display that, under normal circumstances, would work from a 1.5V AAA battery.

Chapter 3

System architecture

In this chapter, the first section introduces the proposed wireless sensor and antenna, describes them with regards to capabilities and morphological characteristics, and also contains information related with the selected controller. The second section briefly introduces and reviews the major software tools that were utilized in hardware and software development. The third and last section of this chapter presents a detailed hardware overview of the sensor node.

3.1 Introduction

The wireless sensor node that was developed in this work is made of a 0.813mm thick RO4003C high frequency laminate with thin layers of copper on both sides. The bottom layer of copper was designed to be a ground plane and therefore only the top copper layer was utilized for soldering hardware parts. Nonetheless, as the layout in figure 3.1 shows (on a 2:1 scale), part of the routing was done on the bottom side of the board. Even though some of the ground plane was removed because of the routing, the portion corresponding to the high frequency frontend (on the right side of the layout) was left intact. In terms of dimensions, the board measures 55.1mm long by 25.6mm wide. Choosing to implement the sensor node in PCB had the advantages of a fast design iteration time and a low development cost, and additionally yielded a very flexible (that is, modular) solution. Even though custom IC implementations were not considered, it is worth to note that they typically consume less power than discrete PCB designs and result in a noticeably smaller form factor.

The developed sensor node does not include any form of internal power source, and depends exclusively on the energy harvested from a nearby RF power source radiating at a frequency of 866.6MHz. Since maximizing the range of the sensor node was one of the goals, the UHF band was chosen over the microwave band due to its potentially higher range. Selecting the operating frequency from within the UHF range consisted basically in calculating the central frequency of the band in which more power could be radiated without exceeding radio regulations, given that more radiated power translates into a longer range. In terms of communications, the sensor node supports the same type of bi-directional communication that passive RFID tags support.

In addition to having a wireless power supply, the sensor node also integrates a programmable 16 bit flash microcontroller from the MSP430 family. In particular, the selected MSP430F2132

device (in the 28TSSOP package) combines a number of characteristics that make it well suited to meet the stringent power requirements imposed by the lack of an internal power source on the sensor node. Rather than having just two modes of operation, one for sleeping (basically shutting down everything) and other for active operation, this microcontroller allows various modules to be independently controlled. For instance, the microcontroller can be placed in a deep sleep mode while leaving a clock signal running (for providing a time reference while the device sleeps). To do this, it is only necessary to simply switch off all modules except the clock source that is supposed to remain active during sleep. At the deepest sleep mode, that is, shutting down all modules and optionally leaving interrupts active to trigger a change of mode, the MSP430F2132 consumes as little as $0.1\mu\text{A}$. A low sleep current consumption (as this one) is of uttermost importance for a sensor node which depends on energy harvesting to charge its energy reservoirs. Apart from having various sleep modes and a low sleep current consumption, the selected microcontroller switches from sleep mode to active mode in less than $1\mu\text{s}$ (according to the manufacturer). This is a very important feature for the sensor node to have, because the faster its controller boots, the more responsive it becomes. In terms of power supply voltage, the microcontroller allows various choices that can influence maximum operating frequency, write access to the integrated flash memory, or both. For instance, in this application the supply voltage was set to 1.8V, which is the lowest level for which the microcontroller still guarantees reliable operation. However, operating at this lower supply voltage limits the operating frequency to 6MHz (with higher supply voltages the device can go up to 16MHz) and also disables flash program and erase operations. In terms of storage capacity, the MSP430F2132 provides 8KB of flash (although higher capacities are available in other devices of the same family, those models tend to consume more power) that can be used for code and data. Finally, it is worth to mention how extremely flexible the clock system of this microcontroller is. Its processor can be fed by a variety of different clock sources, allowing it to be fast when speed is indeed needed (while in communication, for instance) and to save energy when the tasks to be carried out do not require higher clock rates (a example of a task where higher speeds are not normally required is sampling sensors). Still regarding the clock system, it is important to note the the startup time mentioned earlier assumes the processor to be linked to the Digitally Controlled Oscillator (DCO) built into the microcontroller, due to the fact that distinct clock sources exhibit distinct startup times (more information on the MSP430 platform can be found in [38], [39] and [40]).

As can also be viewed in figure 3.1, most of the pins of the microcontroller were routed to either one or another row of larger vias (located on the upper and lower parts of the figure), in an effort to not only facilitate debugging tasks but also to provide the support for experimenting with the sensor node without having to rebuild the layout each time a new hardware component is to be tested. Besides free ports, some important signals from within the circuitry of the sensor node were also connected to these rows, allowing measurements to be made without risking to provoke a short circuit or to damage a copper line. Naturally, the programming interface of the MSP430F2132 was also linked to a set of consecutive pins in one of the rows. Though not shown in the layout, each one of the groups of larger vias is designed to be soldered to a row of right angle pin headers with standard 2.54mm spacing (which corresponds to the same spacing that breadboards utilize).

Figure 3.1: PCB layout of the developed sensor node shown on a 2:1 scale (darker-colored lines refer to the top copper layer whereas lighter areas and lines refer to the bottom layer, the circles represent vias and the remaining drawings refer to the component outlines).

In terms of copper traces, a single width was used for all lower frequency circuitry, whereas several larger widths were used in the design of the RF frontend of the sensor node. In the latter case, the widest line width corresponds to a small line segment that is visible on the right side of the figure that contains the layout. This larger line was designed to accommodate an edge-mount SMA connector, so that an antenna system could be attached to the sensor node, enabling it to wirelessly receive power and exchange data.

As with the sensor node itself, its antenna system was also developed in this work, and the latter consists of a printed dipole-based antenna. The PCB layout of the antenna, which is 148mm long by 8mm wide, is shown (on a 1:1 scale) in figure 3.2. Contrary to the previous board, this antenna was built on a much cheaper FR4 substrate with a thickness of 0.8mm. The reason why a high quality laminate was not used was because in this case neither a tightly controlled dielectric constant nor a very low loss tangent were needed, as these parameters were determined to cause little effect on the performance of the antenna. Therefore, FR4 was considered to be a suitable choice for being cheap, despite its lossy behavior (that tends to become more significant at higher frequencies) and its dielectric constant that varies widely among different vendors (and even between different orders from the same vendor). The most notable feature of this antenna is its overall efficiency in the order of 96% (although it should be noted that this efficiency was obtained in simulation and was not actually the result of a measurement). Similarly to what was done in the sensor node, the antenna was also designed to accommodate an SMA edge-mount connector to be soldered to the larger copper areas that can be seen at its center.

With its antenna attached, the proposed sensor node was found to be able to harvest enough energy to switch itself on and communicate at up to 4.1 meters away from a RF source setup to radiate 2W_{erp} at a frequency of 866.6MHz.

Figure 3.2: Proposed dipole-based PCB antenna designed for use with the sensor node, shown on a 1:1 scale (the rectangular area represents the substrate and the other two structures correspond to the arms of the dipole, with the upper arm being on the top side of the substrate and the lower arm being on the bottom side).

3.2 Description of the tools used

A few distinct major software tools were utilized in developing both the sensor node itself and its antenna system. The Easily Applicable Graphical Layout Editor (EAGLE), the Advanced System Design (ADS), the High Frequency Structure Simulator (HFSS), and the Code Composer Studio (CCS) are briefly introduced and reviewed in the subsections that follow.

3.2.1 Easily Applicable Graphical Layout Editor (EAGLE)

EAGLE layout editor is a schematic capture and PCB layout tool which consists of a schematic editor, a layout editor and an autorouter (which was not used in this work), all embedded in a single user interface which, while not being flawless, was found to be quite simple to use. Due to its simplicity, the interface did not require much time to get used to it. This concept of an unified interface enables changes in the schematic to be immediately reflected into the layout without having to export the netlist and start a layout program in a separate session. In this software, each component is comprised of a symbol (which is a graphical representation to be used in the schematic), a package (which contains the footprint to be used in the layout) and a device, which is a combination of the previous two. It should be noted that existing symbols and packages can be shared between multiple devices, provided the former are pin-compatible. At a higher level of organization, devices can be grouped into libraries. Even though the library concept is good and a large number of libraries is bundled with the software, finding specific components proved to be a time-consuming task because of a poorly designed searching mechanism. Even with this drawback this application still provides an easy to use, powerful and affordable schematic capture and PCB design package.

The layouts depicted in figures 3.1 and 3.2 were made using the Light Edition of EAGLE 5.10, which is an edition that can be used for free with the following limitations (that did not pose any practical problem though):

1. The useable board area is limited to 100mm by 80mm;
2. Only two signal layers can be used;
3. The schematic editor can only create one sheet.

Apart from these three limitations, the Light Edition of EAGLE can do anything the (more expensive) Professional Edition can do, according to the manufacturer, CadSoft Computer. This application is available for Windows, Linux and Mac platforms.

3.2.2 Advanced Design System (ADS)

ADS is an electronic design software for high frequency applications. This software provides an integrated design environment for developing products such as wireless networks, high speed data links, satellite communications, or radar systems. It supports schematic capture, layout (which was used in this work only to verify the consistency of schematics), frequency domain and time domain circuit simulation, and electromagnetic field simulation (which was not used at all during this work). This software application has a somewhat steep learning curve, not only because of an user interface filled with less visible and relatively hard to find advanced options, but also due to a scarcity of tutorial resources. Nevertheless, ADS is indeed a very powerful and versatile software tool, developed and maintained by Agilent.

In the framework of this work, ADS 2009 Update 1 was utilized to create a simulation model to mimic (and optimize) the process of harvesting energy from high frequency radio waves. It is important to notice that this simulation model refers solely to hardware parts and copper traces

within the sensor node itself, and thus does not take the antenna into account. Agilent currently maintains versions of ADS for operating systems based on Windows and Linux.

3.2.3 High Frequency Structure Simulator (HFSS)

HFSS is a full-wave electromagnetic simulator for arbitrary three-dimensional volumetric passive device modeling, which relies on the finite element method to compute the electrical behavior of high frequency and high speed components, such as on-chip embedded passives, integrated circuit packages, PCB interconnects, transmission lines, antennas (and arrays of antennas), biomedical devices and others. The finite element method is a numerical analysis technique utilized to obtain solutions to the differential equations that describe (or approximately describe) a wide variety of physical problems. The underlying premise of the method states that a complicated domain can be subdivided into a series of smaller regions in which the differential equations are approximately solved. By assembling the set of equations for each region, the behavior over the entire problem domain is determined. Users are only required to specify a geometry model using a combination of geometric solids and more advanced techniques such as parametric curves if necessary, material properties and the desired output. The software will automatically generate an appropriate mesh for solving the problem after some iterations. As would be expected, the more sophisticated the simulation model is, the longer the corresponding mesh takes to converge. Simulations normally take a couple of minutes considering very simple conceptual structures, however, once the model becomes more complex (in order to get closer to reality), the simulation times usually increase drastically. Fine tuning the performance of a model using trial and error techniques is therefore a time-consuming task. Nonetheless, this application yields reasonably accurate results and allows virtually any structure to be simulated, regardless of its size or shape.

The antenna designed for use with the sensor node was fully designed using HFSS 10, which was developed by Ansoft. The latter still continues to improve HFSS and has been releasing new versions regularly, for Windows and Linux platforms (as a matter of fact, two major versions were already released since version 10).

3.2.4 Code Composer Studio (CCS)

CCS is a integrated development environment for digital signal processors, microcontrollers and application processors that is developed and maintained by Texas Instruments. The application consists of a suite of tools utilized to develop and debug embedded applications, which includes a source code editor, a project build environment, a debugger, a profiler, simulators and as it would be expected, compilers for all device families. The user interface was found to be well organized and easy to use, perhaps in part because it is based on the Eclipse framework, which is an open source framework used by many embedded software vendors. However, it is also true that this software suffers from some syntax recognition problems (in dealing with macros that represent constants, for instance) and is somewhat prone to random crashes from time to time.

CCS 4.1.2 was used in this work for writing the source code for the MSP430 chip integrated into the wireless sensor, for downloading software onto the device and additionally, for debugging purposes. There are two install images available for CCS, with one of them being a full featured paid version, while the other being a free and code size limited version only supporting a limited

number of devices. The free version was sufficient for this work, since it offers a reasonably high code size limit of 8KB for MSP430 devices. CCS currently requires a Windows-based operating system to run (however, it is worth to note that there is an open source software development toolset known as the MSPGCC that enables MSP430 devices to be programmed and debugged in operating systems other than Windows).

3.3 System overview

A block diagram of the proposed wireless sensor node is shown in figure 3.3. When received by the antenna, an incoming RF signal passes through the impedance matching network, then entering the energy harvester block. In conjunction with the voltage regulator block, the energy harvester rectifies the received energy into a regulated voltage to power the system. Connected directly to the output of the power harvester, the demodulator block extracts a data stream from the incoming carrier wave (in case the latter is modulated). The extracted baseband waveform is then read by the MSP430 microcontroller. Uplink data is sent via the backscatter modulator, which was implemented with a single transistor acting as a binary switch.

In order to be energy-efficient, this sensor node (and in fact, any other node that depends on electromagnetic energy harvesting) has to overcome three major design hurdles related to high frequency analog hardware design, software and sleep power consumption.

The wireless range of a sensor node decreases dramatically if its energy harvester is unable to efficiently gather energy from incoming electromagnetic waves, regardless of finely tuned software routines and a minimal sleep power consumption. On the other hand, an efficient energy harvester becomes irrelevant if there are no suitable software algorithms to properly manage the available energy, independently of how small the sleep power consumption is. Finally, a significant sleep power consumption can also severely cripple the maximum wireless range of the node due to an additional and continuous waste of energy that hinders it from recharging.

Figure 3.3: Block diagram of the proposed sensor node (in which the tips of the arrows indicate in which directions power or data, or both, are allowed to flow).

In an effort to reduce as much as possible the current that the sensor node consumes while sleeping, its hardware parts were carefully selected to have reduced quiescent currents. Also, the number of components was kept to a minimum because, as a general rule, the higher the number of components is, the greater their overall quiescent current will be (the most notable exception to

this rule is perhaps the case of power switches). Following these guidelines yielded a sensor node that consumes as little as between 2 and $5\mu\text{A}$ while sleeping as deep as possible. In this state no clock sources are assumed to be active (that is, the sensor node loses track of time) and therefore the only way to trigger the sensor node to switch into active mode is by sending some data to it, assuming of course its energy reservoir does not deplete in the meantime. In this context, it is useful to note that the chosen MSP430F2132 is able to retain Random Access Memory (RAM) contents even if its power supply voltage drops to a value as low as 1.6V (according to datasheet specifications). Once the sensor node switches to active mode, its current consumption increases to roughly 1mA at 6MHz (mostly because of the microcontroller). However, it should be noted that neither the active current nor the sleep current were actually measured and these estimates were calculated based upon information collected from datasheets.

It is important to note that the developed sensor node uses solely commercial off-the-shelf hardware components. In addition, all component footprints were built from scratch in order to allow the sensor node to be small and simultaneously easy to assemble using manual soldering equipment. Also important to note is the fact that some of the hardware components and circuits used in the proposed sensor node were inherited from original WISP designs.

3.3.1 Energy harvesting and communication mechanisms

Figure 3.4 contains a schematic diagram showing the proposed circuitry for the energy harvesting and communication subsystems for the sensor node. Regarding not only this figure but also other figures containing schematics related to the sensor node:

1. Pin numbers are indicated next to the corresponding devices unless unnecessary;
2. Unless stated otherwise, passive devices assume 0603 packages;
3. Nearly all signal names follow the original MSP430 pin naming;
4. Those who do not follow that convention begin with a number sign.

In order to match the sensor node to an impedance of 50Ω , a matching network comprised of an inductor and a capacitor was utilized. Connected to the matching network is a multi-stage voltage multiplier composed of ten diodes linked to each other in a configuration that is known as the Dickson charge pump (as concluded in [41], this manner of connecting the diodes yields better results compared to the other known basic form of charge pump, in which the nodes of the diode chain are coupled to the input via capacitors in series, rather than in parallel). In order to reduce voltage losses within the voltage multiplier circuit, Schottky diodes were used instead of the more conventional PN junction diodes. Because of the diodes it contains, the voltage multiplier becomes a non-linear device, with an efficiency that depends on its load. Therefore, the front-end of the sensor node must be tuned to provide the maximum output voltage in the presence of the desired load. In addition, the voltage multiplier also acts as an envelope detector, as it follows amplitude changes in the incoming carrier wave.

Continuing the analysis and ignoring the upper circuit branch for the time being, following the voltage multiplier are two diodes and a large storage capacitor (compared to other capacitors

Figure 3.4: Proposed circuitry for energy harvesting and communications (individual hardware parts are highlighted with filled areas and accompanied by relevant information).

in the schematic). When exposed to a RF power source, the sensor node remains sleeping until this capacitor charges up, and once sufficient voltage is obtained, it switches to active mode. The amount of time the sensor node is able to sustain continuous active mode operation increases as the storage capacitor increases. In contrast, a smaller storage capacitor leads to a more responsive system as it takes less time to charge. In order to prevent the storage capacitor from discharging through the power harvester in the absence of external power, a different type of Schottky diode was added to the design. This diode offers a reduced reverse leakage current while maintaining a low forward voltage drop, at the cost of a worst high frequency performance (that in fact does not represent a drawback since the diode only needs to operate at lower frequencies). The other diode is a transient voltage suppressor, which was added to protect voltage-sensitive components from voltage spikes that would otherwise damage them. Under normal operating conditions, this diode behaves as an open circuit. To the right of the large capacitor is a low-dropout voltage regulator, whose function is to generate a regulated voltage of 1.8V that powers not only the MSP430 microcontroller but also the majority of the peripheral circuits. A capacitor was connected to the output pin of the voltage regulator to ensure a smooth and stable output voltage.

The upper circuit branch refers to the demodulator circuit. When there is only energy being transferred to the sensor node, the carrier wave remains at a constant amplitude level. During communication, it drops to zero during short amounts of time to mark the boundaries between bits, with the time elapsed from one mark to the next indicating a logical one or zero. From the energy efficiency point of view, this encoding technique is suitable for the sensor node as it keeps the carrier wave active most of the time, thus maintaining a decent energy transfer even during communication. The data signal at the output of the power harvester is fed through a low-voltage comparator that thresholds this waveform to remove noise and glitches. Unfortunately, despite being fast, this comparator is somewhat power-expensive, and that is why a transistor switch was connected to it. Had this power switch not been added, the sleep power consumption would have been drastically increased. As the comparator has an open-drain output, a pull-up resistor was added to the design. For being open-drain, the output of the comparator becomes sensitive to stray capacitances due to the length of the copper traces used in the layout. In an attempt to alleviate this issue, the copper traces connected to the pull-up resistor were made as short as possible. A diode and a capacitor were connected to the positive input terminal of the comparator in order to create a slowly varying average voltage level, thus providing a dynamic reference for bit detection. A suitable value for this capacitor (which also serves as a power supply for the comparator) was determined experimentally. Lastly, an analog switch was included to convert the relative magnitude of the baseband signal that appears at the output of the comparator into a 1.8V logic level for the MSP430. Appearing at the lower left-most corner of the schematic, the backscatter modulator consists of a single transistor connected to the input terminals of the sensor node. When the transistor is in conducting mode it short circuits the analog front end, changing the load impedance of the antenna but also preventing the sensor node from receiving power and data. In the non-conducting state, the transistor has no effect on the rest of the system, and the power harvesting and data downlink functions occur as if it were not present. In terms of encoding, the same technique that was used for downlink was also used for transmitting bits, only in this case the event that marks the boundary between bits is the backscatter modulator switching from an open state to a closed one and then quickly back.

3.3.2 Supply voltage supervisor

Because of the limited rate at which the energy harvester is able to gather energy from incoming radio waves, the microcontroller should never be programmed to start executing tasks as soon as the voltage it needs to operate builds up at the output of the voltage regulator (with regard to this subject, it should be mentioned that the MSP430F2132 that was integrated into the sensor node was found not to require a supply voltage of 1.8V to boot up, but instead only 1.5V, although the reliability below the 1.8V threshold is obviously questionable, since the device was not designed to operate at that voltage). In case it is, the sensor node will most likely run out of power before completing a task because of the active current consumption that drains its storage capacitor at a rate much faster than the energy harvester can compensate. In order to increase the amount of time the sensor node can stay active (in a continuous fashion), the supply voltage supervisor in figure 3.5 was added.

Figure 3.5: Proposed supply voltage supervisory circuit (once again, individual hardware parts are highlighted with filled areas and accompanied by relevant information).

Once the voltage in the large storage capacitor reaches 2V, this circuit sends a signal to the microcontroller (that triggers an interrupt), allowing the sensor node to switch to active mode with a 0.2V voltage room. Stated another way, the voltage supervisor prevents erratic operation by introducing hysteresis, but does that at a cost of range. The higher that the threshold voltage of the supervisor is set, the longer the sensor node is capable of sustaining a continuous active mode operation, potentially allowing more demanding tasks (that should not or cannot be divided into smaller tasks) to be executed. However, a higher threshold requires the energy harvester to generate a higher output voltage, which in turn requires the sensor node to be placed closer to the RF power source.

The supply voltage supervisor consists of an open-drain voltage detector with a 2V voltage threshold, a resistor and an analog switch identical to that chosen for the demodulator. While the large storage capacitor charges, the voltage detector holds its output down. It is important to note that there is a current flowing through the resistor in this situation that adds to the overall sleep current consumption. Once a voltage of 2V is reached, the output of the voltage detector becomes high impedance, allowing the resistor to pull the output high. The pull-up resistor was chosen high enough so that the overall sleep power consumption would be only slightly affected, even knowing that such design decision ultimately leads to a longer response time. Similarly to what was done with the demodulator, care was taken to keep the copper lines connected to the output of the voltage detector as short as possible.

3.3.3 Non-volatile memory

Since the flash memory built into the microcontroller cannot be written to at 1.8V, the sensor node loses the ability to retain data across power cycles (increasing the supply voltage only to gain write access to the flash memory is not a viable option because of the higher power consumption it conduces to). In order for it to gain that ability back, the non-volatile EEPROM memory module that is presented in figure 3.6 was added.

The schematic diagram of the memory module is comprised of a capacitor, two resistors and an Electrically Erasable Programmable Read-Only Memory (EEPROM) chip of 1KB. As the latter draws a significant current while active, a capacitor serving as a transient buffer was placed (physically) close to it. The two resistors are required by the 2-wire Inter-Integrated Circuit (I2C)

communication protocol that the microcontroller and memory chip utilize to communicate with each other.

Figure 3.6: Non-volatile memory circuitry (as with previous figures, individual hardware parts are highlighted with filled areas and accompanied by relevant information).

3.3.4 Other peripheral components

Figure 3.7 shows the last two hardware components yet to introduce from all those which were included into the sensor node. The first is a low frequency crystal oscillator designed to operate at 32768Hz. It provides a stable and precise yet low-power clock signal that can be utilized as a main clock for the microcontroller (instead of the DCO) to save energy. Nonetheless, it should be noted that low frequency crystal oscillators usually take from several hundred milliseconds to a few seconds to startup (although factors other than frequency also affect the startup time of an oscillator, as discussed in [42]), which is much more than the DCO takes. It is also important to mention that the MSP430F2132 also features an internal low power oscillator that provides a typical frequency of 12KHz, which despite not being very accurate, is still an option.

Figure 3.7: Other hardware components integrated into the sensor node, which, from left to right correspond to a crystal oscillator and a to a single resistor.

Finally, the second hardware component is simply a resistor that is necessary for the 2-wire JTAG programming protocol (also known as the Spy-Bi-Wire) to function properly. The process of downloading content into and debugging the developed sensor node only involves connecting a total of four wires (JTAG plus power and ground) to a programmer that supports JTAG, which in the framework of this work was a EZ430-RF2500 development tool.

Pin interfaces

Figure 3.8 shows which signals are connected to each one of the two rows of pins the sensor node has, and to which positions they go. The first four pins belonging to the connector that is on the right side of the figure constitute the programming interface. Rotating this figure ninety degrees

counter-clockwise and superimposing it on figure 3.1 (which contains the PCB layout) yields the correct pin mapping.

Figure 3.8: Two pin interfaces where a number of signals from within the developed sensor node are routed to, having a combined total of 26 pins.

Chapter 4

System optimization

In this chapter, the first section presents some considerations for selecting a proper load to attach to a simulation model of the RF front-end of the developed sensor node, and briefly mentions the transition from hardware to software, in terms of optimization. The second section explains how the RF front-end of the sensor node and its antenna system were designed. Concluding this chapter, the third section discusses the software layer that was developed for the microcontroller integrated into the sensor node.

4.1 Introduction

Once all of the lower frequency circuits within the sensor node were designed, the sleep current consumption that resulted from the sum of all quiescent currents was taken into consideration to model a load for the energy harvester, designed to mimic the real load the harvester sees while the sensor node is sleeping. Out of all the circuitry downstream from the output of the energy harvester, only the large storage capacitor and the Schottky diode that precedes it were included into the simulation model of the load (for simulating the behavior of the diode, a SPICE model provided by NXP was utilized). All other circuits were replaced by an ideal current source that was placed in parallel with the storage capacitor, and whose value was set equal to the estimated maximum sleep current consumption (that is, $5\mu\text{A}$). While this approach was chosen because it favors range, other different strategies could have been employed if different features were to be maximized instead, for instance, a load based on the active current consumption would sacrifice range for an increased data throughput.

Finally, the last layer of optimization refers to the software routines that are loaded into the sensor node, which are responsible for encoding and decoding bits, and for managing the energy that the harvester is able to gather (and also for sampling sensors, in case those are attached to the system).

4.2 Radio front-end performance and design factors

As the version number printed on the (previously shown) PCB layout suggests, the current sensor node was preceded by a less-capable device. From one version to the other, a couple new signals

were added to the pin interfaces and, more importantly, the RF front-end was re-designed based on different considerations. In the first version, all hardware components were assumed to behave according to the typical values found in their datasheets, in terms of quality factors (which are critical in high frequency design) and quiescent current consumptions. However, this design option led to experimental results that were reasonably different from those obtained in simulation, and as a consequence, the sensor node was found to require considerably more power to startup than what was initially predicted. In contrast, the second approach assumed a worst-case scenario, in which minimum quality factors and maximum quiescent currents were considered, and resulted in much better (and more consistent) practical results.

Considering the antenna system of the sensor, two different versions were also built, however in this case only a couple of small adjustments were made from the first to the second version, in order to compensate a minor frequency deviation that the first prototype exhibited.

4.2.1 Matching network and voltage multiplier (envelope detector)

Even after the functional blocks of the RF front-end and the topology for the voltage multiplier were chosen, a reasonable amount of design issues still needed to be dealt with. Presented in the form of a list, the most important design questions (which represent several trade-offs that can seriously affect not only how well the sensor node harvests electromagnetic energy, but also how it communicates) include:

1. How many stages to use in the voltage multiplier;
2. Which diodes and capacitors to select;
3. Which type of matching network to use and how to dimension it;
4. How to design the copper traces to connect all these hardware parts;
5. Which kind of substrate to use.

The number of stages of the voltage multiplier represents a trade-off that involves size, cost and harvested voltage. Earlier simulations have indicated that the harvested voltage increases in a logarithmic fashion with the number of stages, meaning that for a given input power level, the harvested voltage effectively increases each time a new diode is added to the voltage multiplier, but the effect becomes less pronounced as the number of diodes grows. In addition, more diodes occupy more space and make the sensor node more expensive (as a matter of fact, these diodes were the most expensive components utilized). Considering these trade-offs, a 10 diode multiplier was chosen because adding more diodes showed little improvement in terms of harvested voltage and resulted in a more complicated layout. On the other hand, less diodes led to significant drops in the harvested voltage. The criteria for selecting the diodes was a high detection sensitivity at the UHF band and the existence of a simulation model, features that were found in the HSMS-285x series (both the diode chip itself and its package could be accurately modeled using information contained in [43] and in [44], respectively).

Earlier simulations also showed that the harvested voltage increases as the capacitors within the voltage multiplier increase, also in a logarithmic fashion. However, it should be noticed that

no combinations of different capacitors were ever considered and a single capacitance value for all capacitors was always assumed in simulation. A value of 15pF yielded a only a little voltage drop compared to much larger capacitors and was thus chosen, not only because of that but also because that value was found not to interfere (in a significant manner) with the capability of the voltage multiplier to demodulate incoming bits. A multiplier composed of larger capacitors was verified to take a longer time to react to amplitude changes in the incoming carrier wave (which makes sense because the larger the capacitors, the lower the frequency at which they begin to behave like short circuits).

Unsurprisingly, the output of the voltage multiplier was found to be dependent on variations in capacitance that occur on each capacitor (usually due to limitations in precision related with how they are fabricated or imposed by the materials that compose them) but most capacitors have tolerances that are tight enough to avoid any problem. In contrast, the matching network can reduce the output of the voltage multiplier drastically if the tolerances of its components are not kept very small (in fact, the matching network is the part of the system that is more sensitive to hardware imperfections). Additionally, it was concluded in simulation that utilizing more than two matching elements in an attempt to improve matching was not a viable alternative because the better matching did not compensate the added power losses arising from resistive parasitic elements, especially in the case of inductors (key parameters for selecting RF inductors are given in [45]). Due to the fact that the parasitic resistance in an inductor increases with its inductance, inductors were kept as low as possible (in fact, there are some pairs of elements other than those used in the matching network that yielded a better matching but a lower harvested voltage simply because the inductor had a higher value to which corresponded a higher parasitic resistance). It is important to note that designing the matching network using microstrip lines instead of lumped components would in principle conduce to more precise results (and would also allow non-standard inductance and capacitance values to be used), but this technique was not considered, mainly because the free space within the sensor node that could be potentially used for printing stubs or (other kinds of matching structures) was simply too small.

With regard to copper traces, both lower and higher frequency PCB layouts obviously aimed at the same goal of occupying the least amount of space, but while the major difficulty at lower frequencies was mainly the significant number of electrical connections to make (although details such as placing the crystal oscillator away from frequently switching signal lines in order to avoid crosstalk were taken into consideration), higher frequency copper lines demanded a more careful design because they were intended to be included in simulation. The simulation model that was created to optimize the RF front-end of the proposed sensor node is shown in figure 4.1, in which line lengths and parasitic resistances associated with capacitors and with the only inductor used (whose values were based on information contained in their datasheets) are shown. All microstrip lines up to the output of the voltage multiplier were included in the simulation model exactly as they were printed into the PCB layout (no more lines were included because higher frequency signals end at the output of the multiplier). Higher frequency copper lines were made larger than their lower frequency counterparts, not aiming at an improvement in harvesting performance (as a matter of fact, the same performance could be achieved with both line widths), but because larger lines are less affected by degradations in width introduced during manufacturing, in terms

of percentage. Other lines could be made thinner to save space and to facilitate routing, given the fact that only low frequency (and low current) signals would travel through them.

Figure 4.1: The model developed for simulating the RF front-end of the sensor node, in which microstrip line lengths are shown (in millimeters) and parasitic resistances are either added in the form of resistors, or in the case of the inductor, included directly into the component to which they refer to. Different line widths were utilized in the model, corresponding to 0.7, 0.8, 0.9 and 1.77mm, as well as circular vias with diameters of 1.5mm.

Even though copper traces were designed as short as possible, so that power losses could be minimized, simulations indicated that their electrical size (although small) was non-negligible in terms of harvesting performance. Additionally, the harvested voltage was also found to be quite dependent on the loss tangent of the substrate. Nevertheless, the RO4003C was chosen not only for its low loss tangent, but also because of its accurately defined dielectric constant, since the simulations indicated that variations in this property would conduce to impedance mismatches that in turn would result in significant power losses (a variation in the dielectric constant leads to an impedance mismatch due to the fact that it effectively changes the electrical lengths of all copper lines). Table 4.1 shows the main characteristics of the RO4003C (that were obtained from [46]), and compares them with the typical values for FR4.

Table 4.1: A comparison between the RO4003C and FR4, in terms of dielectric constant and loss tangent characteristics at UHF.

	RO4003C	FR4
Dielectric constant	3.55 ± 0.05	usually from 4 to 5
Loss tangent	0.0021	0.02

Still regarding the simulation model, it is relevant to refer that the widest line segment was calculated so that the result would be a 50Ω line. Moreover, the backscatter transistor was not included into the simulation model due to the lack of a suitable simulation model. However, and since the transistor was experimentally verified to cause virtually no disturbances on the rest of the node (only while acting as an open circuit, of course), no significant problems arose from not including it into the simulation, except for the ability to predict the input impedance of the node with its modulator switched on.

4.2.2 Proposed antenna

Since the traditional half-wave dipole already provides an omnidirectional radiation pattern and a natural impedance fairly close to 50Ω (which is useful, because in principle, the simpler the matching structures are, the less power they dissipate), it seemed like a logical starting point for an antenna. Even though it would be possible to remove the matching network from the sensor node and directly match the antenna to the resulting impedance (which is what happens in fact with passive RFID tags), that approach was not considered because the sensor node was intended to have a 50Ω interface, so that standard antennas, cables, connectors and measuring equipment could be utilized with it without additional impedance transformations.

With regard to simulation models, the proposed antenna started as a simple half-wave dipole made of a pair of copper wires. The wires were later replaced by thin copper traces printed onto a FR4 board, in an effort to improve manufacturing precision (so that prototypes would resemble their simulation models as closely as possible). Furthermore, one of the two arms of the antenna was moved to the opposite side of the substrate, so that an edge-mount SMA connector (whose three-dimensional model was also built), would be easily soldered to them (printing the arms into opposite sides of the substrate was found to have no effect on the characteristics of the antenna, which makes sense, because the thickness of the substrate is actually very small compared to the length of the copper traces). Once the simulation model containing the antenna and the connector was completed, many parametric simulations were run, aiming at the goal of improving efficiency while retaining all the other dipole-like characteristics. In order to achieve that goal, many degrees of freedom within the antenna were explored, which included varying the length and width of its arms, placing copper areas at its tips, changing the areas of copper at its center, sliding one arm over the other, adding parasitic copper shapes near the arms, and so on (although no parasitic structures were actually included in the final antenna, it was interesting to observe the substantial effect they produced on its radiation pattern). Simulations demonstrated that while the behavior

of the antenna is affected by the size of the dielectric it is built on, virtually no changes occur when the loss tangent or the dielectric constant are varied in a significant manner (which was why FR4 could be used without causing any major efficiency loss).

4.3 Software algorithms and power management

Two distinct programs (with different purposes and levels of complexity) were developed for the MSP430F2132. One of them is a very simple program that was only designed to keep the sensor node in a known state during some tests, whereas the other is significantly more elaborate, as it combines communication and energy management functions. Both programs are described in the sections that follow.

4.3.1 Deep sleep mode

The purpose of this program is simply to make the sensor node switch to its deepest sleep mode immediately after starting up. Once the MSP430F2132 begins to execute the code, it disables its watchdog timer. It then sets all unused pins as outputs and assigns them a logical value (in order to avoid floating input pins that can increase power consumption), disables the receiving circuitry of the node, and finally, switches its processor and clock generators off. The sensor node then remains in this state until its energy depletes.

4.3.2 Proposed program for communication and control

Like in the previous program, in this the microcontroller also switches to sleep mode immediately after starting up, however, in this case an interrupt is associated with the pin to which the output of the voltage supervisor is connected to. Once a sufficient voltage is accumulated at the storage capacitor, the supervisor triggers an interrupt that awakes both the processor and the DCO inside the microcontroller.

Once in active mode, the microcontroller increases the frequency of its DCO from the default value (which is slightly above 1MHz) to 6MHz (the change of frequency remains in effect until the microcontroller runs out of energy, which means that the operating frequency does not need to be set each time the microcontroller switches into active mode), switches the receiver on, sets an interrupt to detect incoming rising edges and then switches to sleep mode again. Once the end of an amplitude gap in the incoming carrier wave is caught, the microcontroller awakes, sets a timer (fed by the DCO) running, and stops its processor until the next gap comes, at which time the timer is stopped and its value read. In case the time elapsed from a gap to another exceeds $32\mu\text{s}$ the bit is interpreted as a one, else a zero is assumed. In the meantime, the timer is already counting the time to the next gap. The microcontroller expects 8 bits at a time (that correspond to 9 gaps) to be delivered in this fashion. If by some reason a gap is not detected within less than $55\mu\text{s}$ of the previous one, the microcontroller simply discards all received bits and goes back to the state it was just before receiving the first.

When a sequence is correctly received, the sensor node inverts all bits and sends them back as a reply, after a delay of $100\mu\text{s}$ (this would be a suitable time to execute some action based on

the received data, most likely involving sampling one or more sensors and sending back the sensed data, but as the sensor node itself includes no sensors and none were attached to it, the chosen response at least demonstrates that the device is indeed capable of processing and transmitting data dynamically). During transmission, the backscatter transistor remains short-circuited for $8\mu\text{s}$ between bits to mark transitions, and the bits themselves last either 32 or $16\mu\text{s}$, for ones or zeros, respectively. Additionally and very importantly, the microcontroller switches its processor on only when the state of the backscatter modulator needs to be altered (although the DCO stays active during the entire transmission to provide a clock signal to the timer that triggers the processor to go into active mode as necessary).

It should be noted that communication bit rates can be further increased only by improving software, as both the backscatter modulator and the receiver can operate at higher rates (as a matter of fact, the receiving circuitry within the sensor node was experimentally found to be still capable of detecting transitions in a carrier wave modulated with a square wave having a period as low as $4\mu\text{s}$).

Chapter 5

Experimental results

In this chapter, the first section introduces what were the general objectives of the experimental tests made. The second section describes in detail the tests that were performed on the sensor node, first without and then with its antenna attached. Finally, the third section presents the tests that were conducted on the antenna.

5.1 Introduction

Once built, the developed sensor node and its antenna were subjected to a few tests (separately and together), in order to evaluate their performances and, very importantly, to find out the level of agreement between predicted and practical results, under different combinations of power and frequency conditions.

5.2 Harvested voltage, range and power consumption

In order to enable signals to be injected directly into the sensor node itself, the latter was first tested without its antenna. In this condition, the device was evaluated in terms of its harvested voltage at different input power levels and in terms of reflection coefficient, considering different power levels and also different frequencies. With its antenna attached, the node was then tested in conjunction with an actual RF source, in order to determine its wireless range. With exception of this test, in which the main program was loaded onto the sensor node, all other experiments were conducted with the simpler program (deep sleep mode) running.

5.2.1 Harvested voltage as a function of input power

In this test, the sensor node was connected to a signal generator configured to output a sinusoidal waveform at a frequency of 866.6MHz. The output power of the latter was then slowly adjusted first from -12 to 2dBm and then from 2 back to -12dBm, staying at each measurement point for a couple of seconds and then taking the average voltage harvested within that time window, the latter calculated with an oscilloscope. Gathered data related to both ascending and descending power levels, as well as the predicted curve, are all plotted in figure 5.1.

Figure 5.1: Harvested voltage seen at the storage capacitor within the sensor node in function of increasing and decreasing power levels injected at its input (plotted with filled and empty circles respectively), in comparison to the predicted results (dotted line with borderless circles), and always considering a sinusoid at 866.6MHz. Solid lines were plotted to highlight points located at -11.44 and -8.64dBm.

First of all, the figure shows that experimental results are following the trend of the predicted results, and are even moderately better, something that is not unreasonable since the simulation model considered a worst-case scenario. Additionally, there is clearly a hysteresis behavior in the harvested voltage, which means that for a number of power levels the sensor node can be either turned on or turned off, depending on its power history. Also in the figure, the minimum power the sensor node requires to turn on and the power below which it is no longer able to sustain its sleep mode are both highlighted in the figure.

When at higher power levels (more or less above -2dBm), the effect of the transient voltage suppression diode included into the sensor node becomes clearly visible. Naturally, the predicted curve shows no clipping only because the diode was not included into the simulation model (due to the fact that only lower power levels were of interest).

5.2.2 Frequency and power sweeps

In order to plot the reflection coefficient at the input of the sensor in function of frequency and power, the latter was connected to a network analyser. In both sweeps the network analyser was configured to take measurements only at a certain set of points (either in frequency or in power, depending the type of sweep in question), and to insert a time delay of 5 seconds between consecutive measurements. In other words, and in case of a frequency sweep (the same happens

in a power sweep), the network analyser samples the reflection coefficient at a certain frequency, waits for 5 seconds, moves to the next higher value and takes another sample, waits for another 5 seconds, and continues to execute this algorithm until the maximum frequency set in the sweep is reached, at which time it loops back to the lower end of the sweep. Especially in the case of power, having a known and rather slow sweep algorithm (in order to allow the sensor node to stabilize between measurements) becomes of critical importance in order to help reduce erratic behaviors caused by the sweeps themselves.

Regarding the frequency sweep, figure 5.2 shows that experimental results are quite close to those obtained in simulations, considering a frequency range starting at 766.6MHz and ending at 966.6MHz. Nevertheless, in terms of phase, initially there was a little difference between the two sets of results (in other words, in a Smith chart there was a rotational displacement between predicted and practical results). In order to compensate that deviation (that was probably caused by the SMA connector), a segment of an ideal microstrip line having an electrical length of 9 degrees was inserted at the input of the simulation model, so that phase values would match at 866.6MHz. While this calibration naturally did not interfere with the expected magnitude plots, it did improve (even if slightly) phase results related to the power sweep.

Regarding the power sweep, shown in figure 5.3, the same set of points used for measuring the harvested voltage was also used in this case. In terms of magnitude, it can be concluded that the experimental results tend to follow the predicted curves, except for the discontinuity which is triggered at the moment the sensor turns on. In terms of phase, that discontinuity is even more pronounced. Lastly, the effect caused at higher power levels by the transient voltage suppressor is also somewhat visible in this figure.

5.2.3 Operating distance in practice

In this test, the practical setup consisted of a signal generator, a power amplifier and a patch antenna with circular polarization and a gain of 5.5dBil (which expands as dBi linear), taken from an european commercial UHF RFID system. At first, the signal generator was connected to the input of the power amplifier and setup to output a sinusoid at 866.6MHz. Its output power level was then adjusted, so that the amplifier would generate enough power for the antenna to radiate a 2W_{erp} sinusoid. Finally, the sensor node was slowly moved towards the power source up to the point in which its voltage supervisor triggered a change in operation mode. After repeating the same procedure for several times it was determined that the sensor node would wake up always somewhere around 4 meters away from the source (with variations of basically no more than a few centimeters). Moving the sensor node away from the power source instead, and observing the maximum distance at which it would turn off yielded only slightly better results, in the order of just one tenth of a meter.

The generator was then configured to output an amplitude-modulated signal (centered at the same frequency the previous sinusoid was), in which a sequence of bits was encoded and repeated indefinitely, at regular and fairly spaced time intervals. Under these conditions, the previous tests were repeated, in this case to evaluate at which distances would the sensor start and cease to generate replies (which were determined by continuously monitoring the state of its backscatter modulator). In terms of results, the measured distances were basically equal to those obtained

Figure 5.2: Reflection coefficient measured at the input of the sensor node (represented by filled circles) in function of frequency, plotted against the predicted magnitude and phase curves (dotted lines with borderless circles) and assuming a test power of -8.64dBm. Solid lines were utilized to mark points at 866.6MHz.

before, mostly because the transmitted signal barely changed from one test to another, in terms of average power.

Even though in all tests both the transmitter antenna and the sensor node were maintained always at least 1 meter above the floor, and away from walls and other potential obstacles (and naturally, in line of sight with each other) the results might have been somewhat influenced by unwanted signal reflections, typical of indoors environments (that might have either improved or

Figure 5.3: Reflection coefficient measured at the input of the sensor node (represented by filled circles) in function of power, plotted against the predicted magnitude and phase curves (dotted lines with borderless circles) and assuming a test frequency of 866.6MHz. Solid lines were used to mark points at -8.64dBm.

degraded them, in case of constructive or destructive interferences, respectively). Either way, the distance effectively achieved in practice is below the value of 5.4 meters predicted by the Friis formula, but still, a very reasonable result (calculations were carried out assuming the developed antenna to have a gain of the traditional half-wave dipole and a received power equal to the turn on threshold previously measured for the sensor node).

5.3 Antenna measurements and efficiency considerations

The fact that the wireless range estimated for the developed sensor node was reasonably close to the value that was obtained in practice suggests that the efficiency of the developed antenna is in principle significantly high. However, and since the parameter in question was not measured, this assumption could not be confirmed.

In testing, the antenna was experimentally characterized in terms of its reflection coefficient and radiation patterns in two different planes.

5.3.1 Reflection coefficient

In this test, the antenna was connected to a network analyser (no particular settings needed to be configured) for measuring its reflection coefficient in the same frequency range used in previous tests, that is, from 766.6 to 966.6MHz. Figure 5.4 presents the corresponding results, the results yielded by the uncompensated antenna under the same conditions (plotted solely for comparison purposes) and additionally, the curve that was initially expected.

As it is depicted in the figure, while the uncompensated antenna missed the desired resonant frequency by a significant margin, the compensated antenna did not.

Figure 5.4: Reflection coefficients of the first and second antennas (plotted with empty and filled circles respectively) in function of frequency, compared to the curve (dotted line with borderless circles) expected for the first antenna, magnitude only. The highlighted data value corresponds to a frequency of 866.6MHz.

5.3.2 Gain and radiation diagram

In this test (realized inside an anechoic chamber), an antenna playing the role of a distant (that is, far-field) power source was setup to radiate power to the developed antenna, with the latter being mounted on top of a rotary platform and connected to a power detector. By observing the power displayed on the power detector at different rotation angles, a two-dimensional radiation pattern corresponding to the plane perpendicular to the rotation axis (vertically oriented in this case) could be measured. From top to bottom, the results plotted in figure 5.5 correspond to the tests in which the antenna was horizontally or vertically mounted, respectively. In both cases, the origin of the angles corresponds to the antennas being oriented towards each other. It should be mentioned that before each sweep the antennas were aligned in terms of maximum radiation pattern and polarization.

While the minimums typical of dipole-based antennas were indeed measured at the expected angles, in the top plot there is an unexpected behavior at angles near the origin, and in the bottom plot, deviations to the predicted curve are quite visible. Regarding these deviations, it is possible that the support to which the antenna was attached could have been actually absorbing some of the power that otherwise would go to the antenna. Other than this, no other reasons that would explain these inconsistencies were found.

Figure 5.5: Normalized radiation patterns of the second antenna, measured at the planes that contain the direction of maximum radiation plus either the electric or magnetic field vectors (upper and lower plots respectively, both indicated by filled circles), plotted against expected curves (dotted lines with borderless circles), considering a frequency of 866.6MHz.

Chapter 6

Conclusions and future research

In this chapter, the first section summarily highlights the most relevant features of the developed sensor node. The second section describes important conclusions extracted from this work, points out various research directions that can be followed in order to improve the current system, and ends by providing a brief glimpse about how much important electromagnetic energy harvesting and wireless sensor nodes will most likely be in the future.

6.1 Introduction

It is somewhat ironic that the same hardware component which enabled (and in most cases still enables) sensor nodes to operate wirelessly is now limiting their natural evolution towards smaller sizes and longer lifetimes. In order to overcome this limitation, a new class of sensor nodes that contain no batteries and harvest all their energy from incoming radio waves is now emerging, and it is to this class that the sensor node developed in this work belongs. In addition to being fully passive, the proposed sensor node offers the flexibility associated with having a general purpose microcontroller and a relatively high number of externally available pins to connect to sensors or other devices. In terms of harvesting performance, the results obtained in practice matched those generated by simulations quite reasonably.

When attached to its custom-designed dipole antenna, the sensor node achieves a maximum communication range that is not only comparable to that of conventional RFID passive tags, but also fairly close to what would be theoretically expected. Nevertheless, the device is still largely underexplored and its range can be extended as its hardware (which includes the antenna, as the latter plays an important role in energy harvesting) and software are further improved.

6.2 Main conclusions

Despite being fairly small in size, the developed sensor node involved a considerable amount of trade-offs in its design, with regards to communication rates, cost, latency, operating frequency, size and others, but the most important trade-off of all was definitely between functionality and energy consumption. While the sensor node was not expected in any way to be able to perform demanding tasks (such as demodulating complex waveforms), it was required to execute simple

functions with a high level of efficiency, so that they would not to affect its energy balance, which would result in a significant reduction in range. Additionally, it is very important to mention the need of an accurate simulation model as the sensor node is pushed towards its limits, otherwise optimizations become irrelevant (which was precisely what happened on its first iteration). In addition to hardware, the software that runs inside the controller of the sensor node is of critical importance to ensure the proper conditions for energy harvesting, and more importantly, to avoid unnecessary energy expenditures by guarantee that each and every hardware module is switched on only when it is absolutely required to.

Beyond optimizations, and considering that the power consumption of hardware components continues to drop, and that controllers continue to increase their number of instructions per unit of energy, the lifetime and processing capabilities of wireless sensor nodes (and also range in case of passive nodes) also tend to naturally increase over time.

6.2.1 Further work

Without modifying any hardware within the developed sensor node, the most immediate way to increase its performance is by improving energy management, for instance, by converting some parts of the software to assembly, in order to achieve a fine grained control over communication and sensing (and also by taking advantage of the currently unused resources in some way, such as the crystal oscillator for instance). Additionally, not only sensors but several other devices can be connected to the sensor node as desired (as long as they do not consume too much energy, or else include an energy source of their own). Regarding this subject, it would be interesting to connect a super capacitor to the sensor node in order to provide the latter the ability to operate continuously without requiring its power source to broadcast power in a continuous fashion. It is actually rather easy to imagine applications where sensors are supplied with power just until they accumulate enough energy to sustain themselves through the rest of the day, for instance.

Re-designing certain hardware blocks within the sensor node effectively opens a new array of possibilities that can be explored. For instance, using smaller component packages and mounting them on both sides of the sensor node would yield a relevant reduction in size. Furthermore, it may be worth investigating whether or not its ground plane can be eliminated (and other types of microstrip lines used) without significantly affecting performance. Additionally, and in order to enable further reductions in size, its antenna system would also have to be modified and perhaps integrated into the same substrate board onto which the node is built. While this approach would avoid losing power in connectors and would (in principle) result in a simpler impedance matching, the advantages associated with having a standard 50Ω interface in the node would be obviously lost. In terms of circuitry, the proposed sensor node would benefit from having a faster wake-up system and, more importantly, a programmable supply voltage supervisor (that would allow its threshold voltage to be set dynamically).

Major modifications, such as integrating hardware designed to harvest other forms of energy into the sensor node or altering its operating frequency (or perhaps using distinct frequencies for data and power, if the advantages of such approach outweigh the added complexity in terms of antenna design) may also be taken into account. Finally, it should be noted that optimizing the sensor node can actually go beyond the device itself. Although only sinusoidal signals were used

to evaluate harvesting performance, different signals may stimulate the energy harvester more efficiently (and as shown in [47], it seems that Power-Optimized Waveforms (POWs) are indeed capable of doing that).

6.2.2 Concluding remarks on electromagnetic energy harvesting

Electromagnetic energy harvesting is undoubtedly not a solution for all energy needs, but it does stand as a concept that will increasingly enhance or even replace batteries in many sensor nodes as circuit techniques and transducer efficiencies improve, allowing them to be long-lived and yet small enough to go unnoticed while performing their tasks in the huge amount of applications in which they will most likely participate in the future.

It is not unreasonable to expect that in a couple of decades (perhaps even less) the world will be covered with sensors wirelessly connected to the internet. At that time, the internet itself will change and become a large-scale physical network.

Appendix A

Source code

In this appendix, the first and second sections contain the two programs that were developed for the proposed sensor node. Both programs are written in C and make use of specific libraries provided by the CCS development environment.

A.1 Deep sleep mode

```
/* WORKS WITH KHALIPHA 1.1 */

/* P1 : 1 1 1 1 1 1(TXB) 1 0(RXB) */
/* P2 : 0(XOUT) 0(XIN) 1(RXE) 1 1 0(SUP) 1 1 */
/* P3 : 1 1 1 1 1 0(SCL) 0(SDA) 1 */

#include "msp430f2132.h"

void main(void)
{
    WDTCTL = WDTPW | WDTHOLD;

    /* WATCHDOG TIMER STOPPED */

    P1DIR = 0xFE;
    P1OUT &= ~0xFE;
    P2DIR = 0x3B;
    P2OUT &= ~0x3B;
    P3DIR = 0xF9;
    P3OUT &= ~0xF9;

    /* OUTPUT I/O PINS SELECTED */

    /* TURN OFF */

    __bis_SR_register(SCG1 + SCG0 + OSCOFF + CPUOFF);
}
```

A.2 Proposed program for communication and control

```
/* WORKS WITH KHALIPHA 1.1 */

/* P1 : 1 1 1 1 1 1(TXB) 1 0(RXB) */
/* P2 : 0(XOUT) 0(XIN) 1(RXE) 1 1 0(SUP) 1 1 */
/* P3 : 1 1 1 1 1 0(SCL) 0(SDA) 1 */

#include "msp430f2132.h"

#define RXB BIT0;
#define TXB BIT2;
#define RXE BIT5;
#define SUP BIT2;

volatile unsigned int SEQ00, IDX, ZZ = 0;

void main(void)
{
    WDTCTL = WDTPW + WDTHOLD;

    /* WATCHDOG TIMER STOPPED */

    P2IES &= ~SUP;
    P2IFG &= ~SUP;
    P2IE |= SUP;

    /* TURN OFF AND WAIT FOR THE SUPERVISOR */

    __bis_SR_register(SCG1 + SCG0 + OSCOFF + CPUOFF + GIE);

    BCCTL1 = XT2OFF + RSEL3 + RSEL2 + RSEL0;

    /* MAXIMUM OPERATING FREQUENCY SET */

    for (;;)
    {
        /* TURN ON THE RECEIVER */

        P2OUT |= RXE;
        P1IES &= ~RXB;
        P1IFG &= ~RXB;
        P1IE |= RXB;

        TACCR0 = 0x14A;

        /* TIMEOUT SET */
    }
}
```

```

TACCTL0 |= CCIE;

do
{
    /* TURN OFF AND WAIT FOR THE WAKE UP BIT */

    __bis_SR_register(SCG1 + SCG0 + CPUOFF);

    /* THE CLOCK SOURCE IS NOW ON */

    IDX = 0x80;

    /* UP TO 16 BITS WITH IDX = 0x8000 (8 BITS SELECTED) */

    TAR = 0;
    TACTL |= MC_1;

    /* WAIT FOR THE REST OF THE BITS */

    __bis_SR_register(CPUOFF);

    /* IF THE SEQUENCE WAS INCOMPLETE GO BACK TO SLEEP */
}
while (IDX != 0);

TACCTL0 &= ~CCIE;
TACTL &= ~MC_1;

/* TURN OFF THE RECEIVER */

P1IE &= ~RXB;
P2OUT &= ~RXE;

SEQ00 = ~SEQ00;

/* WAIT FIRST THEN TRANSMIT */

TACCR1 = 0x258;

TACCTL1 &= ~CCIFG;
TACCTL1 |= CCIE;

IDX = 0x80;

/* UP TO 16 BITS WITH IDX = 0x8000 (8 BITS SELECTED) */

TAR = 0;
TACTL |= MC_2;

```

```

/* TRANSMIT */

P1OUT &= ~TXB;
__bic_SR_register(CPUOFF);
P1OUT &= ~TXB;

TACCTL1 &= ~CCIE;
}
}

#pragma vector = PORT1_VECTOR
__interrupt void Port_1(void)
{
    P1IFG &= ~RXB;

    if (TACTL & MC_1)
    {
        TACTL &= ~MC_1;

/* DECODE A BIT (PIE ENCODING) */

        if (TAR > 0xC0) {SEQ00 |= IDX;} else {SEQ00 &= ~IDX;}
        TAR = 0;
        IDX >>= 1;

        if (IDX == 0)
        {
            /* ALL BITS RECEIVED */

            __bic_SR_register_on_exit(CPUOFF);
        }
        else {TACTL |= MC_1;}
    }

/* ENABLE CLOCK SOURCES */

    else {__bic_SR_register_on_exit(SCG1 + SCG0 + CPUOFF);}
}

#pragma vector = TIMER0_A0_VECTOR
__interrupt void Timer_A0(void)
{
    TACTL &= ~MC_1;

    __bic_SR_register_on_exit(CPUOFF);
}

```



```

#pragma vector = TIMER0_A1_VECTOR
__interrupt void Timer_A1(void)
{
    TACCTL1 &= ~CCIFG;
    TAR = 0;

    if (ZZ != 0)
    {
        P1OUT &= ~TXB;

        if (IDX == 0)
        {
            /* ALL BITS SENT */

            TACTL &= ~MC_2;
            __bic_SR_register_on_exit(CPUOFF);
        }

        /* GENERATE THE FIRST PART OF A BIT */

        if (SEQ00 & IDX) {TACCR1 = 0xC0;} else {TACCR1 = 0x60;}

        IDX >>= 1;
    }

    /* GENERATE THE SECOND PART OF A BIT */

    else {P1OUT |= TXB; TACCR1 = 0x30;}

    ZZ ^= 1;
}

#pragma vector = PORT2_VECTOR
__interrupt void Port_2(void)
{
    P2IE &= ~SUP;

    P1DIR = 0xFE;
    P1OUT &= ~0xFE;
    P2DIR = 0x3B;
    P2OUT &= ~0x3B;
    P3DIR = 0xF9;
    P3OUT &= ~0xF9;

    /* OUTPUT I/O PINS SELECTED */

    /* SELECT CLOCK SOURCE */

```

```
TACTL |= TASSEL_2;
__bic_SR_register_on_exit(SCG0 + CPUOFF);
}
```

Bibliography

- [1] I.F. Akyildiz, Weilian Su, Y. Sankarasubramaniam, and E. Cayirci. A survey on sensor networks. *Communications Magazine, IEEE*, 40(8):102 – 114, aug 2002.
- [2] R. Want, K.I. Farkas, and C. Narayanaswami. Guest editors' introduction: Energy harvesting and conservation. *Pervasive Computing, IEEE*, 4(1):14 – 17, jan. 2005.
- [3] Daniel M. Dobkin. *The RF in RFID: Passive UHF RFID in Practice*. Newnes, Newton, MA, USA, 2007.
- [4] Daniel M. Dobkin and Titus Wandinger. A radio-oriented introduction to radio frequency identification. *High Frequency Electronics*, June 2005.
- [5] H. Stockman. Communication by means of reflected power. *Proceedings of the IRE*, 36(10):1196 – 1204, oct. 1948.
- [6] V. Chawla and Dong Sam Ha. An overview of passive rfid. *Communications Magazine, IEEE*, 45(9):11 – 17, september 2007.
- [7] J. Landt. The history of rfid. *Potentials, IEEE*, 24(4):8 – 11, oct. 2005.
- [8] J. Banks. *RFID Applied*. John Wiley & Sons, Inc., New York, NY, USA, 2007.
- [9] Klaus Finkenzeller. *RFID Handbook: Fundamentals and Applications in Contactless Smart Cards and Identification*. John Wiley & Sons, Inc., New York, NY, USA, 2003.
- [10] Harvey Lehpamer. *RFID Design Principles*. Artech House, Inc., Norwood, MA, USA, 2007.
- [11] J.-P. Curty, M. Declercq, C. Dehollain, and N. Joehl. *Design and Optimization of passive UHF RFID Systems*. Springer, New York, 2006.
- [12] Chee-Yee Chong and S.P. Kumar. Sensor networks: evolution, opportunities, and challenges. *Proceedings of the IEEE*, 91(8):1247 – 1256, aug. 2003.
- [13] Yang Yu. *Information Processing and Routing in Wireless Sensor Networks*. World Scientific Publishing Co., Inc., River Edge, NJ, USA, 2007.
- [14] Kazem Sohraby, Daniel Minoli, and Taieb Znati. *Wireless Sensor Networks: Technology, Protocols, and Applications*. Wiley-Interscience, 2007.
- [15] Holger Karl and Andreas Willig. *Protocols and Architectures for Wireless Sensor Networks*. John Wiley & Sons, 2005.
- [16] I. Khemapech, I. Duncan, , and A. Miller. A survey of wireless sensor networks technology. In *Proceedings of the 6th Annual PostGraduate Symposium on the Convergence of Telecommunications, Networking and Broadcasting*, Liverpool, June 2005. EPSRC.

- [17] Shan Liang, Yunjian Tang, and Qin Zhu. Passive wake-up scheme for wireless sensor networks. In *ICICIC '07: Proceedings of the Second International Conference on Innovative Computing, Information and Control*, page 507, Washington, DC, USA, 2007. IEEE Computer Society.
- [18] Chris Townsend and Steven Arms. *Sensor Technology Handbook*, chapter 22, pages 439–450. Elsevier, 2005.
- [19] Zhi Ning Chen. *Antennas for Portable Devices*. John Wiley & Sons, Inc., New York, NY, USA, 2007.
- [20] Constantine A. Balanis. *Modern Antenna Handbook*. Wiley-Interscience, New York, NY, USA, 2008.
- [21] K. Romer and F. Mattern. The design space of wireless sensor networks. *Wireless Communications, IEEE*, 11(6):54 – 61, dec. 2004.
- [22] Amiya Nayak Hai Liu, Miodrag Bolic and Ivan Stojmenovi. *Encyclopedia On Ad Hoc And Ubiquitous Computing: Theory and Design of Wireless Ad Hoc, Sensor, and Mesh Networks*, chapter 13, pages 319–348. World Scientific Publishing Company, 2009.
- [23] Loc Ho, Melody Moh, Zachary Walker, Takeo Hamada, and Ching-Fong Su. A prototype on rfid and sensor networks for elder healthcare: progress report. In *E-WIND '05: Proceedings of the 2005 ACM SIGCOMM workshop on Experimental approaches to wireless network design and analysis*, pages 70–75, New York, NY, USA, 2005. ACM.
- [24] O.B. Akan, M.T. Isik, and B. Baykal. Wireless passive sensor networks. *Communications Magazine, IEEE*, 47(8):92 –99, august 2009.
- [25] S. Roundy, D. Steingart, L. Frechette, P. Wright, and J. Rabaey. Power sources for wireless sensor networks. *1st European Workshop on Wireless Sensor Networks*, pages 19 – 21, January 2004.
- [26] A. Bereketli and O.B. Akan. Communication coverage in wireless passive sensor networks. *Communications Letters, IEEE*, 13(2):133 –135, feb. 2009.
- [27] M.T. Isik and O.B. Akan. Padre: Modulated backscattering-based passive data retrieval in wireless sensor networks. pages 1 –6, apr. 2009.
- [28] R. Want. Enabling ubiquitous sensing with rfid. *Computer*, 37(4):84 – 86, april 2004.
- [29] Alanson Sample Michael Buettner, Ben Greenstein and Joshua R. Smith. Revisiting smart dust with rfid sensor networks. In *Seventh ACM Workshop on Hot Topics in Networks (HotNets-VII)*, October 2008.
- [30] Joshua R. Smith, Bing Jiang, Sumit Roy, Matthai Philipose, Kishore Sundara-Rajan, and Alexander Mamishev. Id modulation: Embedding sensor data in an rfid timeseries. In *Information Hiding*, pages 234–246, 2005.
- [31] D.J. Yeager, P.S. Powledge, R. Prasad, D. Wetherall, and J.R. Smith. Wirelessly-charged uhf tags for sensor data collection. pages 320 –327, apr. 2008.
- [32] A.P. Sample, D.J. Yeager, P.S. Powledge, A.V. Mamishev, and J.R. Smith. Design of an rfid-based battery-free programmable sensing platform. *Instrumentation and Measurement, IEEE Transactions on*, 57(11):2608 –2615, nov. 2008.
- [33] A.P. Sample, D.J. Yeager, P.S. Powledge, and J.R. Smith. Design of a passively-powered, programmable sensing platform for uhf rfid systems. pages 149 –156, mar. 2007.
- [34] Joshua R. Smith, Alanson Sample, Pauline Powledge, Alexander Mamishev, and Sumit Roy. A wirelessly powered platform for sensing and computation. *Proceedings of Ubicomp 2006: 8th International Conference on Ubiquitous Computing*, pages 495–506, September 2006.

- [35] Joshua R. Smith Kevin Fu Hee-Jin Chae, Daniel J. Yeager. Maximalist cryptography and computation on the wisp uhf rfid tag. *Proceedings of the Conference on RFID Security*, July 2007.
- [36] S. Serkan Basat, S. Bhattacharya, Li Yang, A. Rida, M.M. Tentzeris, and J. Laskar. Design of a novel high-efficiency uhf rfid antenna on flexible lcp substrate with high read-range capability. pages 1031–1034, jul. 2006.
- [37] Alanson Sample and Joshua R. Smith. Experimental results with two wireless power transfer systems. In *RWS'09: Proceedings of the 4th international conference on Radio and wireless symposium*, pages 16–18, Piscataway, NJ, USA, 2009. IEEE Press.
- [38] Texas Instruments. *MSP430F21x2 Mixed Signal Microcontroller*, rev. g edition, December 2009.
- [39] Texas Instruments. *MSP430x2xx Family User's Guide*, 2008.
- [40] John H. Davies. *MSP430 Microcontroller Basics*. Newnes, Newton, MA, USA, 2008.
- [41] Feng Pan and Tapan Sammadar. *Charge Pump Circuit Design*. McGraw-Hill, New York, NY, USA, 2006.
- [42] Peter Spevak and Peter Forstner. Msp430 32-khz crystal oscillators. *MSP430 Applications SLAA322*, Texas Instruments, August 2006.
- [43] Agilent Technologies. *HSMS-2850 Series - Surface Mount Zero Bias Schottky Detector Diodes*, 1999.
- [44] Hewlett-Packard. *Application Note 1124 - Linear Models for Diode Surface Mount Packages*, 1997.
- [45] Coilcraft. *Key Parameters for Selecting RF Inductors*, 2008.
- [46] Rogers Corporation. *RO4000 Series High Frequency Circuit Materials*, 2007.
- [47] M.S. Trotter, J.D. Griffin, and G.D. Durgin. Power-optimized waveforms for improving the range and reliability of rfid systems. In *RFID, 2009 IEEE International Conference on*, pages 80–87, 2009.