

Joshua R. Smith *Editor*

Wirelessly Powered Sensor Networks and Computational RFID

 Springer

Wirelessly Powered Sensor Networks and Computational RFID

Joshua R. Smith

Editor

Wirelessly Powered Sensor Networks and Computational RFID



Springer

Editor

Joshua R. Smith
Department of Computer Science
and Engineering
Department of Electrical Engineering
University of Washington
Seattle, Washington, USA

ISBN 978-1-4419-6165-5 ISBN 978-1-4419-6166-2 (eBook)

DOI 10.1007/978-1-4419-6166-2

Springer New York Heidelberg Dordrecht London

Library of Congress Control Number: 2012954295

© Springer Science+Business Media New York 2013

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed. Exempted from this legal reservation are brief excerpts in connection with reviews or scholarly analysis or material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work. Duplication of this publication or parts thereof is permitted only under the provisions of the Copyright Law of the Publisher's location, in its current version, and permission for use must always be obtained from Springer. Permissions for use may be obtained through RightsLink at the Copyright Clearance Center. Violations are liable to prosecution under the respective Copyright Law.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

*For Mom, who helped me get a high school
nerd's dream job, programming at NASA, and
For Maggie, who moved to Seattle*

Preface

Joshua R. Smith

1 Linking Bits and Atoms

Sensors are points of contact between the material world of atoms, mass, and energy and the seemingly immaterial world of information, computation, and cognition. Linking these two domains more tightly yields all sorts of practical benefits, such as improved input devices for computers, more effective medical devices (implanted or worn), more precise agricultural operations, better monitored buildings or bridges, more secure payment systems, and more reliable sensor–actuator control systems. There are many settings in which tighter coupling between digital and physical planes can enhance safety, security, performance, and reliability.

1.1 The Problem of Powering Sensors

In recent years, powering sensor systems has emerged as a key problem that is inhibiting their long-term, wide-scale adoption. Batteries need to be replaced or recharged, which can render impractical many large-scale systems that might otherwise be deployed. Batteries also increase the size, weight, and cost of sensor systems and must be disposed properly at the end of a system’s life. Furthermore, rechargeable batteries can only sustain a finite number of charge–discharge cycles, which limits the lifetime of battery-powered systems even when it is feasible to recharge them.

Joshua R. Smith
Department of Computer Science and Engineering, Department of Electrical Engineering,
University of Washington, Seattle, WA, USA
e-mail: jrs@cs.washington.edu

1.2 A Solution: Power Harvesting

Power harvesting promises to enable sensing and computing systems that can operate perpetually. Solar power is the most mature energy-harvesting technology, but it has drawbacks. For nighttime operation, systems require enough energy storage to operate for many hours without light, which typically means a battery. Solar cells require a direct optical line of sight from the energy-transducing element (a PN junction, typically on glass) to the sun, a significant constraint on the design of solar-powered devices. New forms of harvesting are needed to expand the design space of energy-harvesting systems.

1.2.1 RF Power Harvesting

RF harvesting has a number of desirable properties. Radio waves can propagate through materials such as plastic, wood, gypsum board, concrete, and (to some extent) animal tissue. This allows RF power-harvesting systems to be embedded more deeply and permanently than systems that rely on solar harvesting. For example, an RF harvesting system can be embedded in a wall or roof, in concrete, or implanted inside the human body; the energy-transducing element (the antenna, for RF harvesting) does not have to be on the host object's exterior surface. By eliminating the batteries needed for nighttime operation, the weight of RF harvesting systems can be made very low. One chapter in this book presents an RF-powered device that is light enough to be flown by a living moth.

2 Wirelessly Powered Sensor Networks and Computational RFID

The focus of this book is the design and application of systems that harvest energy from RF signals, either deliberately transmitted or ambient. RF harvesting is a rapidly emerging area of research, and this is the first book to survey wirelessly powered sensor systems. The book maps the space of RF-powered sensing, computing, and actuating systems, through a collection of specific research examples. Radio frequency identification (RFID) tags are the first widely deployed RF-powered sensor systems, and many of the papers build on RFID technology. The research in this book includes RF power harvesting and transmission, computational RFID, wireless networking, and sensing applications.

2.1 Structure of the Book

The chapters of this book fall into five clusters: introduction, hardware platforms, communication, wireless power transfer, systems and applications, and security.

2.1.1 Introduction

The introductory section includes two chapters. The first chapter, “Range Scaling of Wirelessly Powered Sensor Systems,” is an examination of the macro-level scaling trends, in particular ongoing exponential improvements in the energy efficiency of computation, which have enabled the research in this book, and will determine its future as well. The second chapter, “History of the WISP Program,” presents the history of what we believe to be the first far-field RF-powered computational RFID platform (i.e., the first computational RFID to operate at UHF frequencies). This chapter provides historical context and makes connections among many of the chapters in the book.

2.1.2 Hardware Platforms

The hardware section describes the design of three computational RFID platforms. The chapter entitled “The Wireless Identification and Sensing Platform” describes the design of the original WISP. The SOCWISP (for “System on Chip WISP,” a single-chip implementation of every portion of the WISP functionality except for the general purpose computing capability) is described in the chapter “A 9 μ A, Addressable Gen 2 Sensor Tag for BioSignal Acquisition” by Daniel Yeager, Fan Zhang, Azin Zarrasvand, Nicole George, Thomas Daniel, and Brian Otis. The next chapter describes a similar platform (with sensors that are powered and read by UHF signals) produced by Spanish start-up FARSENS; the paper by Roc Berenguer, Ivan Rebollo, Ibon Zalvide, and Inaki Fernandez is entitled “Battery-Less Wireless Sensors Based on Low Power UHF RFID Tags.”

2.1.3 Communication and Tools

This section has four chapters covering a range of topics related to communication as well as tools. In the chapter “Passive RFID-Based Wake-Up Radios for Wireless Sensor Networks” He Ba, Ilker Demirkol, and Wendi Heinzelman show how to use computational RFID tags to wake up larger battery-powered sensor nodes, thereby saving power for the larger node. The chapter “BAT: Backscatter Anything-to-Tag Communication” is a very exciting example of overlaying a more sophisticated communication protocol on top of the base EPC Class 1 Generation 2 (C1G2) protocol. Molina-Markham, Shane Clark, Benjamin Ransford, and Kevin Fu present a protocol in which tags can relay messages to other tags, through the reader. These first two papers were both enabled by the WISP challenge, a program described in the chapter “History of the WISP Program”. In the chapter “Implementing the Gen 2 MAC on the Intel-UW WISP” Michael Buettner and David Wetherall present an implementation for the WISP of the EPC Class 1 Generation 2 MAC (Medium Access Control scheme). The final chapter in this section is on debugging WISPs. Combining the lack of interface with power constraints, debugging code on RF-

powered computers, can be tricky; the chapter “WISP Monitoring and Debugging” by Richa Prasad, Michael Buettner, Ben Greenstein, and David Wetherall presents tools that aim to simplify debugging of RF-powered computing systems.

2.1.4 Cryptography and Security for Computational RFID

One of the surprises when we began giving WISPs away was the degree of interest in the platform from the security community. Apparently there had been no way to implement security algorithms on a passive UHF RFID tag before WISP. The chapter “Maximalist Cryptography and Computation on the WISP UHF RFID Tag,” by Hee-Jin Chae, Mastooreh Salajegheh, Daniel J. Yeager, Joshua R. Smith, and Kevin Fu, presents what we believe was the first implementation of a strong cryptographic algorithm on a UHF-powered tag. The chapter “Security Enhanced WISPs: Implementation Challenges,” by Alexander Szekely, Michael Hoffer, Robert Stogbuchner, and Manfred Aigner, presents an implementation of AES on the WISP. Both of these papers arose from the WISP challenge program.

2.1.5 Wireless Power Transfer Beyond RFID

This section covers alternate approaches to wireless power transfer, not the basic techniques used by the WISP. The section starts with the chapter “Power Optimized Waveforms that Enhance the Range of Energy Harvesting Sensors” by Matthew Trotter and Gregory Durgin. Power optimized waveforms are RF signals that are designed specifically to facilitate efficient harvesting (while not violating any regulations). This chapter was a result of the WISP challenge program. The next two chapters are projects that are described briefly in the “History of the WISP program” chapter as outgrowths of the WISP project, but not directly part of it. The chapter “Wireless Ambient Radio Power,” by Alanson Sample, Aaron Parks, Scott Southwood, and Joshua R. Smith, measures environmental sensing data and then transmits it short distances using a 2.4 GHz radio. The entire thing is powered by ambient RF power harvested from a TV tower. The chapter “Powering a VAD Using the Portable FREED System,” by Benjamin Waters, Kara Kagi, Jordan Reed, Alanson Sample, Pramod Bonde, and Joshua R. Smith, is an example of our high-power (tens of watts) medium-range wireless power transfer technology. The application to implanted heart pumps could improve quality and length of life for heart failure patients.

2.1.6 Systems and Applications

This section presents two very unusual applications of WISP. The first, in Chap. 13, describes the use of WISPs for oceanographic temperature sensing, as part of an undersea neutrino telescope. This project was another that was made possible

by the WISP challenge. The key benefit of wirelessly powered sensing in this application is avoiding the need to penetrate a pressure vessel that must withstand high pressure. The benefit of using wireless power to avoid breaching a significant boundary with a cable in this undersea application is very similar in some respects to the biomedical application, in which wireless power allows us to eliminate a transcutaneous cable. The chapter “RFID-Vox: A Tribute to Leon Theremin,” by Pavel Nikitin, Aaron Parks, and Joshua R. Smith, presents the life story of the inventor, and how he created an early version of RFID technology. The article also uses WISP to present an RFID implementation of another of his famous inventions, the theremin electronic musical instrument.

2.2 Intended Audience

The book should be useful to anyone with an interest in deepening the links between the physical world and information processing systems. Sensor networking researchers can explore perpetually operating battery-free sensor nodes. Researchers and practitioners in the area of RFID will gain an understanding of the opportunities that arise from adding sensing and computing to RFID tags. Security researchers and practitioners will be exposed to the security challenges and opportunities (both physical and digital) presented by sensor-enhanced, computational RFID tags. Power-harvesting researchers will find challenging questions and issues. Engineers interested in using or designing wireless power systems will find valuable material. Researchers in the areas of ubiquitous computing and human–computer interaction (HCI) can learn about the possibilities for wirelessly powering sensors and input devices. Researchers interested in communication protocols will encounter new research questions, for both backscatter communication protocols, and low overhead “burst networks” that harvest small amounts of power over long periods of time, communicate a very short data burst, and then return to power harvesting.

3 Outlook

The types of RF-powered systems described in this book today seem somewhat exotic and are still not always robust. Because wireless power is such a novel capability, it is exciting to be able to use RF power to operate even small workloads at short range. The chapter “Range Scaling of Wirelessly Powered Sensor Systems” of this book argues that the range at which any particular workload can be wirelessly powered is increasing exponentially. If this trend continues, then capabilities that barely work today will become robust and operate at long range tomorrow. Further research progress and continued energy efficiency scaling could transform RF-powered sensing and computing systems from the novelties they are today to mainstream, essential technologies that will be widely relied upon tomorrow.

4 Acknowledgments

I thank Intel, in particular the management of Intel Research Seattle and Intel Labs, for being enlightened enough to allow me to open source the design of the WISP.

My research group at the University of Washington, the Sensor Systems Lab, has been funded by the Intel Science and Technology Center for Pervasive Computing, a Google Faculty Research Award, the University of Washington Commercialization Gap Fund, the Center for Sensorimotor Neural Engineering, via National Science Foundation award number EEC-1028725, and by NSF award ECCS-0824265, “Realizing the internet of things with RFID sensor networks.”

I thank my colleagues at Intel Labs Seattle and University of Washington, in particular James Landay, David Wetherall, and Dieter Fox for their excellent advice and enjoyable collaborations. It has been a pleasure to work with so many excellent students; many of them are co-authors on chapters in this book.

I thank my family (including my parents Anthony and Mary Smith and brother Ethan Smith) for their encouragement over many years. Thanks to my wife, Maggie Orth, and kids Lily Orth-Smith and Dorothea Orth-Smith, for their patience.

Contents

Part I Introduction

Range Scaling of Wirelessly Powered Sensor Systems	3
Joshua R. Smith	
History of the WISP Program	13
Joshua R. Smith	

Part II Hardware Platforms

The Wireless Identification and Sensing Platform	33
Alanson P. Sample and Joshua R. Smith	
SOCWISP: A 9 μA, Addressable Gen2 Sensor Tag for Biosignal Acquisition	57
Daniel Yeager, Fan Zhang, Azin Zarrasvand, Nicole George, Thomas Daniel, and Brian Otis	
Battery-less Wireless Sensors Based on Low Power UHF RFID Tags	79
Roc Berenguer, Iván Rebollo, Ibon Zalbide, and Iñaki Fernández	

Part III Communication and Tools

Passive RFID-Based Wake-Up Radios for Wireless Sensor Networks	113
He Ba, Jeff Parvin, Luis Soto, Ilker Demirkol, and Wendi Heinzelman	
BAT: Backscatter Anything-to-Tag Communication	131
Andrés Molina–Markham, Shane S. Clark, Benjamin Ransford, and Kevin Fu	
Implementing the Gen 2 MAC on the Intel-UW WISP	143
Michael Buettner and David Wetherall	

WISP Monitoring and Debugging 157
 Richa Prasad, Michael Buettner, Ben Greenstein,
 and David Wetherall

**Part IV Cryptography and Security for Computational
 RFID**

**Maximalist Cryptography and Computation on the WISP
 UHF RFID Tag** 175
 Hee-Jin Chae, Mastrooreh Salajegheh, Daniel J. Yeager,
 Joshua R. Smith, and Kevin Fu

Security Enhanced WISPs: Implementation Challenges 189
 Alexander Szekely, Michael Höfler, Robert Stögbuchner,
 and Manfred Aigner

Part V Wireless Power Beyond RFID

**Power Optimized Waveforms that Enhance the Range of
 Energy-Harvesting Sensors** 207
 Matthew S. Trotter and Gregory D. Durgin

Wireless Ambient Radio Power..... 223
 Alanson P. Sample, Aaron N. Parks, Scott Southwood,
 and Joshua R. Smith

**A Portable Transmitter for Wirelessly Powering a Ventricular
 Assist Device Using the Free-Range Resonant Electrical
 Energy Delivery (FREE-D) System** 235
 Benjamin H. Waters, Jordan T. Reed, Kara R. Kagi,
 Alanson P. Sample, Pramod Bonde, and Joshua R. Smith

Part VI Systems and Applications

**PORFIDO: Using Neutrino Telescopes
 and RFID to Gather Oceanographic Data** 251
 Orlando Ciaffoni, Marco Cordelli, Roberto Habel, Agnese Martini,
 and Luciano Trasatti

RFID-Vox: A Tribute to Leon Theremin 259
 Pavel V. Nikitin, Aaron Parks, and Joshua R. Smith

Index 269

Part I
Introduction

Range Scaling of Wirelessly Powered Sensor Systems

Joshua R. Smith

1 Motivation

This volume describes a variety of RF-powered sensor systems, including the wireless identification and sensing platform (WISP, described in Chap. 3 [11]; see also [12, 13]) and wireless ambient radio power system (WARP, described in Chap. 11 [10]; see also [9]). WISP harvests on the order of $100\mu\text{W}$ (100 microwatts) from a 1 W RF transmitter located 2–4 m away; it includes sensors and a low-power backscatter radio. The WARP system harvests about the same amount of power from a 1 MW (1 megawatt) TV tower that can be many kilometers away. The question motivating this chapter is: why did it become possible to build RF-powered systems like these in the mid-2000s? Could they have been built in the 1980s? In the 1990s? This chapter attempts to answer these questions and also to extrapolate into the future: if the observed technology scaling trends continue, what will be possible in 2020, 2030, or 2040?

1.1 Introduction

The density of transistors on integrated circuits has been growing exponentially since the integrated circuit was invented, as famously observed by Gordon Moore [7]. It has recently been pointed out that the energy efficiency of computation has also been increasing exponentially [5, 6]. Here, we argue that energy efficiency scaling implies range scaling for wirelessly powered systems. In other words, the range at which any given computational workload can be wirelessly powered is also

J.R. Smith (✉)

Department of Computer Science and Engineering, Department of Electrical Engineering, University of Washington, Seattle, WA, USA
e-mail: jrs@cs.washington.edu

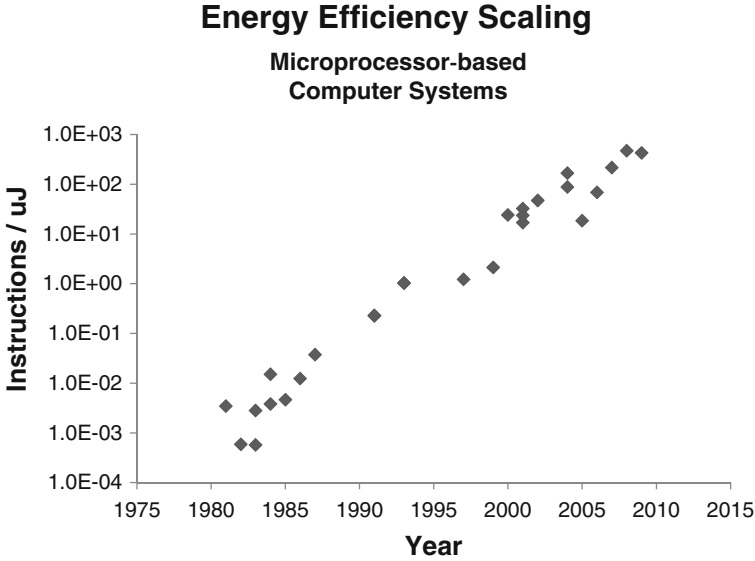


Fig. 1 Scaling of energy efficiency for microprocessor-based computer systems. Data from [5,6]

increasing exponentially, but, due to attenuation of the RF power as it propagates through space, the range scaling exponent is smaller than the underlying energy efficiency scaling exponent. In this chapter we examine the scaling trend, present a range extrapolation for a wirelessly powered sensing workload, and discuss factors that could slow this range scaling trend.

1.2 Energy Efficiency Scaling of Computing Systems

Figure 1 shows the scaling of the energy efficiency of computing throughout the PC era. The raw data is from [5]. This data is for complete computer systems, including everything except the display. From the full dataset, we have selected data points corresponding to microprocessor-based computing (though interestingly the trend continues back through the vacuum tube era) and calculated energy efficiency in units that are more convenient for this chapter, instructions per μJ .

The energy efficiency is computed as follows. Take the number of instructions executed per second and divide by power consumption in watts (Joules per second). The resulting units are instructions per Joule: $\frac{\text{IPS}}{\text{Watt}} = \frac{\frac{\text{Inst}}{\text{s}}}{\frac{\text{J}}{\text{s}}} = \frac{\text{Inst}}{\text{J}}$. Divide by 10^6 to get instructions per μJ .

2 Range Scaling

The power required to execute a fixed workload has historically dropped with improved semiconductor process technology; wirelessly transmitted power attenuates with distance. The point at which wireless power available matches workload power requirement defines the wireless power range. Because power attenuates in a non-linear fashion, range does not scale linearly with energy efficiency. By combining empirical data on energy efficiency scaling with an analytical model of wireless power propagation, we will find an expression for range scaling with time.

2.1 Energy Efficiency of Microcontrollers

In this section we present historical data on the energy efficiency of microcontrollers, which are small embedded computing systems. Microcontrollers tend to prioritize energy efficiency over computing performance, since they are often used in battery-powered applications. This chapter is concerned with the range at which such embedded computing systems could be wirelessly powered.

To gather data on the energy efficiency of microcontrollers, we extracted the following parameters from microcontroller datasheets: instructions per second, power consumption (measured in watts, i.e., Joules per second), and year of microcontroller release. We selected the best performers (in terms of energy efficiency in a particular year) and plotted them in Fig. 2.

By analyzing the computations in terms of instructions per μJ , we are implicitly assuming that when the microcontroller is not active, it is not consuming power.

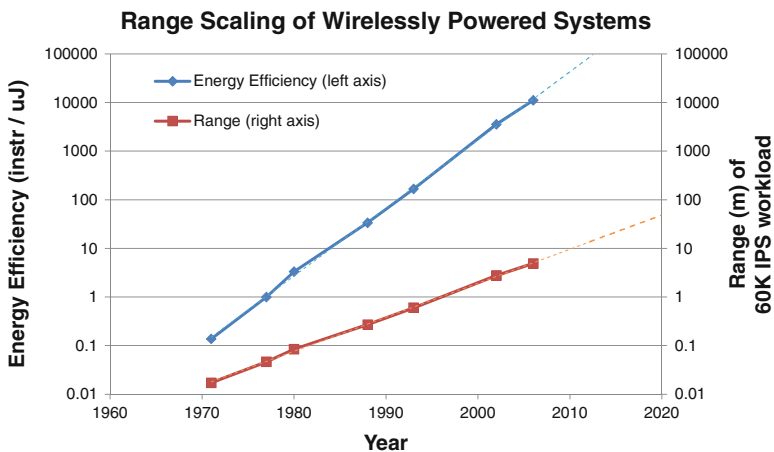


Fig. 2 Scaling of microcontroller energy efficiency (*top trace*) and feasible wireless power range (*bottom trace*)

Since the MSP430 has very low idle power, and transitions efficiently between active and idle, this is not a bad approximation for the systems that motivated this chapter. For energy-constrained systems such as these, the “efficient idling” approximation will presumably get better over time as increased attention is paid to reducing idle power. Even if leakage current rises as a fraction of total power budget, which naively might increase idle power consumption, it will be possible to gate power to most of the system, thereby keeping idle power low.

Note that energy efficiency values are greater in the microcontroller data than in the microprocessor data. For example, in 2006, the microcontroller data shows around 10,000 instructions per μJ , while the plot from [6] only shows about 100 instructions per μJ . Part of the explanation for the discrepancy is that the data presented in [6] and plotted in Fig. 1 is for complete computer systems (including memory, support circuitry such as display drivers, etc), whereas Fig. 2 is for the microcontroller in isolation. Also, the absolute computation rate (in instructions per second) is not constrained to be the same in the two plots, so energy efficiency (in instructions per μJ) is not the entire story. Running the microprocessors at the speeds of the microcontrollers would provide improvements in energy efficiency but would make the absolute computation rate insufficient for the requirements of general purpose computing. Finally, the data points in Fig. 1 were selected because they were commercially available computer systems; the data points in Fig. 2 were selected because they are particularly energy efficient. Despite the differences, it is interesting that the functional form of the scaling curves is the same, that is, a straight line on a semilog plot. The exponential energy scaling behavior appears to be quite generic.

2.2 *Scaling of Workload Power Requirements*

For the purposes of this chapter, we will consider a 6,000 instruction computation, executed ten times per second, for a net workload of 60K instructions per second. This is approximately the computational requirement of the Saturn accelerometer demo, described in the chapter “The Wireless Identification and Sensing Platform” [11] of this volume, or the musical input device described in the chapter “RFID-Vox: a Tribute to Leon Theremin” of this volume [8].

If we use F to denote the energy efficiency of microcontrollers (measured in instructions per μJ), then a fit yields $F \propto 2^{0.46t}$ or $F = c_1 2^{0.46t}$. For a given workload W , a fixed number of instructions per second, the power required is given by $P_{\text{REQ}} = \frac{W}{F}$. With $W = 60,000 \frac{\text{Inst}}{\text{s}}$, then $P_{\text{REQ}} = \frac{60,000 \frac{\text{Inst}}{\text{s}}}{c_1 2^{0.46t} \frac{\text{Inst}}{\mu\text{J}}} = \frac{60,000}{c_1} \frac{\mu\text{J}}{2^{0.46t} \text{s}}$. Combining the constants c_1 and 60,000 into a new constant c , we get

$$P_{\text{REQ}} = c 2^{-0.46t}.$$

We will assume that the transmitted power is fixed at 30 dBm, and the transmit antenna gain is 6 dBm, which places the combination at 36 dBm Effective Isotropic Radiated Power (EIRP), the FCC limit for a UHF RFID reader. For the receive antenna, we use 2 dBi, the value for a dipole antenna.

We use this formula to compute the distance at which our fixed computational workload (60K IPS, as explained above) could be wirelessly powered, for each energy efficiency point from the upper plot. The result is the lower trace in Fig. 2.

2.3 Wireless Power Propagation

To study range scaling, we also need a model of how wireless power attenuates with distance. The Friis transmission formula provides a simple analytical model of radiative wireless power propagation. In this model, power falls with the square of the distance. The Friis transmission formula is

$$P_R = P_T G_T G_R \left(\frac{\lambda}{4\pi r} \right)^2.$$

Here, P_R represents RF power received (after the antenna but before rectification), P_T is RF power transmitted, G_T and G_R are transmit and receive antenna gains, and r is range, the distance from the transmit to the receive antenna.

The rectifier converts the received RF power to DC. Voltage regulation circuitry converts the raw harvested voltage to the value required by the digital circuitry. We will model all the losses associated with rectification and power conditioning as a single multiplicative efficiency constant η whose value is between 0 and 1. P_{DC} is the received DC power: $P_{DC} = \eta P_R$. For harvesting efficiency, we use $\eta = 0.03$, that is, 3%, or approximately -15 dB of harvesting loss. This choice corresponds to the lowest efficiency at which our harvester operates.

Combining these to include power conversion efficiency, we get this expression for the received DC power, as a function of distance:

$$P_{DC} = \eta P_T G_T G_R \left(\frac{\lambda}{4\pi r} \right)^2.$$

It is feasible to power a particular workload at range r only if P_{REQ} , the power requirements of the workload, do not exceed P_{DC} , the received DC power at that range: $P_{REQ} \leq P_{DC}$.

$$P_{REQ} \leq \eta P_T G_T G_R \left(\frac{\lambda}{4\pi r} \right)^2.$$

We can rearrange this inequality to find an expression for feasible range, as a function of required power:

$$r^2 \leq \frac{\eta P_T G_T G_R}{P_{REQ}} \left(\frac{\lambda}{4\pi} \right)^2$$

and

$$r \leq \left(\frac{\eta P_T G_T G_R}{P_{\text{REQ}}} \right)^{\frac{1}{2}} \frac{\lambda}{4\pi}.$$

The energy efficiency scaling shown in the upper trace means that the power requirements P_{REQ} are a function of time. The wireless power range is the r at which the received power equals the required power: $P_{\text{DC}} = P_{\text{REQ}}$. Substituting our process scaling expression for P_{REQ} into the range formula yields

$$r \leq \left(\frac{\eta P_T G_T G_R}{c 2^{-0.46t}} \right)^{\frac{1}{2}} \frac{\lambda}{4\pi}.$$

The maximum range r_{MAX} occurs when r is equal to the expression on the right. Grouping together all the multiplicative constants into a single quantity a , we find this expression for the scaling of r_{MAX} with time:

$$r_{\text{MAX}} = a 2^{0.23t} \approx a 2^{t/4}.$$

This shows that keeping the computational workload and RF power transmitted constant, we can expect the range of far field wireless power systems to double every four years. This range scaling trend is plotted in the lower trace in Fig. 2, the main result of this chapter.

3 Discussion and Limitations

Most of the numerical factors that produce the range scaling trace (the lower trace in Fig. 2) are multiplicative, so changing them results in an upward or downward displacement of the trace on the semilogarithmic plot, but no change in slope. In particular, different workload sizes correspond to different parallel traces. Similarly, a different choice of power harvesting efficiency constant also results in vertical translation of the scaling plot (assuming that the efficiency constant does not change with time or range). Distances on the range scaling plot below one wavelength (about 0.3 m) are meaningless; the distances less than this essentially indicate that wirelessly powering the workload at that time was infeasible.

3.1 Energy Constraints Versus Voltage Constraints

The range scaling trend (the lower trace of Fig. 2) is ultimately derived from an energy bound. In practice, other factors, such as insufficient voltage, could keep the operational range below what the energy constraint alone would allow. However, energy is the more fundamental limit. Energy is a conserved quantity;

it cannot be created. Voltage, on the other hand, can be increased by a number of mechanisms, including transformers, DC to DC converters (boost regulators), diode-based voltage multipliers, or switched capacitor charge pumps. There do not appear to be fundamental limits to our ability to boost voltage; therefore, any limitations on range that arise due to insufficient voltage can in principal be addressed with voltage boosting techniques. Energy constraints, on the other hand, cannot be worked around.

3.2 Energy Constraints Versus Power Constraints

In contrast to energy constraints, which have no exceptions, *power* constraints can be worked around in this sense: if power is harvested at rate P_h and consumed at rate P_c , it is possible for P_c to instantaneously exceed P_h , if previously harvested and stored energy can be used to make up the instantaneous energy deficit. The fundamental constraint that cannot be worked around is the net energy constraint that $E_c = \int_0^T P_c dt < \int_0^T P_h dt = E_h$. This also implies an average power constraint $\bar{P}_c = \frac{1}{T} \int_0^T P_c dt < \frac{1}{T} \int_0^T P_h dt = \bar{P}_h$. Our definition of workload as a certain number of instructions executed *per second* extends the hard constraint on energy into a hard constraint on *average power*. If we did not include computation *rate* in the definition of workload, then a harvester-powered computer could perform a batch computation arbitrarily slowly and use the extra time to harvest as much energy as necessary to perform the task.

3.3 Power Harvesting Efficiency

It is possible that the smaller the input voltage (corresponding to long range), the more difficult it becomes to create an efficient rectifier. If for some reason rectification efficiency necessarily drops as input voltage drops, then the straight lower trace of Fig. 2 would droop downward. However, as discussed in Sect. 3.1, voltage does not obviously place fundamental limits on energy harvesting efficiency, since it is not a conserved quantity. Particularly if we are allowed to tune the system to optimize for a particular (low) input voltage, it may be that low voltage input signals are not associated with any fundamental physical limit on the efficiency with which the energy can be harvested. Clearly if we compare two source signals that have different energy content because of different voltage levels, it will be possible to harvest more energy from the higher voltage, higher energy signal; but there may be no difference in the efficiency with which we can harvest, that is, in the *fraction* of available energy that can be harvested.

3.4 *Energy Efficiency Scaling*

Another risk to the range scaling trend discussed here is that the underlying energy efficiency scaling will not continue. A slowdown in energy efficiency scaling would cause the upper trace of Fig. 2 to curve downward (which in turn would cause the lower trace to droop downward as well). Microelectronic scaling is indeed changing [4]. The rate of decrease in the operating voltage of microelectronics is decreasing, because leakage current becomes more significant as operating voltage approaches transistor threshold voltage. The long-standing “Dennard scaling” phenomenon [3], in which feature size and voltage change in together in a fashion that keeps electric field strength constant, has ended. Dennard scaling allowed simultaneous improvements in transistor density, energy efficiency, and computational performance (speed). Nevertheless, feature sizes, and therefore gate capacitances, are continuing to scale, which should continue to enable improvements in dynamic power consumption for low-end power-constrained devices such as the ones this chapter focuses on. The range scaling phenomenon can continue via improvements in energy efficiency that occur independent of or at the expense of improvements in computational performance.

3.4.1 *Energy Scaling of Sensors and Analog Electronics*

The energy efficiency scaling data discussed in this chapter is from digital systems, in particular microcontrollers. Typical sensor system applications also involve analog electronics and sensors; the energy scaling properties of these analog components are much less straightforward than those of digital microelectronics. Since many transducers are not fabricated using photolithography, the notion of technology scaling for sensors is not in general applicable. Nevertheless, as long as the energy efficiency of some sensor or set of sensors continues to improve with smaller feature sizes, we can expect range scaling to continue for at least these sensor systems. As an example, note that the power consumption of micro-machined accelerometers has been dropping dramatically in the past decade, presumably because of shrinking feature sizes. Again, the energy efficiency scaling trend discussed here should be thought of as a bound on achievable range, not as an average-case guarantee.

3.5 *Path Loss Exponent*

The Friis formula is an example of a more general *path loss* scaling relationship. In the Friis formula, the path loss exponent, which determines how quickly RF power attenuates with distance, is 2. Empirically, path loss exponents of 3 or 4 (even as high as 6 in particularly lossy environments) are sometimes observed. Using a path

loss exponent of 3 or 4, rather than the value of 2 used here, would cause the slope of the range scaling trace to decrease, slowing the predicted range scaling, but not changing the linear functional form observed in the semilog plot. In a space such as a corridor, which can function as a waveguide, path loss exponents lower than 2 are possible. This would yield a steeper range scaling slope.

4 Application Scaling

This chapter has focused on the improvements in wireless power range enabled by energy efficiency scaling. In practice, there is a parallel phenomenon that might be called “application scaling,” in which it becomes possible to wirelessly power new, more demanding applications, at the same range as the previous generation of applications. Long-range UHF RFID represents the first, minimal application of long-range wirelessly powered computing [1]. The next wave was “computational RFID”: the fixed state machine of conventional RFID tags is replaced by a fully programmable microcontroller (plus low-power sensors) [12, 13]. Computing workloads including cryptography [2] have been powered wirelessly. As technology scaling continues, it is becoming possible to power increasingly demanding sensing and computing workloads, such as microphones and cameras [14], sensors for safety and security (for burglar or fire alarms), computer input devices (mice and keyboards), and biomedical sensors (ECG, EMG, or EEG).

5 Conclusion

This chapter argues that the observed exponential improvements in the energy efficiency of microelectronics enables exponential improvements in the range at which a fixed computational workload can be wirelessly powered. Due to attenuation of propagating RF signals, the scaling exponent for wireless power range is lower than that for the underlying energy efficiency trend. If the range scaling trend continues, we can expect an exciting array of new capabilities, as it becomes possible to achieve perpetual operation without batteries in an increasingly large set of devices at increasingly long range from the RF power source.

Acknowledgments Thanks to Dan Yeager, who did the data collection and helped develop the calculations for Fig. 2. An earlier version of these results appeared in his master’s thesis [15] at University of Washington. Thanks to Jonathan Koomey, Shekhar Borkar, and Jan Rabaey for the comments and stimulating questions. Any remaining errors are mine.

References

1. MIT Auto-ID Center. 860–930 MHz Class I RFID Tag RF and Logical Communication Interface Specification v. 1.0.1, Nov. 2002.
2. H.-J. Chae, D.J. Yeager, J.R. Smith, and K. Fu. Maximalist cryptography and computation on the WISP UHF RFID tag. In *Proceedings of the Conference on RFID Security*, July 2007.
3. R.H. Dennard, F.H. Gaensslen, H.N. Yu, V.L. Rideout, E. Bassous, and A.R. LeBlanc. Design of ion-implanted MOSFETs with very small physical dimensions. *IEEE Journal of Solid State Circuits*, SC-9(5):256–268, 1974.
4. H. Esmailzadeh, E. Blem, R.S. Amant, K. Sankaralingam, and D. Burger. In *Proceedings of ISCA*, June 2011.
5. J.G. Koomey, S. Berard, M. Sanchez, and H. Wong. Assessing trends in the electrical efficiency of computation over time. *Final report to Microsoft and Intel*, August 2009.
6. J.G. Koomey, S. Berard, M. Sanchez, and H. Wong. Implications of historical trends in the electrical efficiency of computing. *IEEE Annals of the History of Computing*, 33(3):46–54, March 2011.
7. G.E. Moore. Cramming more components onto integrated circuits. *Electronics*, 38(8), 1965.
8. P.V. Nikitin, A. Parks, and J.R. Smith. RFID-Vox: a tribute to Leon Theremin. In J.R. Smith, editor, *Wirelessly powered sensor networks and computational RFID (this volume)*, New York, 2013. Springer SBM.
9. A.P. Sample and J.R. Smith. Experimental results with two wireless power transfer systems. In *IEEE Radio and Wireless Symposium, RWS '09*, pages 16–18, Jan. 2009.
10. A.P. Sample, A. Parks, S. Southwood, and J.R. Smith. A weather station powered by wireless ambient radio power (WARP). In J.R. Smith, editor, *Wirelessly powered sensor networks and computational RFID (this volume)*, New York, 2013. Springer SBM.
11. A.P. Sample and J.R. Smith. The wireless identification and sensing platform. In J.R. Smith, editor, *Wirelessly powered sensor networks and computational RFID (this volume)*, New York, 2013. Springer SBM.
12. 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. *IEEE Transactions on Instrumentation and Measurement*, 57(11):2608–2615, Nov. 2008.
13. J.R. Smith, A.P. Sample, P.S. Powledge, S. Roy, and A.V. Mamishev. A wirelessly-powered platform for sensing and computation. In *Proceedings of the 8th International Conference on Ubiquitous Computing (Ubicomp 2006)*, pages 495–506, September 17–21 2006.
14. V. Talla, M. Buettner, D. Wetherall, and J.R. Smith. In *Proceedings of the IEEE Topical Conference on Wireless Sensors and Sensor Networks (WiSNet)*, January 20–23 2013.
15. D.J. Yeager. Development and application of wirelessly-powered sensor nodes. Master’s thesis, University of Washington, 2009.

History of the WISP Program

Joshua R. Smith

1 Motivation

This chapter tells the story of the early years of the wireless identification and sensing platform (WISP) project, which created what we believe to be the first far field RF-powered sensing and computing platform. This chapter is about the history of the program: what were the motivations, who was involved, how did one piece of the project lead to the next, where were the dead ends, what other research grew out of it, and what were the impacts? This sort of background can be difficult to extract from the research papers themselves, which typically present self-contained results and do not convey the context. The context and background of the WISP program may be worth reflecting on because it has been such a fruitful research vein and because it bears on meta-research questions such as how to build a community. These meta-research questions are important because they affect the amount of impact that the research ultimately has.

While this present chapter focuses on the history and context of the WISP program, the next chapter “The Wireless Identification and Sensing Platform” is a detailed description of the WISP design and applications [16]. Also, this present chapter discusses just the work of my group and our collaborators. Other chapters of this book contain examples of related work by people who are not collaborators; some of these projects use the WISP, and some use different platforms. Projects on which my group did not collaborate are not discussed in this chapter, because I do not have the context.

J.R. Smith

Department of Computer Science and Engineering, Department of Electrical Engineering, University of Washington, Seattle, WA, USA
e-mail: jrs@cs.washington.edu

I began trying to combine sensing, computing, and RFID at Intel Research Seattle (later renamed Intel Labs Seattle) in 2004, motivated by a problem my colleague Matthai Philipose was grappling with. Matthai was interested in human activity recognition, motivated by eldercare. He had developed a short-range (HF, 13.56 MHz) RFID reader in the form of a bracelet and could detect that a person was using a tagged object when the bracelet reader got close enough to read the object's tag, which typically occurred when the object was being manipulated [13]. In an effort to eliminate the bracelet reader, Matthai and Ken Fishkin had experimented with estimating the motion of RFID tags from changes in the read rate of the (then relatively new) long-range UHF RFID tags [9]. Was there a more direct way to measure tag motion? Clearly an RFID tag with accelerometer sensors would solve the problem, but no such thing existed. At the time it seemed impossible given the power constraints.

1.1 α -WISP

The α -WISP (“alpha WISP”) was my first attempt to solve this problem: I used two mercury switches, in an antiparallel configuration, to multiplex two RFID chips (each with its own unique ID) to one antenna. The antiparallel arrangement ensured that if the tag were held stationary at a particular tilt angle, one switch would be open and the other closed. Each switch was connected in series to an ID chip, and the whole assembly of switches and chips was connected to an antenna. Thus when tilted one way, the α -WISP returns one ID to the reader; when tilted the other way, it returns a different ID. The two IDs can be thought of as two code words that (together, via the choice of one or the other of the code words) encode one bit of sensor information. The α -WISP thus enabled us to use an unmodified commercial RFID reader, and the existing unmodified EPC class 1 generation 1 (C1G1) UHF RFID protocol (which was not designed to transmit sensor data), to communicate one bit of sensor data, as well as many bits of ID data, from a battery-free, RF-powered uniquely identified sensor unit. I described the α -WISP as a one bit accelerometer sensor, because the mercury switches can be thought of as accelerometers with one bit of precision; they can be used to sense tilt because their state is determined by the direction of gravitational acceleration relative to the sensor axis [14, 27]. Later, Anthony Lamarca built a beautiful wooden box with a permanently “implanted” α -WISP that allowed the box's state (open or closed) to be read, along with the box's ID, by an EPC Gen 2 RFID reader. In this new form of α -WISP, a permanent magnet was built into the box top; in the body of the box was an α -WISP with an SPDT magnetic reed switch. The proximity of the magnet in the lid to the magnetic reed switch in the box body determined which of the two IDs the box returned when interrogated by the reader [31].

1.2 π -WISP

With the π -WISP (“pi WISP”), I wanted to send more than one bit of sensor data. I built a 3 axis \times 1 bit accelerometer from three orthogonal mercury switches. Graduate student Bing Jiang designed a small UHF power harvester board that used Schottky diodes to rectify the 915 MHz RF signal emitted by the reader, converting the RF power collected by the harvesting antenna into a DC power output. This RF harvester was used to power the TI MSP430 microcontroller. The microcontroller read the state of the mercury switches and then encoded the data using what I called “ID modulation”: the micro was connected to a gallium arsenide single pole double throw (SPDT) switch that multiplexes two RFID chips to one tag antenna. The mechanical switch of the α -WISP had been replaced by a solid-state, electronically controlled switch. (A gallium arsenide switch was chosen because the switch had to pass ultra-high-frequency RF signals, which are too high in frequency for typical silicon devices to handle.) But because the IDs were now being selected by software (instead of by a mechanical modulator, the mercury switch), the π -WISP could produce an arbitrarily complex time series, in the patterns of ID response. Having previously worked on information hiding (the problem of embedding one signal in another, e.g., a digital watermark in an image), I liked to think of this as “hiding” a stream of sensor data in another stream of data, the sequence of RFID reads. Because the RFID reader (which might be called “the warden” by steganographers) does not notice any strange behavior, we were essentially able to overlay or embed a new protocol (for sensor data) in an old one (for IDs). This embedded (or overlay) protocol was very inefficient: its net data rate was less than one bit of sensor data per second, since the sensor data was encoded in a long stream of ID reads, which included “packet headers” composed of multiple RFID read events. The packet headers were necessary for synchronization: without them (or some other channel sharing scheme in the overlay layer), it would not have been possible for the reader to distinguish bit one of the sensor data from bits two and three of sensor data. Although this protocol for overlaying sensor data in a stream of ID reads was painfully slow, it was exciting that we could now communicate an arbitrarily large amount of sensor data, collected in a battery-free fashion, using apparatus that was not designed to support this functionality [26, 28, 29].

1.3 WISP

Next we built a battery-free, RF-powered EPC C1G1 RFID tag entirely from scratch discrete components. The tag logic was implemented entirely in microcontroller software. Unlike the α -WISP and the π -WISP, the WISP did not include commercial RFID chips. The path from the π -WISP to the WISP was not an entirely

straight line. I had hired new UW graduate student Alanson Sample as an intern to help design a chip, to be fabbed through the new Intel Research Shuttle program. The Research Shuttle was intended to allow researchers to design chips to be built on Intel processes. The chip we designed was the analog front end of an RFID tag: it included a power harvester, a demodulator to extract data from the RFID reader, and a modulator to backscatter information to the reader. The plan was to power the MSP430 via the output of the harvester and (similar to the π -WISP) to use the MSP430 as a software radio to implement the protocol (in this case the full RFID protocol, not an overlay). The result would be a fully programmable RFID tag consisting of two chips, plus whatever sensors were desired.

Soon after the chip was designed and simulated, the Intel Research Shuttle program was canceled. Disappointed, we decided to use discrete components to build an RFID analog front end based on the one that had already been designed, despite the fact that the design had originally been intended to become an IC. While most people think of RFID tags as chips, in fact, a printed circuit board implementation has a number of advantages: one can include low-threshold Schottky diodes (which are not available in many CMOS processes) for efficient power harvesting, one can include capacitor values that would be infeasibly large on an IC, and one can integrate sensors without constraints on fabrication process—mixing and matching is allowed.

Instead of a two-chip tag, the WISP ended up including the microcontroller, passive components, Schottky diodes, and sensors. Some versions included a comparator (in the demodulator) and voltage supervisors (to wake the MSP430 from low-power deep-sleep states). We were uncertain whether the MSP430 would be fast enough to implement the EPC C1G1 protocol. Software engineer Pauline Powledge wrote the first working WISP firmware. The most challenging moment in the entire WISP program was getting the WISP to successfully talk to a commercial RFID reader for the first time. There were numerous advantages to using commercial RFID readers, but ease of debugging tag-reader interactions was not one of them. The problem is that commercial RFID readers are built as black boxes; if the reader does not feel that the tag is behaving properly, the reader just ignores it—it does not provide helpful error messages, or any feedback whatsoever, to the newbie tag designer. At one point, in desperation, we put oscilloscope probes across the antenna of a well-behaved commercial UHF RFID tag, to see how its behavior differed from our tag. We discovered the Generation 1 Alien reader deviated substantially from the published protocol. Once this discovery had been made, we were able to get our first WISP working.

Building on the ID modulation idea used in π -WISP, we allocated certain bits of the tag ID to sensor data. As with the α - and π -WISPs, the reader would simply pass on these sensor bits (the low eight bits of the tag ID, say) without knowing that it was handling sensor data; to the reader it simply looked like a stream of IDs from a quickly changing tag population. This technique was far more efficient than the

ID modulation scheme used by the π -WISP, since a single RFID read even could convey as many bits of sensor data as we were willing to “steal” from the ID space. Suddenly we could send sensor data at what seemed like blistering speeds [17, 30].

1.3.1 Gen 2 WISP

While the WISP was maturing, so were RFID standards. The original EPC C1G1 specification was supplanted by the EPC Class 1 Generation 2 specification. Seong Ho Kim, together with UW graduate student intern Dan Yeager, created the first version of the WISP firmware that supported the Gen 2 specification. The Gen 1 WISP firmware had been written in C. The higher bit rates of the Gen 2 specification required the most time-sensitive parts of the firmware to be written in hand-optimized assembly language to squeeze enough performance from the MSP430. The earlier experiments referenced in this chapter used the C1G1 spec; later ones used “Gen 2.” Michael Buettner did an implementation of the Gen 2 MAC (medium access control) layer, which is described in the chapter “Implementing the Gen 2 MAC on the Intel-UW WISP” of this volume [2].

1.4 Accelerometer WISP

The summer after the first WISP (described in [30]) came to life, I asked Dan Yeager to connect a new, very low power three-axis accelerometer to our latest WISP. The accelerometer, the Analog Devices ADXL 330, consumed only 200 μ A at 1.8 V. Before this accelerometer, it would not have been feasible to power and read an accelerometer using RF signals. We believe that this WISP with accelerometer was the first UHF-powered and -read accelerometer. Having started with a very primitive RF-powered one bit accelerometer, we finally had a real three-axis accelerometer, with eight bits of real accelerometer data for each axis, entirely powered and read by a commercial RFID reader [23, 36]. Later Matthai Philipose, working with Michael Buettner and David Wetherall, used the WISP’s accelerometer for activity recognition, the original inspiration for the project [5].

1.5 WISP Passive Data Logger

In “cold-chain monitoring,” the goal is to verify that temperature-sensitive items, such as vaccines, blood products, or frozen food, have been kept within a required temperature envelope. In this application, and many other “shipping” applications,

it is not feasible to assume that the sensor is near an RFID reader at all times. Dan Yeager and I proposed a Passive Data Logger to address this application space. The Passive Data Logger is a WISP with a large energy store and large memory, likely nonvolatile. The model is that the WISP, mounted on an item whose temperature is to be monitored, accumulates energy while it is waiting at its original location, perhaps a warehouse freezer with a built-in RFID reader. While the item is in transit, the stored energy is used to sense and log data. At the receiving end, the logged data is downloaded via the RFID interface, and the tag begins harvesting power to replenish the energy that was consumed during transit [35].

1.6 Neural WISP

Dan Yeager created a WISP to drive a custom neural amplifier IC designed by members of Brian Otis's group. This effort was successful, although we encountered some interference between sensing and communication in this application. The carrier emitted by the reader is amplitude-modulated to encode downlink data (data that flows from the reader "down" to the tag). The downlink data modulations caused sensor noise. In the application, we solved the problem by time multiplexing between sensing and communication [37].

1.7 Strain Gage WISP

Working with Professor Paolo Feraboli's group and researchers from Boeing, we created a WISP strain gage sensor system targeted at health monitoring for structures such as wings or car bodies made from composite materials. One can imagine permanently embedding battery-free sensors in wings, since they are capable of perpetual operation and can be temporarily energized only when read [10].

1.8 Solar WISP

In the Passive Data Logger [35], the problem of sensor data logging for items in transit (away from a reader) was solved by accumulating energy during time spent near the RF source and then using it later when the tag is no longer near the RF source. In the Solar WISP, Alanson Sample and then-undergraduate Aaron Parks decided to use a secondary energy source, solar, to power the data logger. The elegant result in this project is that a solar cell can be used directly as the RFID antenna: separate structures are not needed. Using the solar cell for both solar and

RF harvesting saves area and cost. The flexible solar cell used in this project showed that the resulting system could potentially take the “sticker” form factor common in RFID tags [20].

1.9 RFID Localization with the LED WISP

RFID localization is a difficult and important problem. If an RFID tag can be precisely localized, it makes capabilities like robotic retrieval of the object much more feasible. For a robot to pick up an object, it needs to know quite precisely where the object is. Localizing RFID tags via reflected RF signal strength is generally unreliable because of multipath. Relatively small changes in the environment, including a person walking by, can dramatically change RF signal strength readings. The RF signals can take multiple paths from the source to the destination; the “ray” along each path has its own phase, affected by the length of the path. The rays sum at the destination location and can interfere constructively or destructively, depending on the detailed geometry of the environment and the paths taken by each ray. Changing the path length (and thus the phase) for one ray can drastically change the resulting signal strength at the destination.

Dan Yeager and I, together with Ali Rahimi, a computer vision colleague at Intel Research Seattle, realized that we could build an LED WISP that could be powered and read by an RFID reader but localized precisely by a camera looking at the LEDs. Dan designed a WISP that included four LEDs, one in each corner. A robot-mounted RFID reader might detect the desired object on a shelf in a warehouse, along with many other tagged objects that are not of interest. With the LED WISP, it should be possible for the reader to command the single tag of interest to flash its LED and hopefully be localized well enough for the robot to retrieve it.

While this idea itself was enough to generate a patent [32], we were not able to get the idea to actually work at that time. Given the power constraints, it turned out to be difficult for the camera to detect the LED. It was only possible to light the LED for a very short period of time, and with a free running camera, not synchronized with the tag, the chance of missing the flash was high.

A couple of years later, at UW, my group returned to the problem. We realized that we needed to synchronize the camera with the RFID reader in order for the idea to work. Since the commercial RFID reader is still essentially a black box that does not provide access to its internal state, we built a “trigger WISP” that monitors the traffic from the reader to the tag of interest. By mimicking the state of the LED WISP (inferred from the reader traffic), the trigger WISP can reliably determine when the LED WISP will flash. The trigger WISP then triggers a camera to capture a frame when the LED is guaranteed to be on. By taking another image with the LED off and differencing the images, the LED can be found very reliably. The team included undergraduate Craig Macomber, who did the smart camera and image processing work, Liang-Ting Jiang, who programmed our PR2 robot, and Alanson Sample, who put the whole system together and did the optical localization work [21].

1.10 SOCWISP

For his master's thesis, Dan Yeager worked with Brian Otis to design a "System On Chip" WISP. This chip includes the EPC RFID analog front end, as well as the digital state machine for an EPC C1G2 tag. A separate application microcontroller can provide data to the SOCWISP via a serial interface. The data becomes the ID that the SOCWISP returns to the RFID reader. One of the remarkable things about this chip was that the first design to be fabricated actually functioned. Part of the reason is that Dan was able to extensively test and debug the digital state machine against a real RFID reader, by implementing the state machine in an FPGA connected to a WISP analog front end. The SOCWISP was so small and light that Dan was able to fly it on a moth and make in-muscle temperature measurements while the live moth was flying. The paper describing SOCWISP appears in this volume as the chapter "SOCWISP: A 9 μ A, Addressable Gen2 Sensor Tag for Biosignal Acquisition" [39]; it was originally published as [38].

1.11 Security

Although the WISP had been conceived as a platform for sensing and computing, it received immediate interest from the security community. Because RFID tags had been black boxes and could not run software, it had not been feasible to implement and test security protocols and encryption algorithms that required nonstandard behavior from UHF RFID tags. The WISP suddenly made this possible.

1.11.1 Encryption

Kevin Fu, a computer science professor then at University of Massachusetts, immediately saw the possibilities of WISP for security research. We gave him some WISPs, and his group became the first outside users of WISP. The first joint publication between the two groups was an implementation of the RC5 block cipher algorithm on the WISP. One school of thought in RFID security is that, because of the limited power and computing cycles, special "minimalist" cryptographic algorithms are required. This paper, entitled "Maximalist Cryptography and Computation on the WISP UHF RFID Tag," introduced a contrary approach, by implementing a full-blown conventional "desktop" cryptographic algorithm on an RF-powered device. The paper was presented at a conference [6], but never published; an updated version appears as the chapter "Maximalist Cryptography and Computation on the WISP UHF RFID Tag" [7] of this volume.

Kevin coined the term computational RFID as a generic term for WISP-like devices, and his group went on to publish many of their own papers on computational RFID. His group has even built their own computational RFID units,

which they named the MOO, for its key feature: the very beefy microcontroller (a larger model MSP430) [40]. The larger micro allows implementation of more sophisticated software, but the increased power consumption limits range.

The security community has continued to innovate using the WISP platform. For example, Pendl, Pelnar, and Hutter implemented elliptic curve cryptography (ECC) on the WISP [12].

1.11.2 Security Through Sensing

One of the challenges in RFID security is that tags can easily be read without their owner's knowledge or permission. Sensors can help with this, by giving the tag's owner a user interface to the tag that can be used to authorize tag responses.

Secret Handshakes

Alexis Czekis and Karl Koscher, two UW CSE graduate students working with Tadayoshi Kohno, proposed using WISP to prototype an RFID access control tag that would be resistant to “ghost and leech” attacks, an attack that is specific to RFID access control tags. This attack relies on the promiscuity of ordinary RFID tags, which will respond to any reader. One of the attackers, the “leech,” draws close to a person who is known to be carrying an RFID access control card for a space that the attackers wish to enter. The leech's RFID reader queries the access control tag and relays the response to the “ghost,” whose programmable RFID tag responds to the RFID reader responsible for access control. This relay attack causes the ghost of the authorized person to appear in front of the access control reader. Note that even if the reader engages in a challenge-response scheme with the tag, suitable versions of this attack will still work.

The solution they proposed was to use the WISP's sensing capabilities to enhance security. Rather than respond promiscuously to any reader, the tag will respond to the reader only after it is moved through a certain gestural trajectory, essentially a gestural password. The accelerometer was used to sense the trajectories, and the microcontroller performed the gesture recognition computations via a template correlation scheme [8].

Capacitive Touch WISP

Another way to add user input to an RFID tag is to use the tag antenna as a capacitive sensor. The paper demonstrating this idea using WISP won the best paper award at IEEE RFID in 2009 (after being flatly rejected the previous year—persistence pays off!) [18]. This technique is described in more detail elsewhere in this volume [16].

1.12 Networking

David Wetherall, Michael Buettner, and Ben Greenstein led an effort to start exploring networking issues raised by WISP sensors [3, 4]. This volume contains a very interesting contribution (Chap. 7, by Molina–Markham et al. [11]) to this space, an overlay to the RFID protocol that allows WISP to WISP communication. This overlay abstraction is implemented at the lower levels by having the RFID reader relay messages from one WISP to another.

2 Commercial Impact

In this section we consider the potential for commercial impact of WISPs and WISP-like devices. It appears that the dream of arbitrarily inexpensive RFID tags will not come to pass. This is largely because the fixed costs associated with each tag, such as dicing, testing, antenna manufacturing, and tag assembly (i.e. mounting the tag IC on the antenna) do not benefit from technology scaling. It seems that these per-tag costs may put a price floor on low-end RFID tags, even if the silicon area required per tag continues to decrease. Even if the silicon area itself were cost free, these tag assembly and test costs would remain.

Given these dynamics, it appears that technology scaling should allow more functionality to be added to RFID tags for little cost above that of the most minimal RFID tag. If this analysis is correct, we would not expect to see the cost of RFID tags drop much below their present prices, but for that price, the tags can become more and more capable. The next few sections consider possible commercial applications for sensing and computing enhanced RFID tags.

2.1 Secure RFID

The ability to implement strong cryptographic algorithms such as RC5 (described in Chap. 15 [7]) and AES (described in Chap. 16 [33]) on passive RF-powered tags will likely be attractive for commercial applications. Using a WISP-like hardware platform, it would be possible to implement a secure RFID tag for applications such as access control or automotive tolling that could communicate securely via a conventional EPC Gen 2 reader. Ciphers such as AES, RC5, or ECC can be implemented in software, avoiding the time and expense of creating custom silicon for these functions.

2.2 *Embedded RFID*

WISP-like circuitry could be embedded in larger systems, such as phones or laptops, to provide some level of functionality, even when the host system is powered completely down. Configuration state for the system could be stored in “dual-ported” nonvolatile memory. With the device fully off, the configuration data or firmware could be read and (after modification) rewritten via the RFID interface. When the system powers up, it would read the data from nonvolatile memory via a conventional wired interface. This could allow configuration edits while the device is off, firmware upgrades while the device is still in its original manufacturer’s packaging. Dirk Haehnel and I received a patent for this idea [25].

The chapter “Passive RFID-Based Wake-Up Radios for Wireless Sensor Networks” of this volume is another nice example of combining a WISP with a larger, battery-powered device in order to save power in the battery-powered device. They use the WISP, which consumes zero standby current, to wake up a conventional battery-powered sensor node. Using a conventional radio receiver to wake a sensor node consumes nonzero standby current consumption. Using a timer to wake the sensor node on a schedule also requires more power than waking from the interrupt that the WISP is used to generate [1].

2.3 *Building Community: Open Source, the WISP Challenge, and the WISP Summit*

It was clear that a platform like the WISP had many more possible applications than we could possibly explore ourselves. We open-sourced the firmware (via the BSD license) and posted all the schematics and design files on the web. With the support and encouragement of David Wetherall and Intel, we started a program, the “WISP Challenge,” to make WISPs available to academic researchers. We solicited applications and awarded WISPs to the best proposals. Our aim was to seed the growth of a community of researchers interested in perpetually powered sensing and computing systems. Several of the chapters in this volume are the result of WISP Challenge awards.

After WISP Challenge users had some time to work with the WISP, our first “customer,” Kevin Fu, suggested we organize a WISP Summit, to exchange information and see what sorts of results the community was generating. We held the first WISP Summit in Berkeley, CA, in conjunction with ACM Sensys 2009 (a major sensor networks conference). Videos of the first WISP Summit are available on the web.

3 Outgrowths

Other exciting projects exploring different forms of RF power harvesting have emerged from WISP, taken on lives of their own, and become major new projects in their own right. These projects are briefly described here; each has a chapter later in this volume.

3.1 *WARP: Wireless Ambient Radio Power*

Having become comfortable with harvesting “planted” RF power that had been deliberately emitted by an RFID reader for the purpose of powering electronic devices, we wondered if it would be possible to harvest “wild” RF power from ambient RF sources, such as TV, radio, or cell phone towers. Alanson found that there was a 1 MW (one million watt) digital TV tower about 4 km Intel Labs Seattle, and the balcony had an excellent view of the tower. Based on Friis’s simple propagation model, we expected to see about 200 μ W on the balcony. Our harvester delivered about 60 μ W of rectified DC, which is about what we expected to see given the efficiency of our harvester, around 25% [15]. The chapter “Wireless Ambient Radio Power” of this volume presents new results on the WARP project [22].

3.2 *WREL: Wireless Resonant Energy Link*

When a group of physicists at MIT published their work on wireless power using magnetically coupled resonators (MCRs), I was the natural person at Intel to engage with it, since I had been working on wireless power for several years by then. Alanson Sample joined me again and began what would become the core of his Ph.D. thesis on wireless power. Alanson and I, together with undergraduate David Meyer, modeled the system using lumped circuit elements. This allowed us to clearly see the effects of mode splitting and opportunities to compensate for it. We built what we believe was the first MCR-based wireless power system that could adapt to changes in transmit–receive distance, coil orientation, or load. We called the system WREL, wireless resonant energy link [19].

3.2.1 **FREED: Free-Range Resonant Electrical Energy Delivery**

The chapter “A Portable Transmitter for Wirelessly Powering a Ventricular Assist Device Using the Free-Range Resonant Electrical Energy Delivery (FREE-D) System” [34] of this volume, by my U.W. graduate student Ben Waters (together with Alanson Sample, Yale heart surgeon Pramod Bonde, undergraduates Kara Kagi

and Jordan Reed, and me), describes our FREED system, which is a very exciting application of MCR-based wireless power. FREED supplies energy to the heart pumps known as LVADs (left ventricular assist devices). The goal is to eliminate the infection-prone “drive line,” a thick cable that protrudes from the abdomen of present-day LVAD patients.

4 The Future

If the range scaling phenomenon described in the chapter “Range Scaling of Wirelessly Powered Sensor Systems” [24] continues, then capabilities that work with insufficient or barely sufficient range today should become feasible at much longer range in the future. But realizing this promise will require research effort.

4.1 *Power Harvesting*

Even if we assume that the power requirements and RF power available scale together (as they will if the technology follows the range scaling trajectory suggested in the chapter “Range Scaling of Wirelessly Powered Sensor Systems”), the voltage presented to the harvester will drop as range increases, unless measures are deliberately taken to boost the input voltage. An open fundamental question, raised in the chapter “Range Scaling of Wirelessly Powered Sensor Systems” is whether received voltage puts any fundamental bounds on the efficiency with which the energy can be harvested. Then there are the open practical questions of how to build the best possible harvester for any particular range (input voltage) and load. Agile harvesting is one exciting area of research: how to design harvesters that can tune themselves for different frequencies, amplitudes, and load levels.

Biological organisms are highly effective energy harvesters. Animals typically use previously stored energy to “fund” their present energy-harvesting efforts. Could RF power harvesters, using powered, active harvesting electronics, benefit from this strategy? Like biological creatures, active harvesters would not be able to let their energy supply drop to zero. Unlike today’s simple energy harvesters, there might be no recovery from a starvation event.

4.2 *Networking*

As RFID tags evolve into RF-powered computers, the applications and usage models will become both more diverse and more complex. This will motivate the development of more sophisticated networking protocols. The BAT scheme, described in the chapter “BAT: Backscatter Anything-to-Tag Communication” [11],

is a step in the right direction. If future tags become full-fledged internet hosts, and RFID readers become access points (APs) connected to the internet, it will be possible to structure RFID systems in a completely new way. In a typical RFID system today, the tag is a simple “license plate” (i.e., a unique ID that is used as a pointer to additional information), and the reader is a straightforward conduit between the tag and a back end computer. It is the back-end computer that runs the application. Because the reader is tightly tied to an application, tags are only useful when they are read by readers that are prepared for those specific tags.

Contrast this with Wi-Fi networks today. As long as a client is authorized to use an AP, the client can communicate with any host on the internet. The client is not restricted to use a particular set of application software specific to each AP. If future highly capable computational RFIDs carry their own application software, and readers become APs that implement general purpose routing protocols, then from any reader in the world a tag could communicate sensor data or other information to its home base (server) or to another tag.

4.3 *Big Power*

While “big data” is currently a hot research topic throughout computer science, I believe that “big power” could be next, at least within the computational RFID community. Of course the notion of “big” is relative; “big power” for a WISP would seem very small to most other communities. The idea is to execute sensing and computing workloads, under RF power, that today seem impractically large, just as the RFID accelerometer seemed impossible when the WISP project started. Workloads such as cameras consume an amount of power that by RF-harvesting standards seems excessive initially; the standby current of the camera may exceed the power harvested from the RF source. One key to achieving “big power” computational RFID systems is that, as explained in the chapter “Range Scaling of Wirelessly Powered Sensor Systems” [24], power is not a conserved quantity (since it can be collected at one rate and spent at another); only energy is conserved. So the “big power” systems we are imagining would harvest energy at whatever (low) rate the source (RF or otherwise) is able to provide; the power-consuming portions of the system (such as the camera) would be completely powered off to avoid wasting the camera’s standby current. Once sufficient energy has been accumulated after a long period of harvesting, the energy is spent suddenly, at high power levels for a short period of time. This approach should make it possible soon to RF power devices such as cameras that so far have seemed far too power hungry to ever be powered and read by an RFID reader.

Acknowledgments I thank James Landay, David Wetherall, Dieter Fox, Anthony Lamarca, and Matthai Philipose for the support and energy they gave to the WISP program at Intel Labs Seattle. Thanks to faculty collaborators Kevin Fu, Brian Otis, Tadayoshi Kohno, Paolo Feraboli, Sumit Roy, and Alexander Mamishev. I had the pleasure of working with many talented students on WISP and

related projects: Alanson Sample, Dan Yeager, Aaron Parks, Justin Reina, Bing Jiang, Seong–Ho Kim, Jeff Braun, Kishore Sundara–Rajan, Michael Buettner, Vamsi Talla, Yi “Eve,” Zhao, Liang–Ting Jiang, Craig Macomber, Jim Youngquist, Ben Waters, Gunbok Lee, and others. I thank all the students for their fantastic work.

References

1. H. Ba, I. Demirkol, and W. Heinzelman. Passive RFID-based wake-up radios for wireless sensor networks. In J.R. Smith, editor, *Wirelessly powered sensor networks and computational RFID (this volume)*, New York, 2013. Springer SBM.
2. M. Buettner and D. Wetherall. Implementing the Gen 2 MAC on the Intel-UW WISP. In J.R. Smith, editor, *Wirelessly powered sensor networks and computational RFID (this volume)*, New York, 2013. Springer SBM.
3. M. Buettner, B. Greenstein, A. Sample, J.R. Smith, and D. Wetherall. Revisiting smart dust with RFID sensor networks. In *Proceedings of the 7th ACM Workshop on Hot Topics in Networks (HotNets-VII)*, 2008.
4. M. Buettner, R. Prasad, A. Sample, D. Yeager, B. Greenstein, J.R. Smith, and D. Wetherall. RFID sensor networks with the Intel WISP. In *Proceedings of the 6th ACM conference on Embedded network sensor systems*, pages 393–394. ACM, 2008.
5. M. Buettner, R. Prasad, M. Philipose, and D. Wetherall. Recognizing daily activities with RFID-based sensors. In *Proceedings of the 11th international conference on Ubiquitous computing*, pages 51–60. ACM, 2009.
6. H.-J. Chae, D.J. Yeager, J.R. Smith, and K. Fu. Maximalist cryptography and computation on the WISP UHF RFID tag. In *Conference on RFID Security (website)*, July 2007.
7. H.-J. Chae, M. Salajegheh, D.J. Yeager, J.R. Smith, and K. Fu. Maximalist cryptography and computation on the WISP UHF RFID tag. In J.R. Smith, editor, *Wirelessly powered sensor networks and computational RFID (this volume)*, New York, 2013. Springer SBM.
8. A. Czeskis, K. Koscher, J.R. Smith, and T. Kohno. RFIDs and secret handshakes: defending against ghost-and-leech attacks and unauthorized reads with context-aware communications. In *Proceedings of the 15th ACM conference on Computer and communications security, CCS '08*, pages 479–490, New York, NY, USA, 2008. ACM.
9. K. Fishkin, B. Jiang, M. Philipose, and S. Roy. I sense a disturbance in the force: Unobtrusive detection of interactions with RFID-tagged objects. *UbiComp 2004: Ubiquitous Computing*, pages 268–282, 2004.
10. F. Gasco, P. Feraboli, J. Braun, J. Smith, P. Stickler, and L. DeOto. Wireless strain measurement for structural testing and health monitoring of carbon fiber composites. *Composites Part A: Applied Science and Manufacturing*, 42(9):1263–1274, 2011.
11. A. Molina-Markham, S.S. Clark, B. Ransford, and K. Fu. Bat: Backscatter anything-to-tag communication. In J.R. Smith, editor, *Wirelessly powered sensor networks and computational RFID (this volume)*, New York, 2013. Springer SBM.
12. C. Pendl, M. Pelnar, and M. Hutter. Elliptic curve cryptography on the wisp uhf rfid tag. *RFID. Security and Privacy*, pages 32–47, 2012.
13. M. Philipose, K.P. Fishkin, M. Perkowitz, D.J. Patterson, D. Fox, H. Kautz, and D. Hahnel. Inferring activities from interactions with objects. *Pervasive Computing, IEEE*, 3(4):50–57, 2004.
14. M. Philipose, J.R. Smith, B. Jiang, A. Mamishev, S. Roy, and K. Sundara-Rajan. Battery-free wireless identification and sensing. *IEEE Pervasive Computing*, 4(1):37–45, 2005.
15. A. Sample and J.R. Smith. Experimental results with two wireless power transfer systems. In *IEEE Radio and Wireless Symposium, RWS '09*, pages 16–18, Jan. 2009.

16. A. Sample and J.R. Smith. The wireless identification and sensing platform. In J.R. Smith, editor, *Wirelessly powered sensor networks and computational RFID (this volume)*, New York, 2013. Springer SBM.
17. 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. *IEEE Transactions on Instrumentation and Measurement*, 57(11):2608–2615, Nov. 2008.
18. A.P. Sample, D.J. Yeager, and J.R. Smith. A capacitive touch interface for passive RFID tags. In *IEEE International Conference on RFID*, pages 103–109, April 2009.
19. A.P. Sample, D.A. Meyer, and J.R. Smith. Analysis, experimental results, and range adaptation of magnetically coupled resonators for wireless power transfer. *IEEE Transactions on Industrial Electronics*, 58(2):544–554, 2011.
20. A.P. Sample, J. Braun, A. Parks, and J.R. Smith. Photovoltaic enhanced UHF RFID tag antennas for dual purpose energy harvesting. In *IEEE International Conference on RFID*, pages 146–153, 2011.
21. A.P. Sample, C. Macomber, L.T. Jiang, and J.R. Smith. Optical localization of passive UHF RFID tags with integrated LEDs. In *IEEE International Conference on RFID*, pages 116–123, 2012.
22. A. Sample, A. Parks, S. Southwood, and J.R. Smith. Wireless ambient radio power. In J.R. Smith, editor, *Wirelessly powered sensor networks and computational RFID (this volume)*, New York, 2013. Springer SBM.
23. J.R. Smith. RFID tag with accelerometer, June 7 2011. US Patent 7,956,725.
24. J.R. Smith. Range scaling of wirelessly powered sensor systems. In Joshua R. Smith, editor, *Wirelessly powered sensor networks and computational RFID (this volume)*, New York, 2013. Springer SBM.
25. J.R. Smith and D. Haehnel. Device configuration with RFID, November 2 2010. US Patent 7,825,776.
26. J.R. Smith and J.A. Landay. Time domain embedding of application information in an RFID response stream, December 14 2005. US Patent App. 11/304,511.
27. J.R. Smith and M. Philipose. Inertially controlled switch and RFID tag, February 26 2008. US Patent 7,336,184.
28. J.R. Smith, K.P. Fishkin, B. Jiang, A. Mamishev, M. Philipose, A.D. Rea, S. Roy, and K. Sundara-Rajan. RFID-based techniques for human-activity detection. *Communications of the ACM*, 48(9):39–44, 2005.
29. J. Smith, B. Jiang, S. Roy, M. Philipose, K. Sundara-Rajan, and A. Mamishev. ID modulation: Embedding sensor data in an RFID timeseries. In *Information Hiding*, pages 234–246. Springer, 2005.
30. J.R. Smith, A.P. Sample, P.S. Powledge, S. Roy, and A. Mamishev. A wirelessly-powered platform for sensing and computation. In *UbiComp*, pages 495–506, 2006.
31. J.R. Smith, A. Lamarca, and M. Philipose. Switch status and RFID tag, August 12 2008. US Patent 7,411,505.
32. J.R. Smith, D. Yeager, and A. Rahimi. Radio frequency identification tags adapted for localization and state indication, July 17 2012. US Patent 8,222,996.
33. A. Szekely, M. Hofler, R. Stogbuchner, and M. Aigner. Security enhanced wisps: Implementation challenges. In J.R. Smith, editor, *Wirelessly powered sensor networks and computational RFID (this volume)*, New York, 2013. Springer SBM.
34. B. Waters, K. Kagi, J. Reed, A. Sample, P. Bonde, and J.R. Smith. Powering a vad using the portable freed system. In J.R. Smith, editor, *Wirelessly powered sensor networks and computational RFID (this volume)*, New York, 2013. Springer SBM.
35. D.J. Yeager, P.S. Powledge, R. Prasad, D. Wetherall, and J.R. Smith. Wirelessly-charged UHF tags for sensor data collection. In *IEEE International Conference on RFID*, pages 320–327, 2008.
36. D.J. Yeager, A.P. Sample, J.R. Smith, and J.R. Smith. WISP: A passively powered UHF RFID tag with sensing and computation. *RFID Handbook: Applications, Technology, Security, and Privacy*, pages 261–278, Boca Raton, FL, 2008. CRC Press.

37. D.J. Yeager, J. Holleman, R. Prasad, J.R. Smith, and B.P. Otis. NeuralWISP: A wirelessly powered neural interface with 1-m range. *IEEE Transactions on Biomedical Circuits and Systems*, 3(6):379–387, 2009.
38. D. Yeager, F. Zhang, A. Zarrasvand, N. George, T. Daniel, and B. Otis. A $9\mu\text{A}$, addressable Gen2 sensor tag for biosignal acquisition. *IEEE J. Solid-State Circuits*, 45(10):2198–2209, 2010.
39. D. Yeager, F. Zhang, A. Zarrasvand, N. George, T. Daniel, and B. Otis. System-On-Chip WISP: A $9\mu\text{A}$, addressable gen 2 sensor tag for biosignal acquisition. In J.R. Smith, editor, *Wirelessly powered sensor networks and computational RFID (this volume)*, New York, 2013. Springer SBM.
40. H. Zhang, J. Gummeson, B. Ransford, and K. Fu. Moo: A batteryless computational RFID and sensing platform. University of Massachusetts Computer Science Technical Report UM-CS-2011-020.

Part II

Hardware Platforms

The Wireless Identification and Sensing Platform

Alanson P. Sample and Joshua R. Smith

1 Introduction

The wireless identification and sensing platform (WISP) is a family of sensing and computing platforms that are powered and read by off-the-shelf EPC class 1 generation 2 (C1G2) UHF RFID readers. Recent WISPs include a fully programmable microcontroller (the TI MSP430) that implements both application logic and the protocol’s MAC layer [4]. Arbitrary sensors can be added to the WISP; its microcontroller reads the sensors through either an analog to digital converter (ADC) or logic input pin and communicates the sensor data to the RFID reader using the C1G2 protocol. Although the initial motivation for WISP was sensing applications, the platform has been received enthusiastically by the security community as well. Because of the platform’s programmability, encryption and other security-related algorithms can be implemented and tested on the WISP. By contrast, most other RFID tags currently are fixed function, non-programmable “black boxes” which provide researchers no ability to control or modify their behavior. The WISP’s programmability and open architecture provides unlimited flexibility for experimentation and research and supports compute-intensive applications such as encryption.

A.P. Sample
Department of Electrical Engineering, University of Washington, Seattle, WA, USA
e-mail: alanson@ee.washington.edu

J.R. Smith (✉)
Departments of Computer Science and Engineering and Electrical Engineering,
University of Washington, Seattle, WA, USA
e-mail: jrs@cs.washington.edu

2 Wireless Sensing and Identification Platform

The WISP is manufactured as a printed circuit board (PCB), which offers a number of benefits when compared to traditional integrated circuit (IC) tag designs. A few of these advantages include low development cost, fast design cycles, and easy debugging and measurement of circuit parameters. The PCB implementation allows the flexibility to physically add and remove sensors and/or peripherals to create devices for new applications. In contrast, IC implementations offer the ability to customize components and decrease power consumption (yielding better range), as well as creating devices with a smaller form factor and at a lower cost when manufactured in high volume.

A block diagram of the WISP is shown in Fig. 1 and is similar in function to traditional IC RFID tags. The antenna is balanced by an impedance matching network and is fed into the RF power harvester. The radio frequency (RF) signal transmitted by the RFID readers is rectified into DC voltage to power the rest of the tag. The demodulator block converts the amplitude shift keyed (ASK) data that is superimposed on the RF carrier into a logic level stream of serial data. This extracted serial data is parsed by the MSP430 microcontroller (MCU) to receive downlink data from the reader. Uplink data is sent via the modulator circuit, which “backscatters” the signal by changing the antenna impedance. Finally, the microcontroller’s internal temperature sensor, as well as any external sensors, is powered and measured by the MCU.

Since the power consumption of the microcontroller, sensors, and peripherals is much greater than that seen in traditional passive RFID technology, the WISP duty cycles between active and sleep mode. In sleep mode, the WISP shuts down and reduces its current consumption to a few microamps and energy is accumulated by harvesting RF power over multiple EPC queries. Once sufficient voltage is obtained, the WISP polls sensors and communicates with the RFID reader.

Figure 2 depicts the WISP platform, made of a four layer FR4 PCB with components on both sides and an integrated dipole antenna. The WISP in its base configuration has several onboard sensors: a circuit for measuring the rectified

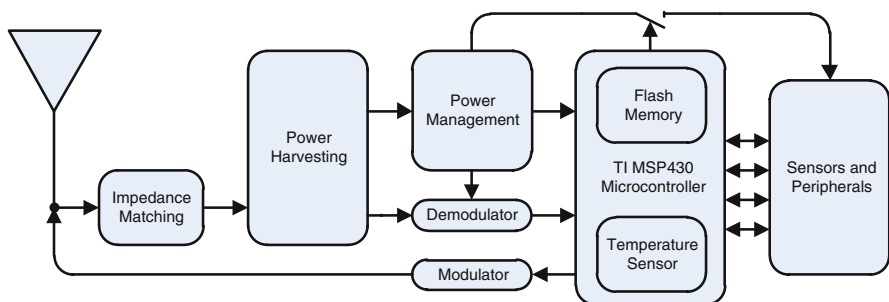


Fig. 1 Block diagram of the WISP




Fig. 2 Wireless identification and sensing platform, model G2.0




Fig. 3 Data logger version of the wireless identification and sensing platform, model number 4.1 DL

supply voltage, a temperature sensor, and a three-axis accelerometer. Small header pins expose all ports of the microcontroller for expansion to daughter boards, external sensors, and peripherals. Finally, a low current surface mount LED is included in the design. Figure 3 shows the data logger version of the WISP which has additional features, such as a larger microcontroller, a real time clock, external EEPROM, and an optional 0.1 F super capacitor for extended lifetime data logging applications.

2.1 Analog Front End

The defining characteristic of far field RFID systems is that tags can be read at a significant distance, generally on the order of 2–10 m. For passive RFID, this requires that the RFID reader transmits sufficient energy to power the tag at large distances. However, due to regulatory limits on the amount of power that can be transmitted and the path loss associated with electromagnetic propagation, there is very little power that actually reaches the tags. Therefore, the power harvesting circuit must maximize the operating distance by converting the very limited incoming RF power to DC power with sufficient voltage to activate the tag. The RF power received by the WISP’s dipole antenna is fed to the analog front end

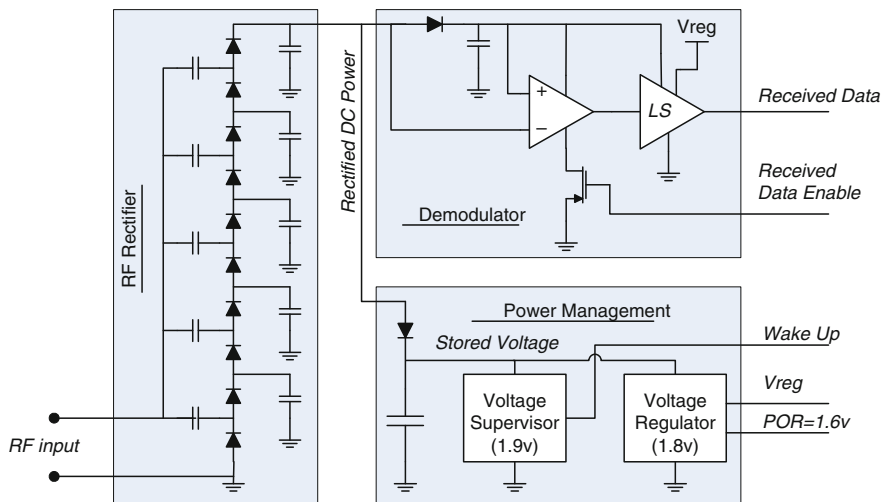


Fig. 4 Schematic of the analog front end

depicted in Fig. 4. A discrete matching network is used to provide the maximum power transfer from the antenna to the rectifier. RF Schottky diodes, specifically designed for 915 MHz low power application, were selected to make a five-stage voltage doubling circuit. This circuit converts the AC input signal to DC power which is fed into a storage capacitor.

For RF rectifiers of this type, the input and output impedances are not well isolated. Further confounding the problem, the output impedance of the rectifier is fairly high, an undesirable trait for any power source. This means that as the load on the rectifier changes the input impedance also changes, resulting in the analog front end becoming mismatched to the antenna. This leads to the problem of selecting values for the impedance matching network when it is not possible to guarantee constant input impedance. To determine the correct values for the matching network the operating cycle of the WISP must be taken into account. First, the WISP is most effective at storing harvested energy when it is in sleep mode, as the current consumption is minimal. Second, the WISP will spend most of its time repeatedly charging up to 1.9 V and then discharging to approximately 1.8 V. Thus, to determine the correct values, the WISP is placed into sleep mode and the impedance matching network is swept with a variable capacitor until 1.9 V is produced for the lowest possible input power. Stated another way, the key parameter for maximizing the read distance of the WISP is minimizing the quiescent current consumption so that the minimum operating voltage of 1.9 V (supervisor threshold) can be rectified with the lowest possible input power.




Fig. 5 Oscilloscope plot of the demodulated data extracting data from the RF waveform transmitted by the RFID reader

2.2 Demodulation and Modulation

The EPC Gen 2 standard defines that reader-to-tag communication uses ASK modulation on a carrier wave in the range of 902–928 MHz. When not transmitting data, the carrier waveform remains at a constant amplitude; when bits are transmitted, the amplitude of the carrier drops to at least 10% of its normal value and the phase of the carrier may be reversed. The duration of the continuous waveform between these low amplitude pulses indicates logical ones or zeros.

Figure 4 shows a schematic of the WISP’s demodulator circuit. The output of the harvester is fed through the diode, which supplies power to the comparator and acts as a reference for the level shifter. A capacitor is used to filter out transients while allowing proper biasing at varying distance and received power levels. When activated, the current consumption of the comparator functions as a constant-current source, pulling current through the diode. In this way, the voltage drop across the diode is used as a detector, where current supplied by the harvester (high amplitude RF modulation) results in positive voltage and a lack of current (low amplitude RF modulation) yields negative voltage. The comparator is used to generate a rail-to-rail logic-level waveform, and the level shifter converts the unregulated logic level to the regulated logic level. It is important to optimize current consumption and speed when choosing a comparator. Further savings can be achieved by disabling the comparator when there is insufficient voltage to start up the MSP430.

An example of a demodulated signal is shown in Fig. 5. This oscilloscope plot shows the 915 MHz RFID waveform and the resulting demodulated signal. Note that

time frame is 20 s per division and thus, the individual cycles of the 915 MHz carrier are not visible. However, the ASK modeled data is visible as gaps in the carrier and enveloped signal.

RFID tags do not actively transmit radio signals. Instead, they modulate the impedance of their antenna which causes a change in the amount of energy reflected back to the reader. This modulated reflection is typically called backscatter radiation. In order to change the impedance of the antenna, a transistor is placed between the two branches of the dipole antenna. When the transistor conducts current, it short-circuits the two branches of the antenna, changing the antenna impedance. In the nonconducting state, the transistor has no effect on the antenna and thus, the power harvesting and data downlink functions occur as if it were not present. This impedance modulation is currently implemented with a 5 GHz RF bipolar junction transistor, which allows for effective shunting of the 915 MHz carrier wave.

2.3 Digital Section and Power Conditioning

Since the power available to RFID tags is extremely limited, careful component selection must be made to minimize current consumption. As advances in IC manufacturing now allow discrete components with less than 1 μA of current consumption and operation at 1.8 V, it is now possible to construct working, wirelessly powered RFID tags with discrete components.

The general-purpose computation capabilities of WISP are provided by an ultra-low power microcontroller. This 16-bit flash microcontroller, the MSP430F1232, can run at up to 4 MHz with a 1.8 V supply voltage and consumes approximately 600 μA when active at those frequency and voltage settings. Of particular interest for low power RFID applications, the MSP430 has various low power modes. Its minimum RAM-retention supply current is only 0.1 μA at 1.5 V. The device provides over 8 kilobytes of flash memory, 256 bytes of RAM, and a 10-bit, 200-kilo-samples-per-second ADC. The low power consumption of this relatively new device is a critical factor in enabling use of a general-purpose microcontroller in passive RFID systems.

Another critical design consideration is operation with uncertain power supply conditions. Because the available RF power varies greatly throughout device operation, supervisory circuitry is necessary to wake and sleep the device based on the supply voltage level. WISP uses a 1.9 V supervisor and a 1.6 V power-on-reset to control device state and reset the microcontroller, respectively. The supervisor provides roughly 100 mV of headroom on the storage capacitor above the 1.8 V of regulator voltage. This serves to buffer the supply voltage from dropping below 1.8 V, due to the large power consumption of the microcontroller in active mode.

3 Firmware and Power Management Algorithm

The WISP is essentially a software defined RFID tag, which uses the MSP430 to implement the EPC C1G2 protocol and performs sensing and computation tasks. There are significant challenges when developing applications on the WISP as compared to battery powered embedded systems. Primarily, there is no guarantee that a given task can be completed before running out of power. Although the voltage supervisor provides headroom above 1.8 V, the rate at which the energy stored in the supply capacitor is consumed is directly affected by the design choices of the programmer. Failure to properly manage sleep cycles when the WISP harvests energy or inefficient coding practices can result in poor performance. The WISP software can be described on three levels. At the lowest level is the power management algorithm, which is responsible for managing the device state, including sleep versus active modes. Built on that is the communication layer, which enables bidirectional communication by sampling downlink data bits, implementing a Gen 2 state machine, and generating uplink data bits. The third level is the application layer where users implement custom functions and encode data in the appropriate EPC packets.

3.1 Power Management Algorithm

Meeting the low power requirements of passive RFID tags requires that the MCU consumes, on average, as little power as possible. As mentioned previously, this is achieved by duty cycling between active and low power sleep states. The key is that the WISP receives a constant amount of power as defined by the Friis path loss for a set distance. When the WISP is in active mode the power consumption far exceeds the power harvested. However, when the WISP is in sleep mode, the total current consumption of all the circuits is a few microamps and there is a net power gain which charges the storage capacitor. Therefore, duty cycling does not simply yield lower power consumption; it represents two different states, power harvesting and active operation.

The state diagram for the power management layer is shown in Fig. 6. State transitions are primarily driven by hardware interrupts from the voltage supervisor, which indicate if there is sufficient energy stored for operation. Initially, the WISP is away from a RFID reader and is in a power down state. When the WISP is brought within range of a reader, it begins to harvest power and the voltage across the storage capacitors begins to rise. At approximately 1.6 V the MSP430 powers up in a reset state and begins executing code. Since this event is not driven by the supervisor, it is important that the code enters sleep mode (LPM4) as quickly as possible in order to repeatedly avoid browning out on start up. Once in LPM4, the WISP waits for sufficient voltage (1.9 V), as indicated by the supervisor interrupt. Next, the state machine transitions to the application layer, which performs user-defined functions,

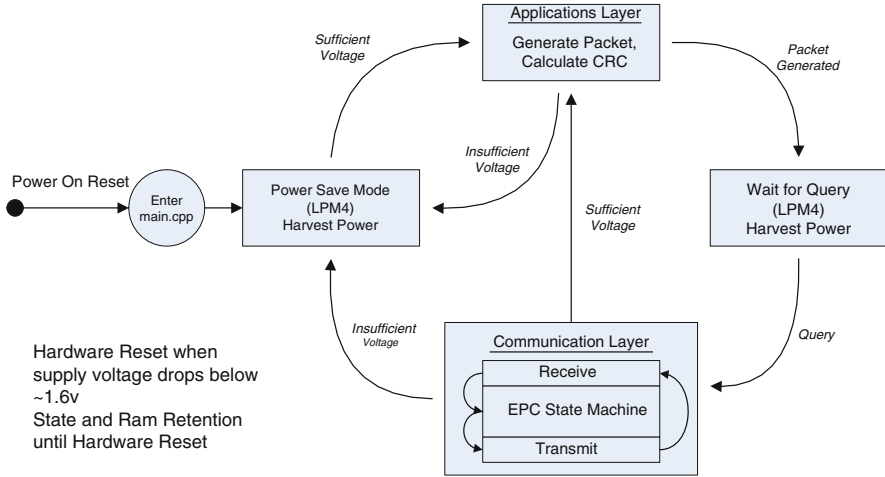


Fig. 6 State diagram of the power management algorithm for the WISP

such as sensor measurements. Here, an EPC packet is generated and the WISP sets up and waits for a commutation interpret which indicates the beginning of an EPC packet. In the communication layer, the WISP processes the incoming data, executes the EPC Gen 2 protocol, and transmits its response. While not shown in Fig. 6, the communication layer often reports the same data twice to increase communication reliability.

3.2 Communication and Application Layers

A considerable challenge when programming the MSP430 involves meeting the timing constraints of the EPC protocol while still maintaining a low clock frequency. RFID tags that have custom state machines are designed at the hardware level to receive and send using the EPC protocol. The general-purpose MSP430 must be carefully tuned to perform EPC communication, both for receiving and transmitting data. In particular, a mix of C and assembly language is used where the C code maintains ease of configurability for the firmware for different sensor applications and the assembly code allows fine-grained control of the timing of the MSP430 for EPC communication.

As previously described, the demodulator envelops and thresholds the phase-reversed amplitude shift keyed (PR-ASK) signal from the reader into a serial data stream representing the data bits 1 and 0 as long and short pulses, respectively. To interpret data from the reader, the MSP430 uses the periodic edge of the waveform as a hardware interrupt, and then during the interrupt service routine resamples the




Fig. 7 Oscilloscope scope plot of the WISP responding to EPC queries along with its rectified voltage

bit line to detect a 1 or 0 during the differentiated part of the waveform. This data is quickly shifted into memory before repeating this process. To detect the end of a transmission, a timer is refreshed during each bit. When bits are no longer received the timer expires, the packet is interpreted and, if appropriate, a response is sent to the reader. A detailed description of how the WISP uses and implements the EPC specification is described in Sect. 2.3. Figure 7 shows a set of EPC queries and responses along with the charge/discharge cycle of the WISP. Since the operating voltage range of the WISP occurs between 1.9 V and 1.8 V the rectified voltage appears to be nearly constant. In actuality, the WISP enters active mode at 1.9 V, consumes the energy in the storage capacitor until approximately 1.8 V, and then enters a sleep state and harvests power until 1.9 V is reached. This duty cycling can be seen in the packet-transmitted plot. Here, the WISP does not respond to every packet sent by the reader, instead it spends most of its time in a sleep state.

Performing application level tasks, such as sensor measurement, is generally done in tight conjunction with the EPC protocol. In this scenario, the completion of a receive/transmit cycle triggers the application layer to immediately take a sensor measurement, generate the desired EPC packet, and setup for a query. This protocol centric approach works well for sensor-driven applications where data is requested from the RFID tag at regular intervals. However, applications which leverage the wirelessly powered computing capability of the WISP benefit from a loose coupling with the communication layer.

4 Power Budget

One of the significant challenges of incorporating microcontrollers, sensors, and peripherals into passive RFID technology is the ability to manage the large power consumption of these devices. For example, the MSP430F1232 running at 3 MHz consumes approximately 470 μA at 1.8 V. The resulting power consumption is significantly larger than typical passive RFID tags. Under these conditions the harvester cannot continuously supply power to the WISP during a single reader query. One method to overcome this challenge is to use a large storage capacitor (on the order of 10 μF) to accumulate charge over multiple EPC queries. Once sufficient voltage is obtained, the WISP can operate in a burst mode, polling sensors and communicating with the RFID reader. This approach of duty cycling is often used in low power applications; however, this presents a challenge for RFID networks when the WISP is not necessarily able to respond to each reader query. The next section examines the issues related to powering the WISP from three perspectives. First is the received RF power required to turn on the device, the second is the operating duty cycle based on input power, and the last is the energy needed in the storage capacitor for active operation of the microcontroller and additional sensors.

4.1 Turn-On Power Requirement

In the presence of the RFID reader, the WISP's RF rectifier will charge the storage capacitor until the power input to the device equals the power lost due to quiescent current.

$$P_{\text{in}} = P_{\text{loss}} \equiv V_{\text{rectified}} \times I_{\text{loss}}. \quad (1)$$

Thus, a key parameter for maximizing the read distance of WISP is minimizing the quiescent current consumption so that the minimum turn on voltage of 1.9 V (supervisor threshold) can be rectified with the lowest possible input power. In order to characterize the system, a network analyzer was used to inject a continuous 915 MHz waveform into the antenna ports of the WISP. Figure 8 shows the resulting plot of rectified voltage and output power versus input power when the WISP is in sleep mode. Rectified voltage was measured with the WISP in sleep mode (only quiescent current draw) and shows the minimum input power needed to start operation. After the 1.9 V supervisor threshold has been met, the rectified voltage continues to increase with input power, until the overvoltage protection diode activates at 5.4 V. In the actual implementation of the WISP, the MPS430 activates at 1.9 V and starts consuming power. Thus, the rectified voltage never rises above the supervisor threshold. Using the minimum input power needed for activation from Fig. 8, the expected operating distance for the WISP can be calculated with the logarithmic form of the Friis path loss formula [Eq. (2)], with a term for polarization loss included.




Fig. 8 Rectified voltage (*left scale*) and output power at 1.9 V (*right scale*) are plotted verse input power measurements using multi-meter and network analyzer for RF signal insertion into the antenna ports

$$P_R = P_T - 20 \log \left(\frac{4\pi d}{\lambda} \right) + G_T + G_R - L_P. \tag{2}$$

The transmit power of the reader $P_T = 30$ dBm (which is equivalent to 1 W). Its center frequency is 915 MHz, corresponding to wavelength $\lambda = 0.33$ m. The transmit antenna gain $G_T = 6$ dBi (this yields an effective isotropic radiated power of 4 W EIRP, the United States’ regulatory limit for this ISM band). The receive antenna gain $G_R = 2$ dBi (the standard gain figure for a dipole antenna), and the polarization loss $L_P = 3$ dB. Loss L_P occurs because only half of the power transmitted from the circularly polarized transmit antenna is received by the linearly polarized receive dipole antenna. Using the operating thresholds of -9.5 dBm from Fig. 8, Eq. (2) predicts a maximum operational range of 4.3 m.

4.2 Duty Cycle

While rectified voltage (rather than power) determines the maximum achievable range, the operational duty cycle (percentage of the time WISP can be active) is determined by the amount of rectified power. In practice, the rectified voltage will typically remain near the threshold voltage (1.9 V). This is due to the operation

of the supervisor, which transitions the WISP from sleep to active mode, resulting in the consumption of power whenever the stored voltage exceeds this operating point. Therefore, it is important to characterize the output power of the harvester at 1.9 V. Figure 8 shows the result of output power verse input power at 1.9 V. This is accomplished by fixing the output voltage at 1.9 V using a power supply and measuring the amount of current that is supplied by the WISP. Then, the duty cycle of WISP (percentage of the time in active mode) is estimated as the ratio of rectifier output power to WISP active power consumption:

$$\frac{P_{\text{out}}}{P_{\text{active}}} = \frac{T_{\text{on}}}{T_{\text{on}} + T_{\text{sleep}}} = \text{Duty Cycle.} \quad (3)$$

In this equation, P_{out} is the output power of the WISP, P_{active} is the active power consumption, T_{on} is the time in active mode, and T_{sleep} is the time in sleep mode. For example, the power rectified at 0 dBm is $310 \mu\text{W}$. Dividing this value by the active power consumption ($1.8 \text{ V} * 600 \mu\text{A} = 1.12 \text{ mW}$) yields a duty cycle of 27 %. This agrees well with experimental values which are presented in Sect. 5.

4.3 Active Energy Consumption

Since the rectifier cannot supply enough power for continuous operation, it is important to quantify the amount of energy that needs to be stored in order to power the WISP during active periods. During one EPC Gen 2 communication cycle, the complete WISP (not just the microcontroller) consumes on average $600 \mu\text{A} * 1.8 \text{ V} = 1.08 \text{ mW}$. A single query takes 2 ms including reader and tag communication. Using the expression for the energy stored in a capacitor ($E = \frac{1}{2}CV^2$, with $C = 10 \mu\text{F}$), the amount of voltage headroom needed above 1.8 V is 116 mV, resulting in a total minimum voltage threshold of 1.91 V for a complete packet transmission. It should be noted that the MSP430 will operate down to 1.7 V, even though this value is below the specified supply voltage. However, operation is not guaranteed; it has been observed that the digitally controlled oscillator (DCO) can begin to slow down. Thus, it is not recommend that the designer rely on the extra 100 mV of headroom below 1.8 V. In the case of the previous example, the use of 16 mV out of specification headroom (1.90–116 mV) has proven to give reliable results.

The same method for calculating the required stored energy can be used when selecting sensors for the WISP platform. Sensor tasks and packet generation are generally done prior to the EPC query. However, it is reasonable to assume that when performing sensor applications the MCU will exhibit similar current consumption. Inequality (4) expresses an energy feasibility condition for a particular sensor; the energy required to read the sensor must not exceed the usable stored energy. This expression can be used to calculate the capacitor size and voltage headroom required

to operate a particular sensor, which in turn determines the range at which the sensor can be operated.

$$V_{\text{dd}}(I_S + I_W)T \geq \frac{1}{2}C(V_{\text{rec}}^2 - V_{\text{dd}}^2). \quad (4)$$

The current consumption for the sensor and WISP are I_S and I_W , respectively; C is the capacitance of the storage capacitor and T is the total time of active operation. The rectified voltage is V_{rec} and V_{dd} is the required operating voltage. Assuming that the sensor has the same voltage supply as the WISP, $V_{\text{dd}} = 1.8$ V. The left-hand side of inequality (4) represents energy consumed by the sensor and WISP during one measurement. The right-hand side represents usable stored energy above V_{dd} , the minimum operating voltage of WISP. Inequality (4) makes it clear that the limiting factor when selecting sensors is not only the current consumption (which determines power) but also the total required execution time of the sensor and WISP (energy rather than power).

5 Experimental Results

Figure 9 shows experimental results of the WISP performance: rectified output voltage, tag responses per reader query, and the rate of tag-to-reader packet errors are plotted versus received power (dBm). The experimental setup consisted of an EPC Gen 1 RFID reader driving a 6dBi circularly polarized patch antenna. The reader's antenna and WISP were placed one meter apart and one meter above the ground to minimize multipath effects. An adjustable attenuator inserted between the reader and its antenna was used to vary the power transmitted to the WISP. Finally, Eq.(2) is used to calculate the path loss over the one meter separation between the WISP and RFID reader. Thus, the WISP received power is defined as reader transmit power (1 W), minus variable attenuator, minus transmission path loss. It should be noted that the 1 W source represents peak output power of the RFID reader, while the average output power (not considered here) is highly dependent on reader transmission rate and the specific implementation of the EPC Gen 1 protocol.

To measure rectified output voltage, the WISP is placed in its low power state and voltage is averaged over a ten second interval using an oscilloscope. This is necessary to account for the variation in output power as the reader implements the EPC protocol. The resulting plot shows the WISP turns on with a peak received power level of -5.9 dBm, which is significantly more than the average power level of -9.5 dBm measured with the network analyzer in Fig. 9. In order to verify that this difference in turn-on threshold is caused by lower average power in the experimental setup, the RFID reader was replaced with a 915 MHz, 1 W continuous wave source and the turn-on power was found to be -8.7 dBm. The 0.8 dBm difference between the continuous source and the network analyzer is thought to be due to impedance mismatch between the dipole and the analog front end of

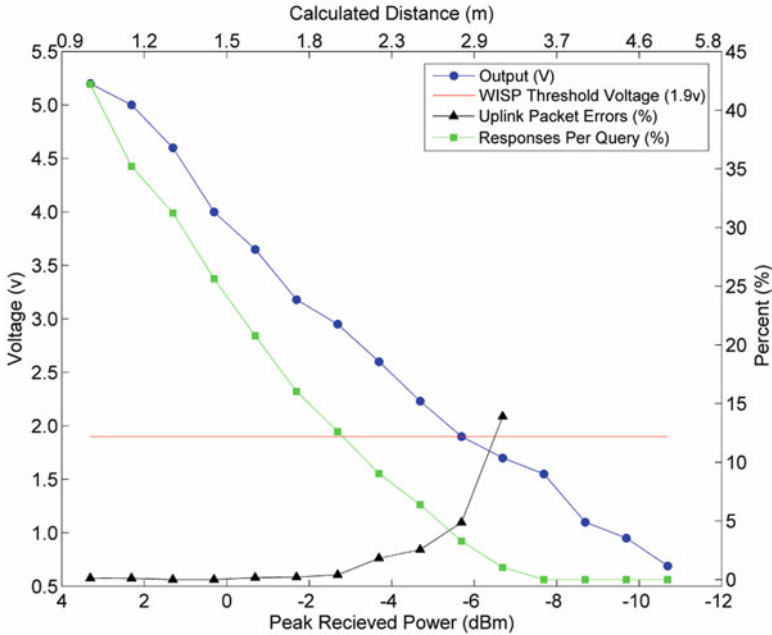


Fig. 9 WISP performance: harvested voltage, uplink packet errors, and responses per query as a function of input power

the WISP as well as antenna non-idealities. The 2.8 dBm difference between the continuous source and the RFID reader is then due to lower average power output by the reader.

The plot of tag responses per query shows the number of successful tag responses received by the reader normalized over the total number of queries made. This is equivalent to the operating duty cycle of the WISP and, as expected, is proportional to received power. The response rate drops to zero at -7 dB because there is insufficient voltage for operation. At 0 dBm input power, Sect. IV.B predicted an operational duty cycle of 27% using Eq. (3), which is close to the experimental value of 25% from Fig. 9. The reason that duty cycle (unlike turn-on voltage) is not diminished by the lower average power of the RFID reader is because duty cycle is normalized to the query rate of the reader. In other words, responses per query excludes times in which the reader is not transmitting.

The uplink packet error represents the percent of query responses made by the tag that are not correctly received by the RFID reader. Due to the limited data interface with the RFID reader selected for the experiment, the number of reader rejected uplink packets is not directly available. To collect this data, the WISP counts the number of query responses it has made and reports the current tally as data encoded in each uplink packet. When the RFID reader application software receives gaps in

the running tag response tally an error is recorded. Figure 9 shows that as received power decreases to the point at which sufficient voltage can no longer be rectified for operation, the uplink packet error rate increases. It is theorized that this system instability is due to the brown out state of the MSP430, along with the ring oscillator, used as the system clock, becoming detuned as the 1.8 V regulator drops out.

6 Sensors and Peripherals

Several types of sensors have been successfully integrated into the WISP platform: light, temperature, push-buttons, rectified voltage, and 3D acceleration. As a rule of thumb, sensors that operate as resistive transducers are good candidates for the WISP. These devices typically require a small amount of current and when placed in a voltage divider configuration they can be easily be measured with the MSP430's ADC. One example is the measurement of rectified voltage, which is easily accomplished with a voltage divider that scales the signal down to the range of the 1.8 V ADC. Alternatively, active sensors are much more demanding in terms of current consumption and required operating voltage. Real-world environmental noise often requires band-limiting the sensor signal to a few kilohertz or less. Such a filter may have a long time constant. Energy considerations should also be examined using Eq. (4) to choose a sufficient capacitor size for the system. For example, powering a 500 μ A sensor for 10 ms requires a great deal more energy than RFID communication; a 50 μ F capacitor charged to 1.9 V would be needed to provide enough energy to measure the sensor.

In order to enable environmental monitoring applications, the WISP was enhanced with temperature sensors [4]. The MSP430 does have an on-chip temperature sensor. However, its accuracy and power consumption is poor compared to signal chip solutions and the LM94021 low power analog temperature sensor from National Semiconductor was added. Figure 10 shows a time series of temperature measurements reported by WISP using the external temperature sensor. An inverted can of compressed air was used to generate a low temperature impulse. After the WISP's temperature sensor had recovered for about 30 s, a heat gun was used to generate a high temperature impulse. The LM94021 has 1.8°C accuracy and a Fluke thermal probe was used as reference for the system. The maximum error recorded between -20°C and 50°C by the WISP was 2°C.

Another example of environmental sensing has been demonstrated using the WISP with a photo resistor as a light sensor [6]. The WISP was mounted by suction cup to the inside surface of an exterior window in an office environment, with the sensor oriented inward. The light level was measured by the WISP over a 13 h period.

Figure 11 shows data collected using the three-axis accelerometer on the WISP. The Analog Devices ADXL330 MEMS accelerometer draws 200 μ A at 1.8 V. Due to the relatively high current consumption of these devices, continuously powering them would cripple the range of WISP. To overcome these high power requirements,

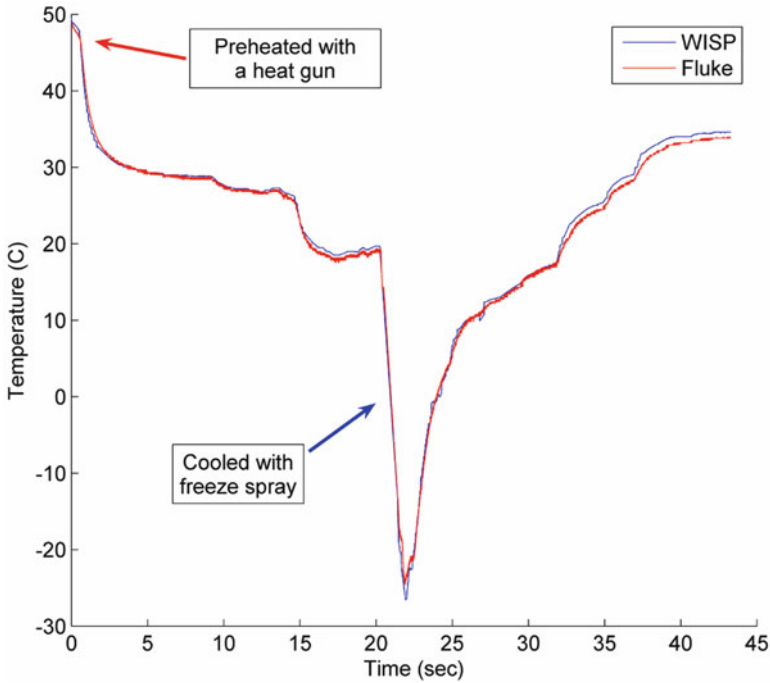


Fig. 10 Cold impulses are applied to WISP and a Fluke thermal probe and plotted over time

the sensor is only powered for a short period of time, just long enough to take a measurement. Provided that the sensor and conditioning electronics can stabilize sufficiently quickly, this allows for a wide range of sensors to be measured over UHF RFID. Powering this accelerometer, the WISP is able to provide accelerometer measurements at rates of approximately 1 to 50 samples per second, depending on range. After the measurement is taken and the data packed into the EPC ID, the WISP calculates the correct CRC. Then the “ID” is reported to the RFID reader and the information is then decoded in real time by the computer. Although the WISP’s accelerometer data rate is presently too low for some applications, it provides good absolute orientation data that is already suitable for some gaming and input device applications.

Figure 12 shows a demonstration of the WISP plus three axis accelerometer being used as a wirelessly powered, battery-free input device which was first demonstrated in [1]. Panels (a–f) show the WISP tilted at various orientations with respect to gravity and corresponding images of the planet Saturn in the background. In this application the WISP is interrogated by an EPC Gen 2 RFID reader which is connected to a host computer. When the WISP has enough power it measures its orientation relative to gravity using the three-axis accelerometer. This data is reported back to the RFID reader as EPC IDs, which are passed to the host computer.

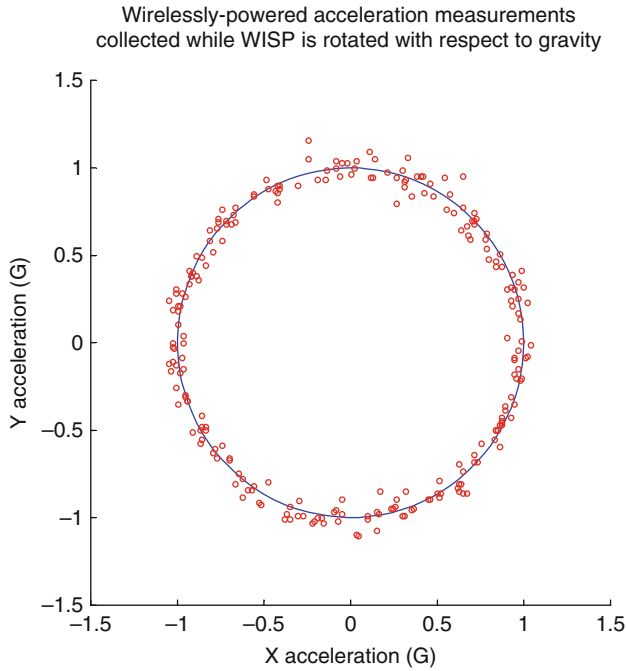


Fig. 11 Received measurement of the acceleration of earth’s gravity along two axes using the WISP enabled with a three-axis accelerometer

The host then decodes the packet, computes the WISP’s tilt relative to gravity from these acceleration measurements, and displays the corresponding orientation of the planet Saturn on the screen.

7 WISP Applications

The WISP is intended to be a research vehicle that allows people both inside and outside of the RFID community to explore new applications and usage models for RFID. Traditionally, RFID tag designers have been specialists in IC design. They have generally focused on innovating CMOS circuit blocks, such as RF rectification, power management, and low power state machines, with the goal of increasing tag read range. The process of manufacturing these custom-IC tags presents a significant barrier to entry when considering the high cost of software, servers, chip fabrication, and specialized testing equipment, not to mention the long fabrication cycles.

In contrast, WISPs are PCB-based, flexible platforms that allow rapid, low-cost prototyping of tag hardware and software. The full-featured microcontroller allows for fast code development with debugging support. Sensors and peripherals

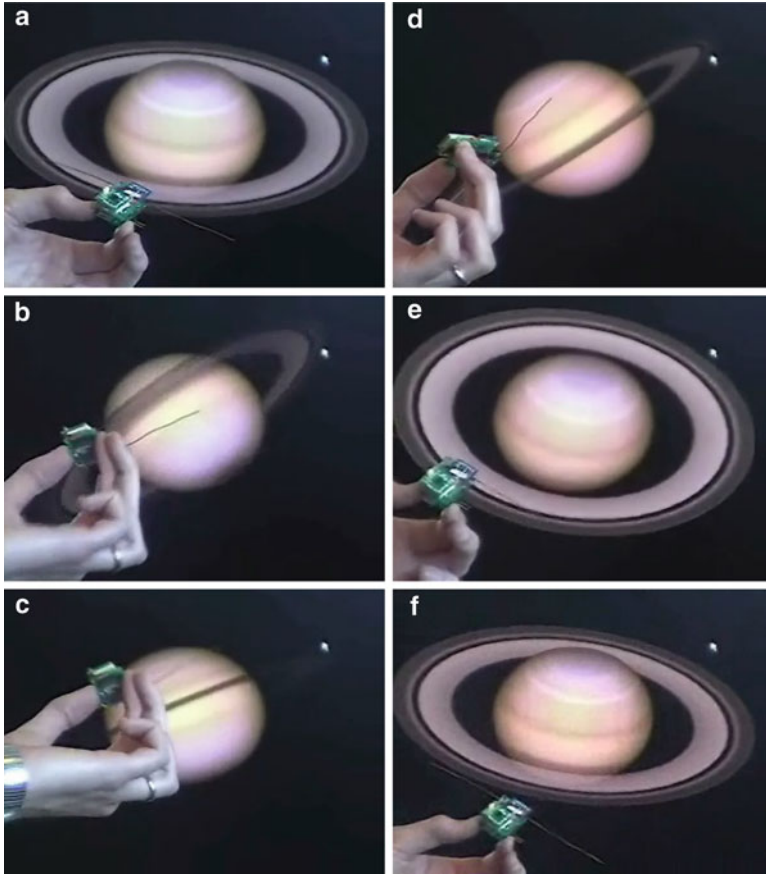


Fig. 12 Illustration of WISP accelerometer as wirelessly powered, battery-free input device. In images (a–f), the WISP is tilted at various orientations with respect to gravity. The WISP samples the X, Y, and Z channels of its on board accelerometer, encodes this data in an EPC ID (with dynamically computed CRC), and reports the ID to the reader. The reader decodes the packet and computes the WISP’s apparent direction of gravity from these acceleration measurements. For demonstration purposes, the planet is transformed to the same tilt angle relative to gravity as the WISP

can be easily added via the exposed headers or by using an optional daughter board. Testing equipment generally consists of an RFID reader and an oscilloscope. When compared to IC tags, probing and debugging circuit elements is easy and straightforward, as many of the signal lines are exposed by the PCB design.

We hypothesized that implementing the WISP as a flexible, PCB-based platform would allow people from a wide variety of fields to develop RFID technology. Since the completion of the first few prototypes, there have been a number of research efforts using the WISP to create new RFID applications. The following sections

summarize a few of the areas being investigated. First, security applications are discussed. Second the WISP is modified to create a passive data logging device capable of operation away from RFID readers. Third, the WISP's reconfigurability is used to rapidly prototype an RFID tag with a touch interface.

7.1 Security Applications

Conventional wisdom states that strong cryptographic algorithms are unrealistic for RFID considering the computational constraints and power issues of IC tags. As a result, various lightweight cryptographic protocols have been proposed and implemented. However, many of these protocols have serious vulnerabilities and were subsequently hacked or exploited. The computational power and flexibility of the WISP enables the realization of stronger, more conventional cryptographic techniques designed to enhance both privacy and security.

In [2], the WISP was used to demonstrate RC5-based symmetric cryptography for use on UHF RFID tags. The particular RC5 variant implemented uses a 32-bit word, 12 rounds, and a 16-byte secret key, which is stored in flash. While there were practical challenges in implementing RC5 on such a resource-constrained platform, the authors showed that with careful implementation strong cryptography is within the scope of UHF RFID. Additionally, their choice of RC5 was partly because RC5 can be efficiently implemented in both hardware and software, so their work can be used as a basis for IC implementations.

Even when strong cryptography is used, RFID is still susceptible to "man in the middle" attacks. For instance, RFID is widely used for access cards where an RFID-enabled employee badge uses a cryptographically strong challenge/response mechanism to open doors to a secured building. An attacker in this scenario does not need to break the encryption but only needs to generate the correct response to the RFID reader challenge. On the other hand, RFID-enabled credit cards have no encryption and the information transmitted via RFID is virtually identical to that printed on the card. Gathering this information no longer requires the conscious act of removing the card from a wallet and swiping the card through a magnetic reader. Consequently, thieves can steal the card information wirelessly, even while the card remains securely in the cardholder's wallet or purse.

In [3], the three-axis accelerometer on the WISP was used to implement a secret hand shake based authentication system to protect against "ghost and leech" and skimming attacks. When the user wants to authenticate a transaction or gain building access, they first perform a gesture with the card which unlocks the card and enables communication. The gesture could be a figure eight or any unique movement that the card would not experience in everyday activity. Only if this handshake is correct will the WISP unlock and transmit its ID to the reader. This approach leverages not only the computational power of the WISP but also its sensing capabilities to provide a level of security that is not possible using standard IC tags.




Fig. 13 WISP data logger with operational supercapacitor, model 4.0 DL

7.2 *Passive Data Logger*

There are many examples that demonstrated novel uses of passive RFID technology, which benefit from wireless, battery-free operation. However, these systems are inherently limited by the requirement of tag proximity to a reader for power and finite wireless range, due to RF path loss over distance. One particularly interesting class of applications involves a tagged item that travels between two reader-equipped locations but does not have reader proximity during transit. For example, this situation occurs during cold chain transport of food and chemicals between warehouses. One may be interested in tracking the temperature or vibration of goods during transit where there is no reader coverage.

To enable these applications, the authors in [7] have proposed a new tag device called a passive data logger (PDL). A PDL is a battery-free RFID tag with a large capacitor for energy storage. The PDL seamlessly recharges its capacitor when it is near a reader and uses the stored energy to measure attached sensors and log data to nonvolatile memory (NVM) when it is away from a reader. As a proxy for cold chain monitoring, a refrigerated milk container was instrumented with a WISP-PDL and monitored throughout its consumption. For this study, the WISP-PDL sampled and logged data in 10 s intervals and consumed $1.8 \mu\text{A}$ on average from a 1.8 V supply. Over the course of 24 h, the temperature and fill level of the carton was measured and written to memory. At the end of the study, the data was read from the WISP-PDL using the Gen 2 Read command showing the complete history of the milk carton. As the refrigerator acted as a faraday cage, the WISP-PDL harvested energy only when removed from the refrigerator but continued to sense when not directly powered by a reader. Building off of the work in [7], a fully integrated WISP-PDL has been implemented is shown in Fig. 13. The 0.1 Farad red super capacitor is clearly visible on top of the WISP. Additional features, such as a 1.8 V external 8k EEPROM and a 350 nA real time clock, have been added to expand the capabilities of the PDL. Presently this platform is still under development.

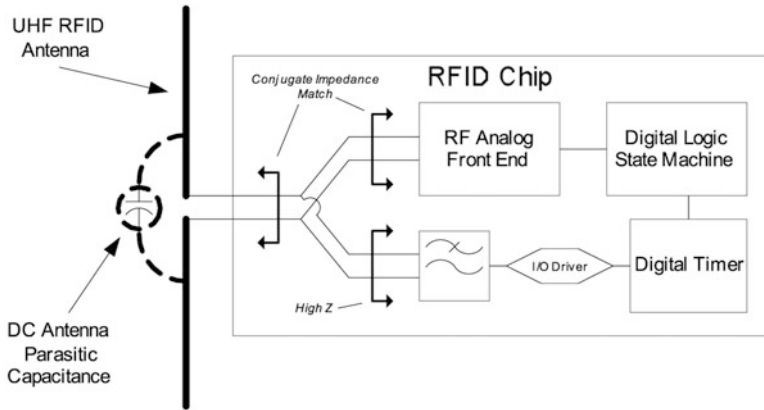


Fig. 14 Block diagram of the capacitive touch sensor-enhanced RFID tag. The antenna forms a sensing capacitor at low frequencies and its RC discharge rate is measured with the digital timer

7.3 Capacitive Touch Interface

The authors in [5] present a novel method for incorporating a capacitive touch interface into existing passive RFID tag architectures without additional parts or changes to the manufacturing process. This approach employs the tag’s antenna as a dual function element in which the antenna simultaneously acts as both a low-frequency capacitive fringing electric field sensor and also as an RF antenna. To demonstrate the feasibility of this approach the WISP was chosen because of its flexible and accessible design, which allows for modification of both firmware and low-level hardware as well as easy access to the RF antenna ports. The key is to take advantage of the frequency separation between the RFID reader’s carrier signal (HF or UHF bands) and the low-frequency RC time constant of a capacitive sensor (DC to LF). By separating these two signals, the antenna will operate as a dual band device: the UHF signal only interacts with the radiating structures and the low-frequency electric fields only interact with the capacitive sensing circuitry.

Figure 14 shows a block diagram of an RFID tag enhanced with a capacitive touch circuit. The UHF antenna forms a parasitic capacitor at DC (depicted as dotted lines), which is used as the touch-sensing element. The antenna is simply connected to the RFID chip through the two standard RF pads. As with traditional tags, the antenna’s impedance (at the RF design frequency) is matched to the complex conjugate of the analog front end of the RFID chip for maximum power transfer. In order to insure proper operation of the analog front end, a low-pass filter presents a high impedance path blocking the RF signal from the capacitive measurement circuits. Utilizing a first-order RC low-pass filter, resistance in the order of 200 kΩ and capacitance in the order of 5 pF meets the design criteria sufficiently.




Fig. 15 Image of the passive UHF RFID tag with a capacitive touch input and LED output for feed back

As an example of how an RFID antenna can form a capacitive sensing element, consider a simple dipole antenna. Inherently, at DC, the two halves of a dipole form a capacitor with energy stored in fringing electric fields. This capacitor can be charged by injecting a DC voltage across the two ports of the antenna. Even though the two halves of the antenna are at a different DC potential, the AC characteristics of the antenna are unchanged, and thus the dipole will continue to radiate energy (and through reciprocity receive energy).

The operation of an RFID tag enhanced with the touch sensors is as follows: Once the RFID tag is powered on and operating, the I/O port of the capacitive sensing circuit charges the positive side of the antenna. This DC signal is blocked from interfering with the analog front end of the IC due to the rectifier's AC coupling capacitors.

After the positive side of the antenna capacitor is charged to the regulated voltage of the IC, the direction of the port is changed from an output to an input and a digital timer is started to measure the rate of discharge. In order to register a user touch, the time at which the voltage on the capacitor reaches the threshold of the input port is compared to a calibrated or hard-coded threshold time value. Having the ability to change the trigger threshold for a touch event allows a single tag IC to be used in multiple antenna inlays with different absolute DC capacitances.

An EPC C1G2 tag has been prototyped to demonstrate the feasibility of capacitive sensing through the tag antenna. Sensor measurements can be reported to a commercial RFID reader in a variety of forms: sensor measurements can be encoded in the tag's ID, used to prevent/allow ID transmission, or retrieved through a read operation to user memory. Finally, the application layer decodes and displays tag sensor information reported by the reader.

The tag prototype is shown in Fig. 15. The antenna consists of copper foil on an FR4 substrate. The antenna is laminated with an insulator to prevent resistive loading of the capacitive sensor. Three fins between the dipole branches increase the sensitivity of the capacitive sensor. Finally, the WISP along with the RC filter, is connected to the antenna for capacitive sensing, signal processing, and communication.

8 Conclusion

In order to explore RFID applications beyond simple barcode replacement, this chapter argues the need for a reconfigurable, computation and sensor-enhanced device, capable of being wirelessly powered and interfacing with standard RFID technology. It is believed that such a device will lead to the discovery and rapid implementation of innovative RFID applications that cannot be developed by traditional CMOS tag design alone. Furthermore, the cost structure of traditional RFID applications is focused on manufacturing tags that can be sold at the lowest possible price, as a replacement for barcodes. However, it has been shown that by adding even a small amount of increased functionality to RFID tags, new markets can become available and the profitability of the tags can be substantially increased.

This chapter presents the design and performance of the WISP, a programmable, sensor-enhanced, passive UHF RFID tag. The WISP is powered by a standard, commercially available UHF RFID reader and implements the Electronic Product Code (EPC) Class 1, Generation 2 protocol.

The first half of this chapter describes the implementation of the major functional blocks of the WISP: RF harvester, downlink demodulator, power management, microcontroller, and uplink modulator. Additionally, the power management algorithm is presented along with a general overview of how the WISP uses the EPC Gen2 protocol to communicate data to and from the reader. This is followed by a detailed analysis of the power budget and performance of the WISP, which shows an ideal range of 4.3 m. Finally, several low power sensors that have been integrated into the WISP are presented: light, temperature, and acceleration.

More generally, WISP has proven the feasibility of powering devices with relatively large power consumption (such as a microcontroller and sensors) using only the RF energy transmitted by a standard RFID reader. The WISP is the first of a new class of battery-free, wireless sensing and computational devices.

Acknowledgments Figures 8–10 are ©2008 IEEE. Reprinted, with permission, from [4]. Figures 14 and 15 are ©2009 IEEE. Reprinted, with permission, from [5].

References

1. M. Buettner, B. Greenstein, R. Prasad, A.P. Sample, J.R. Smith, D.J. Yeager, and D. Wetherall. Demonstration: Rfid sensor networks with the intel wisp. In *6th ACM Conference on Embedded Networked Sensor Systems*, 2008.
2. H.-J. Chae, D.J. Yeager, J.R. Smith, and K. Fu. Maximalist cryptography and computation on the WISP UHF RFID tag. In *Proceedings of the Conference on RFID Security*, July 2007.
3. A. Czeskis, K. Koscher, J.R. Smith, and T. Kohno. RFIDs and secret handshakes: defending against ghost-and-leech attacks and unauthorized reads with context-aware communications. In *Proceedings of the 15th ACM conference on Computer and communications security, CCS '08*, pages 479–490, New York, NY, USA, 2008. ACM.

4. 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. *IEEE Transactions on Instrumentation and Measurement*, 57(11):2608–2615, Nov. 2008.
5. A.P. Sample, D.J. Yeager, and J.R. Smith. A capacitive touch interface for passive rfid tags. In *IEEE International Conference on RFID*, pages 103–109, April 2009.
6. J.R. Smith, A.P. Sample, P.S. Powledge, S. Roy, and A.V. Mamishev. A wirelessly-powered platform for sensing and computation. In *Ubicomp*, pages 495–506, 2006.
7. D.J. Yeager, P.S. Powledge, R. Prasad, D. Wetherall, and J.R. Smith. Wirelessly-charged uhf tags for sensor data collection. In *IEEE International Conference on RFID*, pages 320–327, April 2008.

SOCWISP: A 9 μ A, Addressable Gen2 Sensor Tag for Biosignal Acquisition

Daniel Yeager, Fan Zhang, Azin Zarrasvand, Nicole George,
Thomas Daniel, and Brian Otis

1 Introduction

Compelling applications in both the scientific and medical monitoring of biosignals have created a demand for wireless, unobtrusive sensors to collect this data. Example biosignals include temperature, blood pressure, heart rate, blood glucose level, and neural activity. In scientific applications, measurement of these biosignals helps researchers study complex biological systems, the effect of various diseases, and research treatments. In clinical settings, these signals are used by a doctor or patient to either detect disease at onset or help administer treatment.

Measurement of biosignals presents several challenges. Most importantly, the sensor must be unobtrusive to the user. This involves minimizing the size and weight of the sensor as well as maximizing the sensor lifespan. A means of wireless data collection is necessary for scientific research, where data should be available in real time, and for implantable medical biosensors where data is otherwise inaccessible. A number of solutions have been proposed including use of a small battery to power the sensor or an inductive link to power and communicate with the sensor. Unfortunately, battery-powered sensors suffer from short lifespan due to the size and weight constraints of the battery. Inductively coupled devices suffer from short wireless range (on the order of cm).

We propose the use of passive radio frequency identification (RFID) technology to address many of the challenges of biosignal sensors. Most importantly, passive RFID allows wireless, battery-free operation at meter ranges (see Fig. 1). This enables wearable and implantable biosensors with an unlimited lifespan, small size,

D. Yeager (✉) • F. Zhang • A. Zarrasvand • B. Otis
Department of Electrical Engineering, University of Washington, Seattle, WA, 98195 USA
e-mail: yeagerd@ee.washington.edu

N. George • T. Daniel
Department of Biology, University of Washington, Seattle, WA, 98195 USA




Fig. 1 Biosignal acquisition system

and sub-gran weight. Many passive sensor tags reported to date have employed simple ring oscillator temperature sensors with no protocol or addressability. However, realistic applications demand accurate processing of μV -level biosignals and compatibility with industry standard RFID protocols.

We present a fully passive 900 MHz RFID tag IC with addressability, full EPC Class 1 Generation 2 (Gen2) protocol compatibility, a $1.25 \mu\text{V}_{\text{rms}}$ integrated noise chopper-stabilized micropower sensor interface amplifier, and an 8b ADC. The communication range is 3 m with an off-the-shelf RFID reader, enabling previously impossible recording scenarios like in-flight recording from small insects. A significant improvement in performance beyond the state of the art was achieved by utilizing a novel self-calibrating on-chip frequency reference, subthreshold digital logic, a low-noise chopper amplifier, 8b ADC, and a unique chip ID generator.

2 Prior Work

A plethora of RFID-related research has appeared in the literature, predominately since Karthaus' [1] influential and widely cited paper in 2003. The recent technological feasibility of RFID, coupled with numerous application spaces and widespread commercial adoption, has galvanized academic research in this field. A number of full tag implementations have been presented in the literature, and they are grouped by feature set in Table 1.

2.1 Passive Tags with Sensors

The first category, which describes this work, represents fully passive tags with sensors, addressability, and Gen2-compliant protocol. This combination of features enables practical deployment because (a) low-cost COTS (commercial, off-the-shelf) readers can be used, (b) sensor data can be associated with the person

Table 1 Comparison of published tag features

Type	Author	Protocol	NVM	ID	Sensors	Passive
A	This work	G2	–	Y	8b ADC	Y
A	Sample [2]	G2 ^a	Y	Y	10b ADC	Y
B	Kim [3]	G2 ^b	Y	Y	Temp	N
B	Kocer [4]	N	–	3b	5b ADC	N
C	Nakamoto [5]	ISO	Y	Y	–	Y
C	Pillai [6]	ISO	Y	Y	–	Y
C	Karthaas [1]	ISO ^c	Y	Y	–	Y
C	Barnett [7]	G2	Y	Y	–	Y
D	Shen [8]	N	–	N	8b ADC	Y
D	Cho [9]	N	–	N	Temp	Y
D	Shenghua [10]	N	–	N	Temp	Y

^aPartial protocol support

^bSensor data not sent by Gen2 protocol

^cSpecific protocol not stated in publication

or animal that the tag is attached to, and (c) useful signals can be amplified, digitized, and retrieved from the sensor. Although [2] includes all of the required features for a biosensor tag, its use of a commercial microcontroller requires duty cycling to satisfy its high-power requirements. This significantly reduces the data rate as the wireless range increases.

2.2 Tags with Batteries

The second category comprises tags with batteries. These tags lack the key advantage of an RFID solution: they are limited by battery life. Furthermore, batteries add cost, weight, and safety concerns to realistic biosensing scenarios.

2.3 Tags Lacking Sensors

The third category includes tags with addressability but no sensors. These tags target conventional RFID applications and provide a useful design reference for the rectifier and communication circuitry.

2.4 Tags Lacking Addressability

The last category includes tags with sensors but no protocol/addressability. It is unclear how one would deploy this type of tag because without an ID, the sensor

data they report cannot be associated with the object to which the tag is attached. Furthermore, none of these tags include an ADC to digitize practical biosignals and are thus limited in their application.

2.5 *Tag Components*

In this section, we describe prior research in components necessary for a Gen2-compatible RFID tag.

2.5.1 **Oscillators**

Gen2 RFID tags require a local oscillator to clock the digital core. In particular, the tag backscatter specifications necessitate either a precise 1.28 MHz clock or a higher clock rate with a programmable divider. Most designs employ a current-starved ring oscillator, and efforts focus on reducing variation and power consumption. Ring oscillators suffer from high sensitivity to supply voltage, bias current, temperature, mismatch, parasitics, and process variation. Several oscillators with frequencies ranging from 1 MHz to 2 MHz have been presented; however, they all require manual trimming to achieve the precision required by the EPC protocol [11, 12]. Other works suffer from poor performance or excessively high frequency, which in turn increases the digital core power consumption [13, 14]. We demonstrate in Sect. 4.6 that an ultra-low-power tag clock can be synthesized via a ring oscillator, without trimming or tuning, by using a programmable divider and timing in the EPC protocol.

2.5.2 **Communication Logic**

Minimizing the power consumption of each tag circuit block, including the digital core that implements the tag's communication protocol, is critical in achieving high tag sensitivity. Zalvide et al. [15] presents simulation results comparing performance gains from various power reduction strategies for a Gen2 digital core. Unfortunately they target just one allowable link frequency and consequently may not work with commercial readers (which also only support certain subsets of the protocol). Ricci et al. [16] presents simulation results for a 2 μ W Gen2 digital core with cryptography and some analysis on clock frequency selection. They target a 2.0 MHz clock frequency, which cannot be guaranteed by the tag oscillator over process and temperature variation.

2.5.3 Rectifiers

There have been a number of in-depth analyses of UHF and microwave rectifier design. Much of the literature focuses on optimization of conventional Dickson charge pumps (voltage-doubling ladders) [17–20]. Mandal and Sarpeshkar [21] discuss fundamental physical relationships that link the operating bandwidth and range to technology-dependent quantities like threshold voltage and parasitic capacitances. Several papers have investigated nonconventional topologies. For example, Nakamoto [22] tune transistor V_t to maximize rectifier power conversion, achieving 36.6% efficiency. This work is replicated using floating gate PMOS transistors while optimizing for sensitivity instead of efficiency [23]. They achieve the best results to date: 1 V output voltage is reported at -22.5 dBm input power. This work is similar to patented “adaptive silicon” technology which uses floating gate V_t tuning [24].

3 System Architecture

The system architecture is shown in Fig. 2. One of the challenges of this work is the integration of sensitive instrumentation amplification onto a severely power-constrained platform that also suffers from tremendous supply voltage fluctuations and electromagnetic interference (EMI). Accurate signal amplification and digitization require precise supply and reference voltages. Ultra-low-power linear regulators, bandgap reference, and bias current generation provide a stable bias and supply for the chip (Sect. 4). Sensor input signals (e.g., EEG, EMG, thermocouple) are first amplified with the on-chip low-noise chopper-stabilized amplifier (Sect. 5). An 8b SAR ADC then digitizes the sensor data. The sensor data is associated with

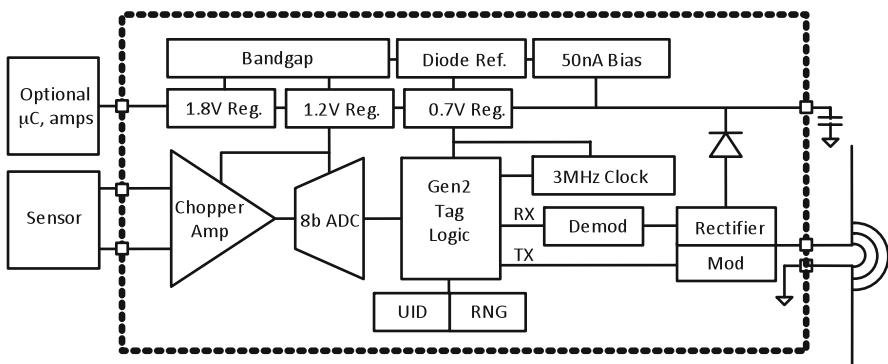


Fig. 2 Block diagram of the system

a person or animal by means of a unique tag ID (UID). This UID leverages process variation in the start-up configuration of an SRAM, thus eliminating the need for nonvolatile memory [25]. Random numbers (RN) are required in the Gen2 protocol for anticollision and (weak) encryption of reader-to-tag data. These are generated by sampling the (unpredictable) clock phase at the downlink baseband edges and passing it through an LFSR. Finally, the on-chip controller logic encodes the RN, UID, and ADC data into a Gen2-compatible packet in response to reader commands (Query/ReqRN, Ack, and Read, respectively) (Sect. 6). The UID and sensor data are available for real-time use by a PC through an Ethernet connection to the reader.

4 Analog Core

The analog core is responsible for generating accurate bias currents and voltages for the chip, which in turn enables accurate biosignal amplification and digitization. Total measured power consumption for the analog blocks is $1.2\ \mu\text{A}$. The various circuit blocks are described in the following sections:

4.1 Rectifier

The RF rectifier employs a 6-stage voltage-doubling charge pump topology. High sensitivity and efficiency are achieved by using zero- V_t diode-connected NMOS devices. Measured performance is plotted in Fig. 3. An off-chip L-match network transforms the impedance to $50\ \Omega$; alternately, this matching network can be easily absorbed into the antenna as in commercial designs.




Fig. 3 Measured rectifier efficiency versus output voltage and input power

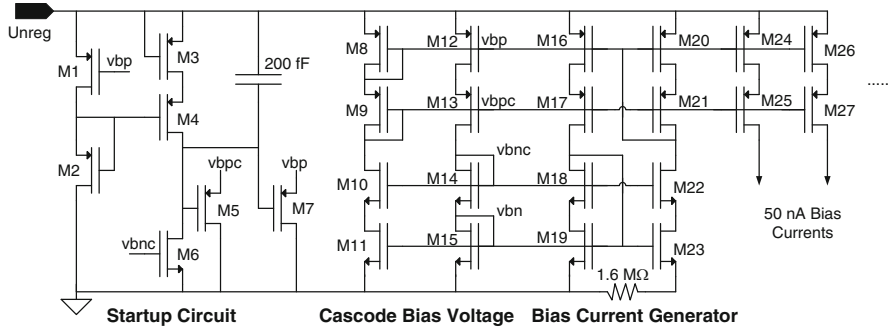


Fig. 4 Bias current generator schematic

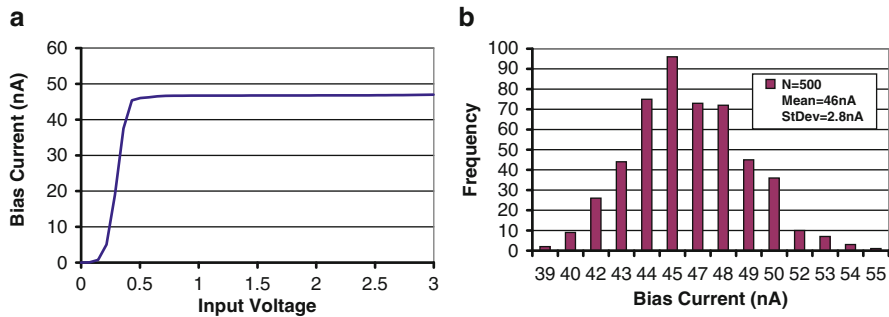


Fig. 5 Bias current generator performance (a) Simulated bias current versus supply voltage (b) Simulated Monte Carlo bias current distribution (process and mismatch)

4.2 Bias Current Generation

Bias currents for the chip are generated by a 45 nA V_{gs}/R reference, shown in Fig. 4. Despite the use of thick-oxide devices (which allow up to 3.6 V unregulated supply voltage), the bias generator starts reliably at 0.6 V. The high output impedance of the cascode devices maintains a constant output current of 45 nA from 0.6 V to 3.6 V as shown in Fig. 5a. Process and mismatch variation is minimized through use of relatively large device sizes as well as precision resistors. Simulated Monte Carlo variation is 2.8 nA (6%) without trimming (see Fig. 5b). Measured bias current consumption for the full analog core corresponded well with the anticipated current based on simulation of the bias current reference.

Many start-up circuits for bias current references consume significant static current. The start-up circuit shown in Fig. 4 uses MOS-bipolar pseudo-resistors (M2, M3) as an area-efficient means of creating very large resistances (e.g., >100 G Ω), which minimize static current. A 200 fF capacitor provides fast transient start-up when the circuit is powered on and prevents oscillation of the start-up circuitry. In addition, supply noise will not cause disturbances in steady-state operation because

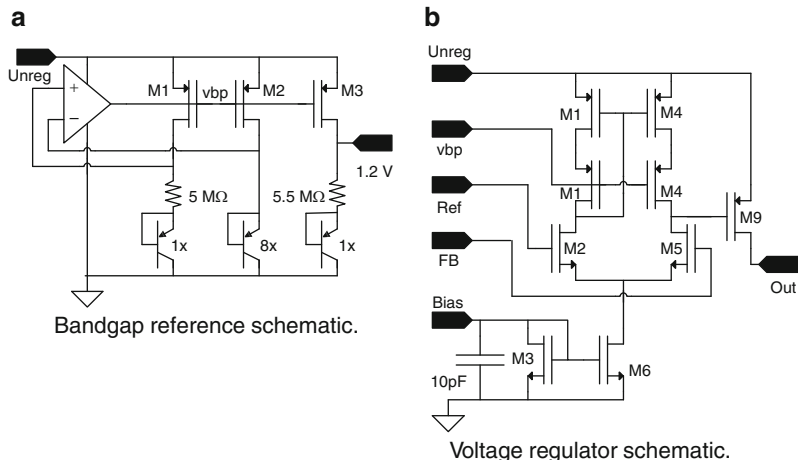


Fig. 6 Bias voltage and supply generation schematics (a) Bandgap reference schematic (b) Voltage regulator schematic

the gates of M5 and M7 are pulled to ground. Transistors M1, M2, and M4 ensure that the gates of M5 and M7 are fully discharged in spite of off-resistance variation across process corners.

4.3 Bandgap Reference and Supply Regulation

A stable reference voltage is obtained through an ultra-low-power bandgap reference, shown in Fig. 6a. The total supply current is 220 nA, and the measured output voltage is stable to within 4 mV of the nominal 1.2 V across 0–100°C.

Three low-drop-out linear regulators provide stable supplies for the 0.7 V digital core, the 1.2 V analog core, and an auxiliary 1.8 V supply for any off-chip ICs, respectively (Fig. 7). The 1.2 V and 1.8 V regulators utilize the bandgap voltage reference voltage, which provides a precise, temperature-independent voltage reference for the biosignal amplifier and ADC. The 0.7 V regulator is generated by sensing the transistor V_t . This creates a CTAT supply voltage that automatically compensates for process and temperature variation in the digital blocks.

The regulators employ a straightforward single-stage op-amp with 60 dB gain and 100 nA current consumption. The 0.7 V and 1.2 V regulators use unity-gain feedback, while the 1.8 V uses resistive feedback. On-chip mega-ohm resistors are used in 1.8 V regulator feedback network to limit the quiescent current. Low precision (but well-matched) resistors can be afforded because the ratio, not the absolute value, of two resistors sets the accuracy of the output voltage.

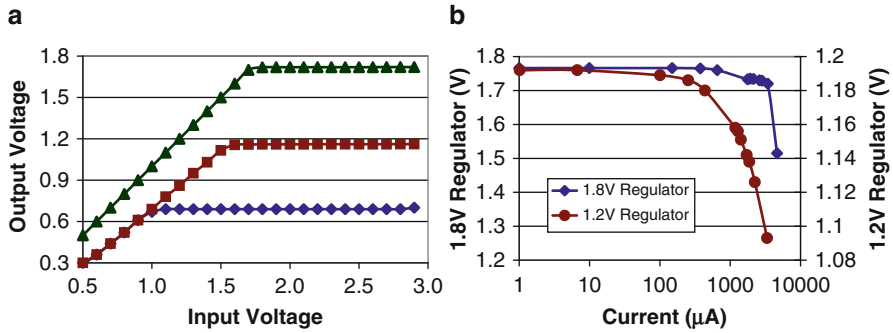


Fig. 7 Measured regulator performance (a) Line regulation for the 1.8 V (top), 1.2 V (middle), and 0.7 V (bottom) regulators (b) Load regulation

4.4 Demodulator

Reader-to-tag communication employs pulse-interval encoding (PIE). Specifically, the duration of the positive pulse width determines whether each bit is a zero or one. Positive pulses are delineated by a short negative pulse of duration PW, which is as small as 1.66 μ s. This sets the lower limit on the demodulator bandwidth to approximately 1.2 MHz.

The digital core measures and converts the positive pulse durations into data and clock. However, the received signal (the output of the rectifier) is not a suitable digital signal due to the finite time constants in the rectifier. The purpose of the demodulator is to recover a logic-level (rail-to-rail) signal from the rectifier output. The demodulator inputs are the rectifier output and a low-pass-filtered version of the rectifier output. The low-pass filter has an approximately 16 kHz bandwidth to filter out the negative pulses.

The demodulator schematic is shown in Fig. 8. The first stage provides differential to single-ended conversion. The second stage significantly boosts the signal swing in order to achieve logic-level voltages. Thick-oxide buffers with a low supply voltage (0.7 V) prevent crowbar current that could result from the limited slew rate of the low-power comparator. The differential pair is biased with 75 nA, and the common-source stage is biased with 300 nA which only draws current when the reader carrier is on. The unity-gain bandwidth of demodulator is set to 2 MHz. This ensures reliable recovery of the input signal (up to 1.2 MHz) under a wide range of temperature and process corners.




Fig. 8 Analog demodulator that is used to recover a logic-level baseband signal from the rectifier output

4.5 Modulator

RFID tags communicate with the interrogator by either absorbing or reflecting the RF carrier sent by the interrogator. A single transistor switch is used to modulate the tag reflection coefficient. A thick-oxide (IO) high- V_t device is used to prevent breakdown when the tag is near the reader (at which point greater than 10 dBm can be expected at the tag antenna). Experimental results demonstrate that the tag is downlink limited and the return loss with the modulator enabled is -0.6 dB.

4.6 Oscillator

Many tags use a 1.5 MHz clock, which requires the tag oscillator PVT stability to meet Gen2 timing specifications ($\pm 15\%$ for 640 kHz uplink) as the integer divider residual exceeds the allotted tolerance. Resistor trimming [12], bias current tuning [11], phase locking, and quartz references have been proposed to compensate for PVT variation but are prohibitive due to cost, power, and size constraints. We propose a 3 MHz temperature-stabilized ring oscillator, shown in Fig. 9, which lowers the divider residual such that PVT compensation can be performed by the integer divider. The oscillator consumes 260 nA from the 0.7 V digital supply. We take three approaches to improve stability. First, large device size and careful layout limit process variation to 13%. Second, a novel temperature compensation shown in Fig. 9 tunes the oscillator bias current by measuring and compensating the V_t temperature coefficient (measured, Fig. 10a). Third, the divider residual is centered at zero, which reduces the peak residual by a factor of 2 (Fig. 10b). As temperature

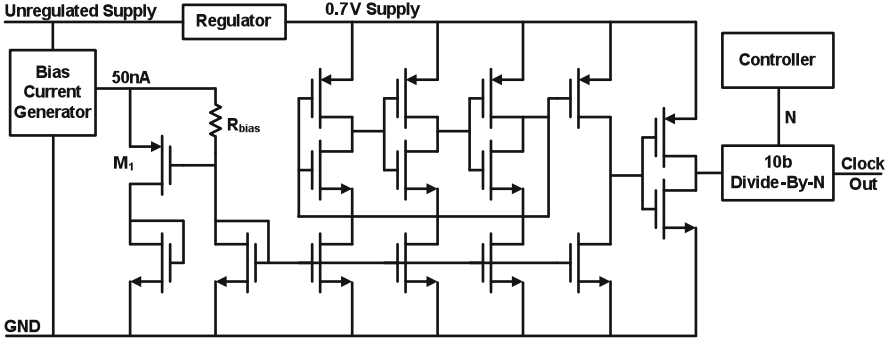


Fig. 9 Three-stage ring oscillator schematic and prescale divider

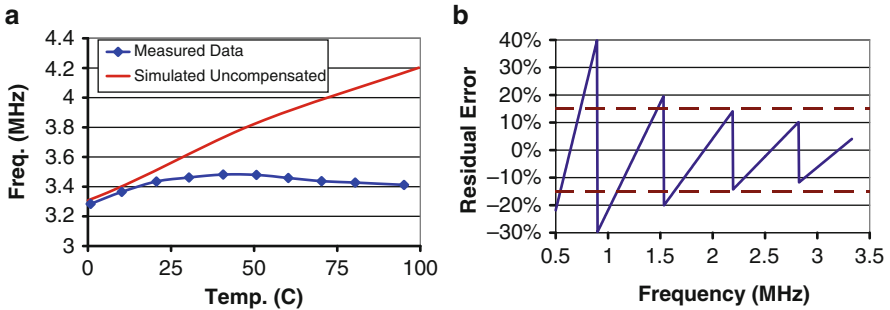


Fig. 10 Oscillator performance versus temperature and divider operation. The red dashed lines denote the acceptable residual error as per the EPC specification (a) Measured and simulated oscillator frequency versus temperature (b) Divider output residual error versus input (oscillator) frequency

decreases, f_{osc} decreases. The negative $\Delta V_{gs}/\Delta T$ coefficient increases the current through R_{bias} , providing the oscillator increased current to compensate for the temperature variation. Careful design of the transistor inversion coefficient and value of R_{bias} results in a first-order temperature coefficient cancellation (measured, Fig. 10a).

4.7 ADC

An 8-bit successive approximation register (SAR) ADC digitizes the amplified biosignals. Virtually no static current is consumed through the use of a discrete-time comparator. Complete testing and characterization of the ADC is presented in [26].

4.8 Unique ID Generation

In order to associate sensor data from a tag with the object that the tag is attached to, the tag must have a unique ID. Traditionally, the tag is manually programmed with a specified ID by the interrogator, which is stored in a nonvolatile memory (NVM) such as flash, EEPROM, or resistor fuse array. However, these NVM require large voltages and currents, which in turn degrade tag performance due to increased power consumption. We instead employ lithographic uncertainty and random dopant fluctuation to create a unique, random ID for each tag [25]. Specifically, the tag reads the power-on state of a 128-bit SRAM array. This technique allows reliable identification of ICs without explicit programming steps. On average, 95% of the bits are stable, and unstable bits are easily filtered out via application-level software. Note that unstable ID bits do not impair Gen2 protocol compatibility due to the use of random numbers as tag handles.

5 Biosignal Amplifier

5.1 Chopper-Stabilized Low-Noise Amplifier Design

Our multipurpose sensing tag is designed for a variety of sensor interfaces, including biosignal detection, thermocouple readout, and gas detection. These applications demand an extremely low noise floor ($< 2 \mu\text{V}_{\text{rms}}$ input referred) under a relatively low bandwidth ($< 1 \text{ kHz}$). When the signals of interest fall below a few hundred Hz, the dominating circuit noises shift from the thermal noise to $1/f$ and popcorn noise [27]. Excess low-frequency noise can undermine the system's signal-to-noise ratio (SNR) and cause errors in the measurement. As a result, we use a chopper-stabilized topology to suppress $1/f$ noise and offset that plague submicron CMOS processes.

Closed-loop chopper stabilization has been adopted recently [27–29] to suppress gain and sensitivity errors, as well as to prevent saturation due to amplifier offset. Among the recent implementations, [27] provides the best figure of merit so far. AC feedback is employed to ensure all signals entering the amplifier are well above $1/f$ noise corner. However, separate active input-biasing circuitry is used, and higher supply voltage is required due to single-ended approach.

As shown in Figs. 11 and 12, we employ a fully differential closed-loop architecture to ensure sufficient linearity and supply rejection. Operating transistors in the subthreshold region enable the use of a power-efficient telescopic-cascode op-amp topology under low supply voltages. Signal up-conversion occurs at the gate of the input transistors, which are biased in weak inversion to maximize the transconductance. We introduce a novel dual-feedback technique to simultaneously set the mid-band gain of the amplifier through C_{fb} and bias the amplifier's input node through high-resistance pseudo-resistors. Chopper switches are included in

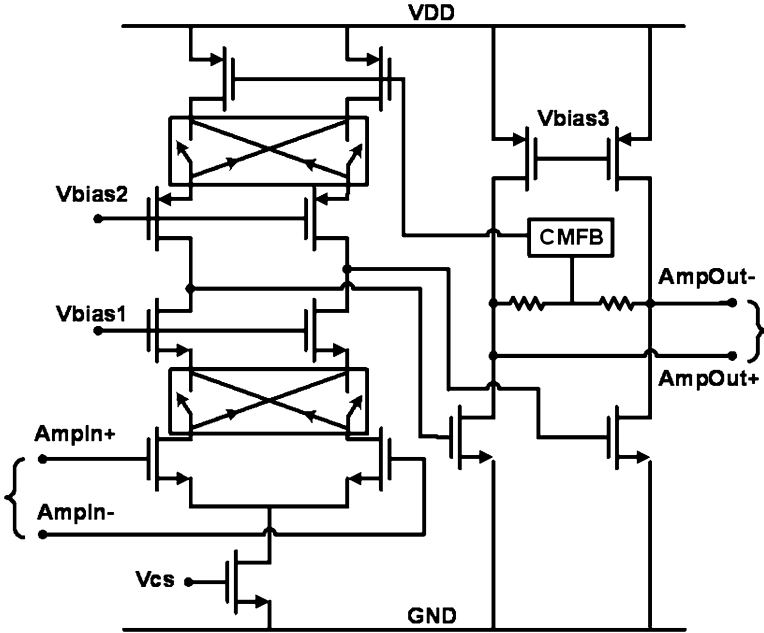


Fig. 11 Schematic of fully differential chopper-stabilized low-noise amplifier. Compensation capacitors and nulling resistors are not included for simplicity

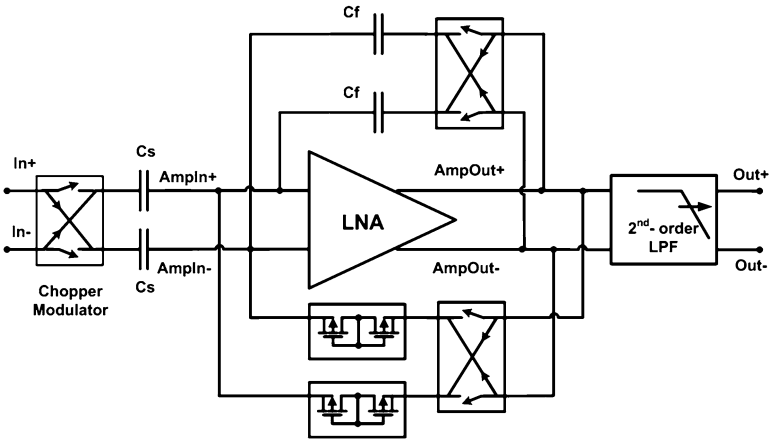


Fig. 12 Chopper-amplifier architecture

both the signal and biasing paths to not only guarantee negative feedback around the amplifier but also avoid using additional input-biasing circuitry as in [27]. We realize the chopper modulator with a minimally sized CMOS transmission gate to minimize charge injection. The input capacitance (C_{in}) is 15 pF. When modulated

with a 10 kHz chopper clock, the input impedance ($1.06\text{ M}\Omega$) is high enough to avoid loading the electrodes for biomedical applications. The ratio of C_{in} and C_{fb} establishes a 40 dB mid-band gain. C_{fb} is sized slightly smaller (140 fF) to take into account the addition of parasitic and switch capacitances.

Two additional sets of chopper switches are added in the first stage of the amplifier: one set of switches is placed at the drains of the input transistors to demodulate the ac signal down to baseband and modulate the input offsets up to the chopper frequency; another pair is placed at the drains of the PMOS current source to modulate their flicker noise up to a higher frequency. At the output of the amplifier, the signal returns to baseband while the offsets and flicker noise are modulated up to high frequency and then filtered by the amplifier's 2nd stage. The 2nd stage is implemented as common source to increase the output swing under low supply voltages. The output is then fed back to the summing node at the input of the amplifier after being modulated up to the chopper frequency. In order to avoid large passive devices, we implemented continuous-time tunable Gm-C filters to reduce ripple at the output of the amplifier. The input-referred noise from the ripple filter (Gm-C filters) is designed to be negligible. The six achievable bandwidths of the Gm-C filters are logarithmically spread between 150 Hz and 400 Hz. The tunability of the filters is realized through digital control of the transconductor current.

5.2 Low-Noise Amplifier Measurement Results

Figure 13a plots the gain magnitude with chopping on and off. The mid-band gain for both cases is approximately 38.5 dB. When the chopper clock is off, the amplifier operates as a conventional AC-coupled amplifier and has a high-pass corner of 0.2 Hz. Amplification is preserved down to DC when chopper stabilization is enabled. The tunable low-pass corner is set to 230 Hz in this measurement.

Figure 13b illustrates the input-referred noise of the amplifier with chopping on and off. Low-frequency spot noise is reduced by more than a decade when the chopper is enabled. The measured integrated noise from 0.05 Hz to 100 Hz is $1.25\text{ }\mu\text{V}_{\text{rms}}$ when the chopper switches are on, compared to $4.46\text{ }\mu\text{V}_{\text{rms}}$ when the chopper switches are off.

6 Digital Core

A high-level block diagram of the digital core is shown in Fig. 14. The *Receive* block performs clock and data recovery (CDR) on the PIE signal. A *Packet Parse* block decodes EPC commands and stores relevant information for the *Controller* and *Prescaler*. The *Controller* block decides what packet to send after receiving a packet

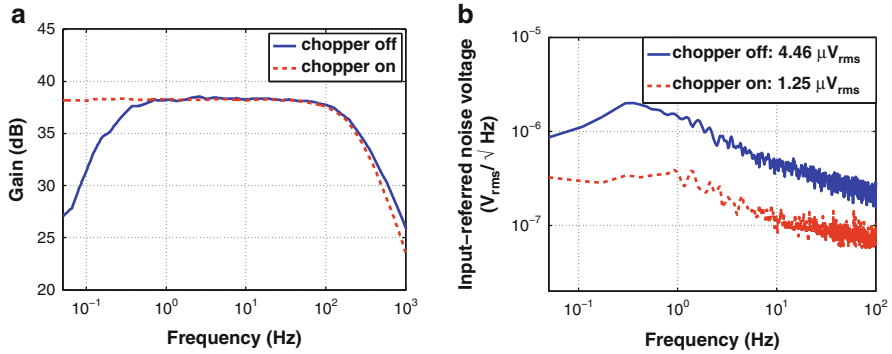


Fig. 13 Measured transfer function and noise plot of the low-noise chopper-stabilized biosignal amplifier (a) Chopper-amplifier gain magnitude plot (b) Chopper-amplifier noise plot

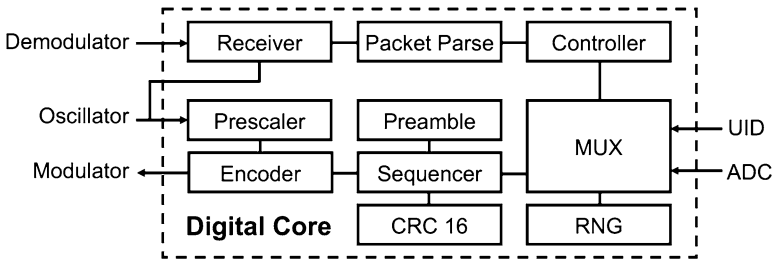


Fig. 14 EPC Gen2 digital core with sensor data interface

from the reader and enables the transmitter logic if appropriate. The *Sequencer* constructs an EPC-compliant packet including preamble and CRC, and the *Encoder* converts the data bit stream to Miller or FM0 encoding.

7 Sensor Data Protocol

Interoperability with COTS readers requires Gen2 protocol compatibility. Figure 15 illustrates how the protocol works and how sensor data is retrieved. The Query and ReqRN commands implement anticollision in the protocol. These commands require random numbers, which are generated by sampling the (unpredictable) clock phase at the downlink baseband edges and passing it through an LFSR. The tag identifier (ID) is queried via the Ack command, and the ID is generated by the UID, as described in Sect. 4.8. Finally, sensor data is returned through the Read command. An example measured ID returned by one of our tags is shown in Fig. 15.

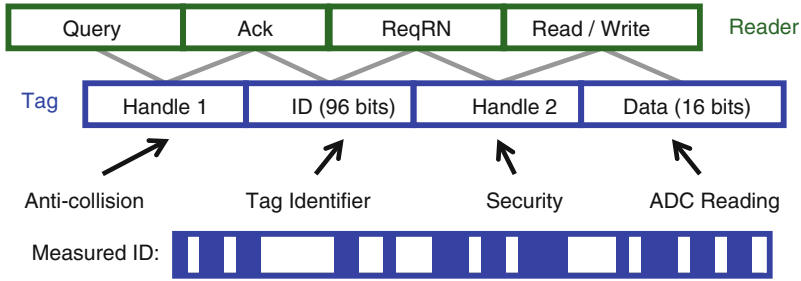


Fig. 15 EPC Gen2-compliant sensor data protocol

8 Performance

This system was fabricated in a 0.13 μm CMOS process. The die photo is shown in Fig 19. A ground shield made of the 4 μm thick top aluminum metal covers the entire die to reduce the impact of EMI, capacitive, and light sensitivity on the extremely high-impedance internal instrumentation nodes. First-silicon functionality was ensured through extensive FPGA testing of the digital core with a commercial RFID reader through a passive RFID analog front end composed of discrete COTS ICs (see [2] for a schematic). Mixed signal simulations were performed to verify the communication interfaces between the digital core and the analog blocks (UID, ADC, and demodulator).

Tables 2 and 3 summarize and compare the performance of this work to other tags presented in the literature. For a comparison of tag features, refer back to Table 1. This work compares favorably to other sensor tags without resorting to duty cycling [2, 4], loss of Gen2 compatibility [8], or elimination of the ADC [3]. Drawing comparisons to tags lacking sensors and/or addressability (all remaining tags) is less meaningful, but [1, 7] clearly achieve the best RF sensitivity for a fully passive tag. In part, our use of a linear regulator, ADC, and more complex state machine simply limits the achievable sensitivity.

9 In Vivo, In-Flight Biosignal Acquisition

The ability to monitor in vivo biosignals such as temperature and motor patterns in an organism without altering the kinematic output is invaluable to many fields of biology and is particularly important in understanding the control of movement in humans and other animals. Currently, there is only a general understanding of how temperature affects muscle performance. Though the temperature dependence of muscle contraction has already been established from in vitro preparations, the ability to correlate locomotor and kinematic performance with body temperature has thus far been hindered by obtrusive sensors [30, 31].

Table 2 Sensor tag performance

<i>System</i>	
Current consumption	9.2 μ A
Unregulated voltage	1.8 V–3.6 V
RF sensitivity	–12 dBm
Peak rectifier efficiency	37%
IC area	2.0 mm ²
<i>Components</i>	
Analog core	1.2 μ A
Reference oscillator	260 nA
Digital core	6.0 μ A
ADC and UID	500 nA
Biosignal Amp	1.2 μ A

Table 3 Comparison of published (measured) tag performance

Author	Sensitivity(dBm)	Active	Sleep(μ A)	Tech.(μ)	Clock
This work	–12	9 μ A	6	0.13	3 MHz
Sample [2]	–9 ^a	800 μ A	2	PCB	3 MHz
Kim [3]	–5.3 ^b	15 μ A	0.5	0.25 ^c	-
Kocer [4]	–12.3 ^a	660 μ A	0.9	0.25	-
Nakamoto [5]	–6.2 ^d	87.8 μ W	-	0.25	-
Pillai [6]	19 ^a	12 μ A	0.6	1.0	-
Karthaus [1]	–17.8	1.5 μ A	-	0.5 ^c	300 kHz
Barnett [7]	–14	2.75 μ A ^b	-	0.13 ^c	1.28 MHz
Shen [8]	–8.2 ^b	10.2 μ A ^c	-	0.35 ^c	-
Cho [9]	–12.9 ^b	3.4 μ A	-	0.25	330 kHz
Shenghua [10]	–20.5 ^b	0.9 μ W	-	0.18	-

^aDuty cycling required at this input power

^bCalculated based on 20% rectifier efficiency and 50% modulation losses

^cProcess includes Schottky diodes

^dNo power numbers reported, but authors claim 36.6% rectifier efficiency and 4.3 m range given 4W EIRP

^eSimulation result, no measurements reported

The hawkmoth, *Manduca sexta*, provides a well-documented model system in which to develop an understanding between biosignal output and behavior. Like other endotherms, *Manduca* utilizes the heat released from synchronous isometric muscle contractions to generate an elevated core temperature [32]. Increased muscle temperature allows these insects to increase their wing beat frequency and thus produce greater mechanical power output [33].

We used our sensing tag with a COTS RFID reader to record the in-flight temperature of the dominant flight muscles of *Manduca* (the dorsolongitudinal muscles: DLM_{1 sensu} [34]). The copper-constantan thermocouple was inserted into the DLM₁, approximately 3 mm below the dorsal aspect of the cuticle (Fig. 16). Using the instantaneous readout of our wireless sensor, we can correlate the rate of heat production during flight with the observed kinematic response. Both rapid and




Fig. 16 Schematic of the cross section of *Manduca* dorsolongitudinal muscle temperature measurement




Fig. 17 Measured temperature versus time from a tethered moth

large scale changes in body temperature will reveal how *Manduca* thermoregulates according to its energetic requirements.

Results from tethered and untethered temperature measurements show that the temperature of the DLM_1 increased with time, stabilizing around 29°C for the tethered moth and $30\text{--}35^\circ\text{C}$ for the untethered moth (Figs. 17 and 18). Because the behavior of the moth is uncontrolled, the untethered data show the moth resting at the 4 min and 7 min time marks. The tethered data represent one continued period of activity by the moth. The temperature is different between the two trials due to different implant depths of the thermocouple as well as different characteristics of the moth.

Figure 19 shows the PCB used for this experiment, which measures less than 1 cm^2 and weighs 0.25 g. Additional thermocouple gain is achieved using on-board 1.8 V micropower op-amps that are powered from our on-chip 1.8 V regulators. A photo of the moth wearing the system is shown in Fig. 20. Including the antenna and thermocouple, the entire system weighs 0.35 g. Full system specifications are listed in Table 4. The power of our moth-worn system is two orders of magnitude lower than an active radio-based tag [35]. This work enables the first long-term in-flight recording of an insect by removing both the wires and batteries from the recording equipment.




Fig. 18 Measured temperature versus time from an untethered moth




Fig. 19 Chip and board micrographs



Fig. 20 Photograph showing an untethered moth with the RFID measurement sensor

Table 4 RFID sensor specifications

Sample rate	100 Hz
Resolution	4.7 μV / 0.11 $^{\circ}\text{C}$
Peak noise	± 1.5 bits = $\pm 7 \mu\text{V}$

10 Conclusions

For many wirelessly deployed sensors, periodic battery replacement is costly at best and infeasible at worst. Two key features of RFID make it attractive for wireless sensor deployment: power is wirelessly and deliberately supplied to the tag, and backscatter modulation allows for nearly zero power tag-to-reader communication. Coupling this technology with biomedical applications promises exciting advances in medical delivery. Specifically, Gen2 compatibility, tag addressability, and the ability to interface with useful sensors will enable these biosensors to be deployed as medical and research tools.

Our multipurpose sensing tag was designed for a variety of sensor interfaces such as EMG (electromyograms), thermocouple readout, and gas detection. These applications demand an extremely low-noise floor ($< 2 \mu\text{V}_{\text{rms}}$ input referred) under a relatively low bandwidth ($< 1\text{kHz}$). Furthermore, the tag lifespan should not be limited by battery life. This chapter demonstrates the feasibility of in vivo, untethered, in-flight temperature recording of insects. The key novelty is that this long-term wireless recording was previously unattainable due to the size, weight, and lifespan of conventional sensors.

Acknowledgments This work was supported in part by NSF ECS Award 0824265, the Komen Endowed Chair, the ONR MURI grant to TLD, and Intel Labs Seattle. This chapter is ©2010 IEEE reprinted, with permission, from IEEE Journal of Solid-State Circuits (JSSC), Vol. 45, No. 10, 2010.

References

1. U. Karthaus and M. Fischer, "Fully integrated passive UHF RFID transponder IC with 16.7- μW minimum RF input power," *IEEE Journal of Solid-State Circuits*, vol. 38, no. 10, pp. 1602–1608, Oct. 2003.
2. A.P. Sample, D.J. Yeager, P.S. Powlledge, A.V. Mamishev, and J.R. Smith, "Design of an RFID-based battery-free programmable sensing platform," in *IEEE Transactions on Instrumentation and Measurement*, 2008.
3. S. Kim, J.-H. Cho, H.-S. Kim, H. Kim, H.-B. Kang, and S.-K. Hong, "An EPC Gen 2 compatible passive/semi-active UHF RFID transponder with embedded FeRAM and temperature sensor," in *Solid-State Circuits Conference, 2007. ASSCC '07. IEEE Asian*, Nov. 2007, pp. 135–138.
4. F. Kocer and M. Flynn, "A new transponder architecture with on-chip ADC for long-range telemetry applications," *IEEE Journal of Solid-State Circuits*, vol. 41, no. 5, pp. 1142–1148, May 2006.

5. H. Nakamoto, D. Yamazaki, T. Yamamoto, H. Kurata, S. Yamada, K. Mukaida, T. Ninomiya, T. Ohkawa, S. Masui, and K. Gotoh, "A passive UHF RF identification CMOS tag IC using ferroelectric RAM in 0.35- μ m technology," *IEEE Journal of Solid-State Circuits*, vol. 42, no. 1, pp. 101–110, Jan. 2007.
6. V. Pillai, H. Heinrich, D. Dieska, P. Nikitin, R. Martinez, and K. Rao, "An ultra-low-power long range battery/passive RFID tag for UHF and microwave bands with a current consumption of 700 nA at 1.5 V," *Circuits and Systems I: Regular Papers, IEEE Transactions on*, vol. 54, no. 7, pp. 1500–1512, July 2007.
7. R. Barnett, G. Balachandran, S. Lazar, B. Kramer, G. Konnail, S. Rajasekhar, and V. Drobny, "A passive UHF RFID transponder for EPC Gen 2 with -14 dBm sensitivity in 0.13 μ m CMOS," in *Solid-State Circuits Conference, 2007. ISSCC 2007. Digest of Technical Papers. IEEE International*, Feb. 2007, pp. 582–623.
8. H. Shen, L. Li, and Y. Zhou, "Fully integrated passive UHF RFID tag with temperature sensor for environment monitoring," in *ASIC, 2007. ASICON '07. 7th International Conference on*, Oct. 2007, pp. 360–363.
9. N. Cho, S.-J. Song, S. Kim, S. Kim, and H.-J. Yoo, "A 5.1- μ W UHF RFID tag chip integrated with sensors for wireless environmental monitoring," in *Solid-State Circuits Conference, 2005. ESSCIRC 2005. Proceedings of the 31st European*, Sept. 2005, pp. 279–282.
10. Z. Shenghua and W. Nanjian, "A novel ultra low power temperature sensor for UHF RFID tag chip," in *Solid-State Circuits Conference, 2007. ASSCC '07. IEEE Asian*, Nov. 2007, pp. 464–467.
11. F. Cilek, K. Seemann, D. Brenk, J. Essel, J. Heidrich, R. Weigel, and G. Holweg, "Ultra low power oscillator for UHF RFID transponder," in *Frequency Control Symposium, 2008 IEEE International*, May 2008, pp. 418–421.
12. R. Barnett and J. Liu, "A 0.8 V 1.52 MHz MSVC relaxation oscillator with inverted mirror feedback reference for UHF RFID," in *Custom Integrated Circuits Conference, 2006. CICC '06. IEEE*, Sept. 2006, pp. 769–772.
13. C. Klafp, A. Missoni, W. Pribyl, G. Holweg, and G. Hofer, "Analyses and design of low power clock generators for RFID TAGs," in *Research in Microelectronics and Electronics, 2008. PRIME 2008. Ph.D.*, 22 2008–April 25 2008, pp. 181–184.
14. F. Song, J. Yin, H. Liao, and R. Huang, "Ultra-low-power clock generation circuit for EPC standard UHF RFID transponders," *Electronics Letters*, vol. 44, no. 3, pp. 199–201, 31 2008.
15. I. Zalbide, J. Vicario, and I. Velez, "Power and energy optimization of the digital core of a Gen2 long range full passive RFID sensor tag," in *RFID, 2008 IEEE International Conference on*, April 2008, pp. 125–133.
16. A. Ricci, M. Grisanti, I. De Munari, and P. Ciampolini, "Design of a 2 μ W RFID baseband processor featuring an AES cryptography primitive," in *Electronics, Circuits and Systems, 2008. ICECS 2008. 15th IEEE International Conference on*, 31 2008–Sept. 3 2008, pp. 376–379.
17. T. Umeda, H. Yoshida, S. Sekine, Y. Fujita, T. Suzuki, and S. Otaka, "A 950 MHz rectifier circuit for sensor networks with 10 m-distance," in *Solid-State Circuits Conference, 2005. Digest of Technical Papers. ISSCC. 2005 IEEE International*, Feb. 2005, pp. 256–597 Vol. 1.
18. G. De Vita and G. Iannaccone, "Design criteria for the RF section of UHF and microwave passive RFID transponders," *Microwave Theory and Techniques, IEEE Transactions on*, vol. 53, no. 9, pp. 2978–2990, Sept. 2005.
19. R. Barnett, S. Lazar, and J. Liu, "Design of multistage rectifiers with low-cost impedance matching for passive RFID tags," in *Radio Frequency Integrated Circuits (RFIC) Symposium, 2006 IEEE*, June 2006, pp. 4 pp.–.
20. J. Yi, W.-H. Ki, and C.-Y. Tsui, "Analysis and design strategy of UHF micro-power CMOS rectifiers for micro-sensor and RFID applications," *Circuits and Systems I: Regular Papers, IEEE Transactions on*, vol. 54, no. 1, pp. 153–166, Jan. 2007.
21. S. Mandal and R. Sarpeshkar, "Low-power CMOS rectifier design for RFID applications," *Circuits and Systems I: Regular Papers, IEEE Transactions on*, vol. 54, no. 6, pp. 1177–1188, June 2007.

22. H. Nakamoto, D. Yamazaki, T. Yamamoto, H. Kurata, S. Yamada, K. Mukaida, T. Ninomiya, T. Ohkawa, S. Masui, and K. Gotoh, "A passive UHF RFID tag LSI with 36.6% efficiency CMOS-only rectifier and current-mode demodulator in 0.35 μm FeRAM technology," in *Solid-State Circuits Conference, 2006. ISSCC 2006. Digest of Technical Papers. IEEE International*, Feb. 2006, pp. 1201–1210.
23. T. Le, K. Mayaram, and T. Fiez, "Efficient far-field radio frequency energy harvesting for passively powered sensor networks," *IEEE Journal of Solid-State Circuits*, vol. 43, no. 5, pp. 1287–1302, May 2008.
24. R. Glidden, C. Bockorick, S. Cooper, C. Diorio, D. Dressler, V. Gutnik, C. Hagen, D. Hara, T. Hass, T. Humes, J. Hyde, R. Oliver, O. Onen, A. Pesavento, K. Sundstrom, and M. Thomas, "Design of ultra-low-cost UHF RFID tags for supply chain applications," *Communications Magazine, IEEE*, vol. 42, no. 8, pp. 140–151, Aug. 2004.
25. Y. Su, J. Holleman, and B. Otis, "A digital 1.6 pJ/bit chip identification circuit using process variations," *IEEE Journal of Solid-State Circuits*, vol. 43, no. 1, pp. 69–77, Jan. 2008.
26. J. Holleman, A. Mishra, C. Diorio, and B. Otis, "A micro-power neural spike detector and feature extractor in 13 μm CMOS," in *IEEE Custom Integrated Circuits Conference, 2008. CICC 2008*, 2008, pp. 333–336.
27. T. Denison, K. Consoer, W. Santa, A. Avestruz, J. Cooley, and A. Kelly, "A 2 uw 100 nv/rthz chopper-stabilized instrumentation amplifier for chronic measurement of neural field potentials," *IEEE Journal of Solid-State Circuits*, 2009.
28. K. Makinwa and J. Huijsing, "A wind sensor with an integrated low-offset instrumentation amplifier," in *Electronics, Circuits and Systems, 2001. ICECS 2001. The 8th IEEE International Conference on*, 2001.
29. R. Yazicioglu, P. Merken, R. Puers, and C. Van Hoof, "A 60 uw 60 nv/root hz readout front-end for portable biopotential acquisition systems," *IEEE Journal of Solid-State Circuits* 2007.
30. A. Bennett, "Thermal dependence of muscle function," *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, vol. 247, no. 2, p. 217, 1984.
31. R. Josephson, "Contraction dynamics of flight and stridulatory muscles of tettigoniid insects," *Journal of Experimental Biology*, vol. 108, no. 1, p. 77, 1984.
32. B. Heinrich, "Thermoregulation in endothermic insects," *Science*, vol. 185, no. 4153, p. 747, 1974.
33. R. Stevenson and R. Josephson, "Effects of operating frequency and temperature on mechanical power output from moth flight muscle," *Journal of Experimental Biology*, vol. 149, no. 1, p. 61, 1990.
34. Y. Kondoh and Y. Obara, "Anatomy of motoneurons innervating mesothoracic indirect flight muscles in the silkworm, *Bombyx mori*," *Journal of Experimental Biology*, vol. 98, no. 1, p. 23, 1982.
35. D. Daly, P. Mercier, M. Bhardwaj, A. Stone, Z. Aldworth, T. Daniel, J. Voldman, J. Hildebrand, and A. Chandrakasan, "A pulsed UWB receiver SoC for insect motion control," *IEEE Journal of Solid-State Circuits*, vol. 45, no. 1, p. 153, 2010.

Battery-less Wireless Sensors Based on Low Power UHF RFID Tags

Roc Berenguer, Iván Rebollo, Ibon Zalbide, and Iñaki Fernández

1 Introduction

Nowadays, there is an undeniable and unstoppable trend to use RF identification (RFID) in a number of applications, such as supply chain management, public transportation, and access control [1]. The use of this technology entails a number of advantages over barcode technologies such as tracking people, items, and equipment in real time, non-line of sight requirement, long reading range, and standing harsh environments.

Recently, the combination of RFID with sensory systems has extended the applications of RFID to environmental monitoring [2, 3] or to health-care applications [4]. Those existing sensors, such as [2, 3], usually operate at the 13.56 MHz and 134.2 kHz frequency bands, respectively. However, these sensors have the inconvenience of a limited reading range (a few centimeters) and high cost. To overcome this limitation, RFID sensor development has been focused on RFID tags using the UHF bands (868 MHz, 900 MHz, and higher) as shown by recently reported designs [5–7] which offer temperature monitoring and higher reading

R. Berenguer
Centro de Estudios e Investigaciones Técnicas de Gipuzkoa (CEIT)
and Tecnun - University of Navarra, Po Manuel de Lardizábal 15,
20018 Donostia - San Sebastián, Spain
e-mail: rberenguer@ceit.es

I. Rebollo (✉) • I. Zalbide
FARSENS S.L., Paseo Mikeletegi, 54,
20009 Donostia-San Sebastián, Spain
e-mail: ivan.rebollo@gmail.com; ibon.zalbide@farsens.com

I. Fernández
Tecnun University of Navarra, Po Manuel de Lardizábal 13,
20018 Donostia-San Sebastián, Spain
e-mail: ifernandez@tecnun.es




Fig. 1 Architecture of an RFID tag with external sensor

distances than previous ones. On the other hand, the increasing popularity of the standard EPC C1G2 [8] (ISO 18000-6C [9]) makes it possible the use of a universal communication standard compatible with many readers from many vendors. Therefore, there is a strong motivation to supply optimized RFID tags able to support different sensor types, such as temperature sensors and pressure sensors. and able to use standardized communication protocols such as EPC C1G2.

Figure 1 shows the typical architecture of an RFID tag. The antenna receives the signal emitted by the reader. In order to achieve the maximum power transference from the antenna to the voltage multiplier, a matching network is needed. Typically this matching network is implemented together with the antenna. The voltage multiplier rectifies the incoming signal charging the supply capacitor C_{SUPPLY} . This capacitor is used to supply power to the rest of the tag. The analog front end provides the signals that the rest of the tag requires to work properly, such as regulated voltages, clock, and reset signals. It also is in charge of demodulating the incoming amplitude shift keying (ASK) signal and modulating the tag answer. The digital core communicates with the EEPROM and, when present, the implemented sensor. It also realizes the required actions to answer the reader queries using the EPC C1G2 standard.


The present contribution focuses on reducing the power consumption of both the analog front end and the digital core. As a result a low power front end is implemented and a power-effective architecture is proposed for the digital core of an EPC C1G2 sensor tag. Finally a complete sensory system is also presented showing successful communication up to 2.4 m between the tag sensor and the reader.

2 System Constraints

In the present section, the main limitations involving the maximum communication distance between tag and reader are presented. Three limitations are given for the reader-to-tag link and another one for the tag-to-reader link.

Figure 2 shows a simplified equivalent circuit of the complete RFID sensor, for analysis purposes, where the tag+sensor is represented by its input impedance (R_{IC} and X_{IC}).

Fig. 2 Equivalent circuit for the whole RFID tag



2.1 Previous Definitions

1. Power available (P_{AV}) at the input of the tag antenna

$$P_{AV} = \frac{P_{EIRP}}{4\pi r^2} \frac{\lambda^2}{4\pi} G \quad (1)$$

where G is the tag antenna gain, r is the distance between the reader and the tag, λ is the wavelength, and P_{EIRP} is the effective isotropic radiated power.

2. Input power ($P_{RF,IN}$) [10]

$$P_{RF,IN} = P_{AV}(1 - \rho^2) = P_{AV} \frac{4X^2}{R_A^2 + 4X^2} \quad (2)$$

where ρ is the reflection coefficient, R_A is the impedance of the antenna, and X is the reactance introduced by the load modulator as Fig. 2 shows.

3. Backscattered power (P_{BS}) [10]

$$P_{BS} = P_{AV} \frac{4(R_A^2 + X^2)}{R_A^2 + 4X^2} \quad (3)$$

4. Quality factor (Q)

$$Q = \frac{X_{IC}}{R_{IC}} \quad (4)$$

where R_{IC} and X_{IC} are the components of the complex input impedance of the tag + sensor (Fig. 2).

2.2 Forward Link Constraints (Reader \rightarrow Tag)

The constraints regarding the reader-to-tag link are related to the minimum input power ($P_{RF,IN}$) to ensure a correct operation and the minimum required voltage (V_{MIN}) at the input of the rectifier. Both constraints assume a constant continuous wave signal at the input of the tag.

2.2.1 Minimum Input Power

The reader feeds the tag with an RF continuous wave. The incident power must be big enough to afford the whole power consumption of the analog and digital blocks of the tag and the external sensor. Equation 5 provides a relation between power consumption and maximum range [10]:

$$P_{AV} \frac{4X^2}{R_A^2 + 4X^2} \geq \frac{1}{\eta} (P_{ANALOG} + P_{DIGITAL} + P_{SENSOR}) \quad (5)$$

where P_{ANALOG} , $P_{DIGITAL}$, and P_{SENSOR} are the power consumption of the analog front end, digital core, and the external sensor, respectively, and η is the efficiency of the rectifier circuit.

From (5) and considering free-space power losses, the communication range is given by

$$r_P^2 \geq P_{EIRP} \frac{\lambda^2 G}{4\pi} \frac{X^2}{R_A^2 + 4X^2} \frac{\eta}{P_{ANALOG} + P_{DIGITAL} + P_{SENSOR}}. \quad (6)$$

From (6), for a given P_{AV} , an increment on the efficiency (η) implies an increment on the communication range. This fact should be considered when designing the voltage multiplier, which is the most critical block in the efficiency.

Besides, as shown in (6), if $P_{ANALOG} + P_{DIGITAL} + P_{SENSOR}$ increases, r_P de-creases. Therefore the minimization of the power consumption of the analog and digital blocks is mandatory in order to achieve longer communication distances for a given power consumption of an external sensor.

2.2.2 Minimum Voltage at the Input of the Voltage Multiplier

Not only a minimum input power is necessary at the tag antenna. It is also necessary a minimum voltage amplitude (V_{MIN}) at the input of the voltage multiplier to guarantee a proper operation. This way the voltage multiplier operates with good efficiency and is able to achieve, at its output, the required supply voltage both for analog and digital circuits.


The circuit shown in Fig. 3 is obtained from the one shown in Fig. 2, where

$$R_{IC-P} = R_{IC}(Q^2 + 1). \quad (7)$$

The quality factor of the matching network (Fig. 3) is defined as

$$Q_{MN} = \sqrt{\frac{R_{IC-P}}{R_A} - 1}. \quad (8)$$

Fig. 3 Serial to parallel transformation of the input impedance of the tag



From Fig. 3, (9) can be obtained:

$$\frac{V_2}{V_1} = \frac{1}{1 - \frac{j}{R_{IC-P}C\omega} - \frac{j}{LC\omega^2} + \frac{C_{IC-P}}{C}} \quad (9)$$

where V_1 and V_2 are, respectively, the voltage at the input and the output of the matching network.

In Fig. 3, R_{IC-P} and C_{IC-P} are the equivalent parallel values of R_{IC} and X_{IC} of Fig. 2. L and C are the values of the components of the matching network.

As mentioned earlier and shown in Fig. 3, the antenna and the tag can be simplified as a voltage source and a real impedance (antenna) followed by a complex impedance (rest of the tag and sensor). A matching network has been implemented in order to provide maximum power transference which implies an increase of the circuit efficiency. Additionally, the voltage before and after the matching network is increased. The voltage is multiplied by the quality factor of the matching network as shown in (10):

$$\frac{V_2}{V_1} \cong 0.7Q_{MN}. \quad (10)$$

Equation (10) is valid whenever R_{IC-P} is much bigger than R_A . Anyway, it is quite a restrictive simplification. Besides, it is multiplied by 0.7 in order to make it more restrictive. If a matching condition is achieved ($R = R_A$), then


$$V_1 = \sqrt{P_{AV}R_A}. \quad (11)$$

From (1) and (11), the range can be calculated as

$$r_V \simeq \frac{0.7Q_{MN}\lambda\sqrt{P_{EIRP}R_A G}}{4\pi V_{MIN}}. \quad (12)$$

Equation (12) provides the relation between the necessary voltage at the input of the voltage multiplier ($V_2 > V_{MIN}$) and the required distance between tag and reader. An increment of V_{MIN} implies a reduction of the communication range. Therefore, a minimum V_{MIN} is desired. But this parameter depends on the threshold voltage of the employed devices and therefore it is a process dependent parameter.

Fig. 4 Modulated RF signal during downlink communication



This constraint is not usually included in the system analysis of the RFID bibliography, although it could be the limitation of the maximum achievable reading range.

2.2.3 Modulation Limitations

When the communication starts the input signal from the reader is modulated and the assumption of constant continuous wave (CW) is no more realistic. The modulation creates time intervals in which no input power is present. During these time intervals, the supply capacitor gives the required energy to the rest of the tag. Thus, it is interesting to study the behavior of the supply capacitor and see the limitations from the energetic point of view.

The EPC C1G2 communication protocol specifies that the forward link (Reader \rightarrow Tag) communication shall be ASK with a modulation depth of 90%, and the backward link (Tag \rightarrow Reader) communication uses ASK or phase-shift keying (PSK) backscattering.

During forward link communication the RF envelope is modulated with pulses of duration PW (Fig. 4). High PW is required in order to obtain a stable communication. However, low PW is required to minimize the time periods with no input power. The designer has to trade off the value of PW.

Assuming that the power received during PW is negligible, the supply capacitor will supply the energy required by the rest of the tag. This will produce an energy discharge during PW.

A similar effect occurs when the tag replies to the reader backscattering the received signal. The modulation in the backscattered signal is produced deadapting the antenna at the working frequency. When that happens, the tag can communicate with the reader but cannot receive all the energy of the input signal.

The energetic behavior of the supply capacitor can be studied simplifying the charging operation of the tag to an RC circuit, R_{EQ} and C_{EQ} . The current generated by the incoming power of the reader charges C_{EQ} . The speed of the charge process depends on the incoming power, the power consumption of the circuitry connected to C_{EQ} , and the value of C_{EQ} . C_{EQ} is the supply capacitor plus the parasitic

capacitances connected to the supply node and R_{EQ} represents the resistance of the system to the charge process of C_{EQ} . Two states have to be analyzed:

1. C_{EQ} receives energy

When a voltage is applied to C_{EQ} , this capacitor accumulates energy. The maximum accumulable energy ϵ_{MAX} is

$$\epsilon = \int_0^Q \frac{q}{C} dq = \frac{1}{2} \frac{Q^2}{C_{EQ}} = \frac{1}{2} C_{EQ} V^2. \quad (13)$$

The energy accumulation in C_{EQ} is given by

$$E_C = \frac{1}{2} \frac{q^2}{C_{EQ}} = \epsilon_{MAX} \left(1 - e^{-\frac{t}{R_{EQ}C_{EQ}}} \right)^2. \quad (14)$$

Assuming that the tag needs to have $E_{C_{EQ}} = \gamma \epsilon_{MAX}$ to start working, the settling time of the tag can be calculated from (14):

$$t_{CHARGE} = -R_{EQ}C_{EQ} \ln(1 - \sqrt{\gamma}). \quad (15)$$

The EPC C1G2 standard defines the maximum settling time to be 1,500 μ s. Thus, the selected C_{EQ} should maintain t_{CHARGE} below this limit. As t_{CHARGE} depends on $R_{EQ}C_{EQ}$, and each design has its own equivalent R_{EQ} , the value of C_{EQ} shall be selected for each design.

2. C_{EQ} receives no energy

When the tag is not receiving any energy from the input signal, C_{EQ} supplies the required energy, discharging the capacitor. As the accumulated energy decreases, the voltage generated between the terminals of C_{EQ} is reduced as well.

Figure 5 shows the voltage in C_{EQ} when the input power is modulated. We can observe that for great values of C_{EQ} , the voltage reduction caused by the modulation is minimal. On the other hand, if C_{EQ} is small, the modulation causes a fast decrease in the supply voltage. In order to maintain the supply voltage upon a limit, a minimum value for C_{EQ} is required, C_{MIN} .

To sum up, the communications range is also restricted by the settling time $t_{CHARGE} < 1,500\mu$ s and the value of the capacitance $C_{EQ} > C_{MIN}$. The values of t_{CHARGE} and C_{MIN} are dependent of the specific implementation of the tag.

If $C_{EQ} > C_{MIN}$ but $t_{CHARGE} > 1,500\mu$ s, the physical communication range of the tag would yet be given by (6) and (12), but the settling time given in the standard would not be satisfied. The system will work correctly if the reader accepts extended settling time. This problem can also be solved by software sending some dummy commands at the beginning of the communication, increasing the actual acceptable settling time.





Fig. 5 Behavior of C_{EQ} when input signal is modulated

Fig. 6 BPSK constellation



2.3 Backward Link Constraints (Tag \rightarrow Reader)

The main constraint concerning the tag-to-reader link is the maximum probability of error allowed at the input of the reader.

The backward link is performed using a binary phase-shift keying (BPSK) modulation. For a BPSK receiver architecture, the probability of error can be calculated from the analysis of its constellation (Fig. 6) as

$$P_{\text{error}} = Q\left(\frac{d}{2\sigma}\right) = Q\left(\frac{2A \sin(\theta)}{2\sigma}\right) = \frac{1}{2} \operatorname{erfc}\left(\frac{A \sin(\theta)}{\sqrt{2}\sigma}\right) \quad (16)$$

where σ is the standard deviation of the noise, A is the amplitude of the backscattered signal, and θ is the phase of the signal.

Equation (16) does not accurately represent the real value of the probability of error because it does not take into account neither the phase noise nor the thermal noise. Equation 17 gives a more accurate value for the probability of error [10]:

$$P_{\text{error}} = \frac{1}{2} \left\{ \operatorname{erf}\left(\frac{A \sin(\theta)(2 \cos(\varphi) - 1)}{2\sigma}\right) \operatorname{erf}\left(\frac{A \sin(\theta)}{2\sigma}\right) \right\} \leq 10^{-3} \quad (17)$$

where φ is the phase-noise-related term.

The value of 10^{-3} is considered enough to achieve an appropriate communication.

Next section will present the design of the different blocks of the analog front end of the tag, taking into account the already-mentioned constrains.

3 Analog Front End

Figure 1 shows the typical RFID sensor node architecture. It is formed by three main blocks: the analog front end, the digital module, and the sensor. The last one can be either integrated on chip, typical of temperature sensors, or off chip, typical of pressure sensors, humidity sensors, and accelerometers that communicate with the digital module. The digital module controls the communication flow with the reader and when necessary with the sensor. The EEPROM stores the sensor ID and when necessary calibration values for the sensor.

As mentioned before, the present section is focused in the analog front-end design. Figure 7 shows the analog front-end block diagram. The RF incoming signal is rectified and multiplied in the voltage multiplier (VM) block. This block charges the supply capacitor (C_{SUPPLY}) and stores the required energy to operate. It also provides the required V_{DD} voltage value for proper operation of the analog front end. A voltage limiter is also implemented in order to avoid possible damages in the circuit, due to voltage surges whenever reader and tag are very close.

Two series voltage regulators, 1.4 V and 2.1 V, are also implemented. The first one (1.4 V) is intended to supply a regulated voltage to the digital modules, while the second one (2.1 V) is intended to supply a regulated voltage to the clock generator (480 kHz), the charge pump, the EEPROM, and the sensor. By using the voltage regulator, the PSRR requirements are relaxed for the clock generator and the sensor. In order to demodulate the ASK incoming signal, a demodulator is needed and so, it is connected directly to the antenna. A load modulator is also implemented to backscatter the signal; thus communication between tag and reader is possible.

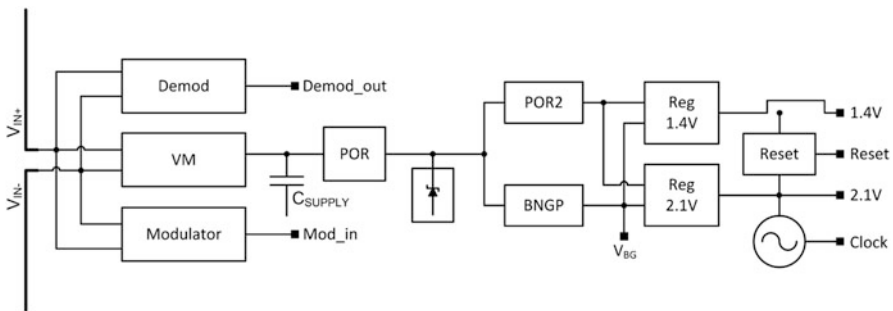


Fig. 7 Simplified block diagram of the analog front end




Fig. 8 N-stage Greinacher topology with the supply capacitor connected

Two power-on reset (POR) blocks and some control signals are added to the analog front end in order to manage power consumption during start-up of the RFID sensor node. During start-up, all blocks remain switched off until a supply voltage level of 1.6 V is detected by POR_1 block. At that moment, only the bandgap reference block is started. The POR_2 block activates the rest of the blocks once the bandgap reference voltage is stable, avoiding erroneous start-up of the different blocks and optimizing the power consumption of the analog front end during start-up. Once the voltage regulators are stable and clock signal is generated, the front end sends a reset signal to the digital module. Finally, using the ASK_EN signal, the ASK demodulator can be switched off during transmission mode in order to save power. The key blocks of the front end are the rectifier (voltage multiplier), the voltage limiter, the ASK demodulator, and the load modulator. Those blocks have been optimized for low-power operation. The descriptions of these blocks are detailed in the following sections.

3.1 Matching Network and Voltage Multiplier

The goal of the VM circuit is to charge the supply capacitor and provide the supply voltage to the other analog front-end blocks. This capacitor will supply the necessary energy to the rest of the tag. As it has been mentioned previously, high efficiency and Q are desired in order to achieve higher distances.

Greinacher topology (Fig. 8) has been selected because high multiplication factors can be reached with high substrate capacitance, typical of integrated circuits. Other topologies are too complex for RFID systems and its multiplication factors are smaller [11].

Using the Greinacher topology, the output voltage can be obtained approximately with the following formula:

$$V_{OUT} \approx 2K_N N (V_{IN} - V_{FWD}) \quad (18)$$

where V_{OUT} is the output voltage, V_{FWD} is the voltage drop on the diodes, V_{IN} is the voltage at the input of the multiplier, N is the number of stages and K_N is a constant dependant on the number of stages, ($K_N \approx 2$).

A low V_{FWD} value is desired in order to obtain a high value of V_{OUT} and maximize the communication range (12). This can be implemented using Schottky diodes in the voltage multiplier design [12].

For environmental monitoring applications, usually, long reading distances are required, so the whole tag has to be optimized in order to achieve the maximum communication range as explained earlier. Although both the voltage multiplier and the ASK demodulator are connected directly to the matching network output, the voltage multiplier impedance is much smaller than the ASK demodulator; therefore the input impedance of voltage multiplier is critical.

The limitations that depend on the voltage multiplier are given by (6) and (12). The optimization of the voltage multiplier depends on the most critical of both constraints (r_P and r_V). In the voltage multiplier design, different parameters will be optimized depending on which constraint is the real limitation. In some cases, higher r_P will be required and in others r_V is the real limitation that has to be increased. So a classification based on the different constraints can be made.

3.1.1 Power Constraint Is Critical

In the case of $r_P < r_V$, the reading range is limited by the power consumption of the sensor RFID tag. Two alternatives are given in order to obtain higher r_P .

The first one is obtaining lower power consumption of the analog and digital blocks. This is a task to be performed by digital and analog designers.

The second choice deals with the design of the voltage multiplier itself. It consists on increasing conversion efficiency.

In order to obtain higher efficiency, a minimum number of stages is desired. The main reason of this increase is that the Q of the input impedance is also increased and so the Q of the matching network (Q_{MN}). Therefore, a matching network with higher Q_{MN} entails higher V_2 . With higher input voltage at the input of the voltage multiplier, the diode performs better and so its efficiency is much higher.

Furthermore, to obtain the maximum efficiency, it is important to use the optimum diode area. Higher area means lower forward voltage (and higher V_{OUT}). But on the other hand higher losses to substrate are obtained. Thus, a compromise to obtain the optimum diode area should be achieved.

3.1.2 Voltage Constraint Is Critical

This constraint, $r_P > r_V$, applies to the input voltage of the voltage multiplier. From (10) and (11), the output voltage of the matching network can be defined as

$$V_2 \cong 0.7Q_{MN}\sqrt{P_{AV}R_A}. \quad (19)$$

So, using (8) in (19),

$$V_2 \cong 0.7 \sqrt{P_{AV}(R_{IC-P} - R_A)}. \quad (20)$$

As it has been mentioned before, higher input voltage at the voltage multiplier (V_2) is desired. To achieve this goal, (20) shows that the antenna impedance has to be minimized and R_P maximized (7). Therefore, both R_{IC} and Q have to be maximized.

3.1.3 Design Steps

As mentioned previously, the R_{IC} and Q have to be maximized to improve reading range. Every time a new stage (N) is added, as it is connected in parallel to the previous one the input impedance decreases. In the same way, as the value of capacitors C_S are increased, as they are connected in parallel, the input impedance is also reduced. On the other hand, in order to maximize the quality factor Q , a high C_S value and a low C_P value are required, in a minimum number of stages situation. Consequently, a trade-off between high and low C_S value should be considered in order to maximize Q and R_{IC} at the same time.

Besides, the input impedance of the matching network varies with the input power: for low input power levels, diodes perform worse (the efficiency decreases) and the input impedance changes. As the matching network assures the maximum power transference at a fixed load, if the load suffers a variation on its impedance, a mismatch between both blocks is generated and so power losses are obtained. Therefore, lower efficiency is obtained.

In order to obtain the optimum power transference, the following steps have to be followed:


- Determine the minimum input voltage (V_{MIN}) necessary for the correct performance of the tag without the matching network.
- Calculate the input impedance.
- Calculate the matching network to obtain the maximum power transference between the antenna and the tag.

This methodology guarantees the correct performance of the tag for the maximum distance. Shorter distances entail higher input power at the tag, which produces a variation of the input impedance. Therefore, a mismatching between the antenna and the tag is obtained. But power losses are widely compensated with the increase of the input power.

Both constraints demand to use a minimum number of stages, but they have to be set taking into account the input power (that the matching network transforms in input voltage) and the output voltage that is necessary to be achieved. Besides, there is a minimum value for C_P and C_S in order to guarantee the power transference between stages.

Therefore, depending on which one is the constraint, r_V or r_P , C_S and C_P have to be determined. Also, the number of stages has to be established in order to achieve

Fig. 9 Voltage limiter architecture



the required voltage, taking into account that a minimum number of stages are desired.

The designed rectifier, in the present analog front end, consists on the Greinacher topology. The optimum voltage multiplier performance is achieved with four stages, $C_S = 1$ pF and $C_P = 800$ fF. With these parameter values, the rectifier provides good efficiency (approx. 35% for an input power around -3 dBm) and a combination of low minimum input required voltage (V_{MIN}) and high quality factor ($Q_{MN} > 6$), improving power and voltage constrains, respectively [13, 14].

3.2 Voltage Limiter

Figure 9 shows the architecture of the voltage limiter (VL). A differential amplifier (M_1 - M_4) compares the bandgap reference with a divided voltage (V_X) of the supply voltage ($V_{SUPPLY} \approx V_{POR}$). The VL can only be activated when POR signal is also activated. When V_X is higher than V_{BG} , M_6 is enabled and so M_8 is switched on, which carries most of the current, preventing the supply voltage from increasing. In the designed analog front end, a limiting voltage of 3 V has been fixed in order to avoid possible damages in the circuits. Since the bandgap voltage is very stable against temperature and process variations, either the limiting voltage.

3.3 ASK Demodulator

The circuit schematic of the ASK demodulator is shown in Fig. 10. It consists on an envelope detector (a rectifier in Greinchaer's topology), a low-pass filter, implemented with diodes D_5 - D_6 and capacitors C_F , and a differential amplifier.


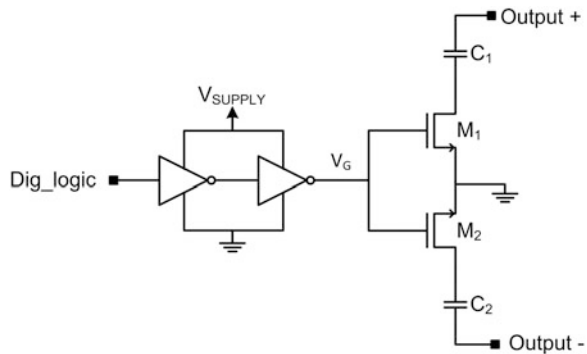


Fig. 10 ASK demodulator architecture

Fig. 11 Load modulator architecture



When the incoming signal is modulated, V_{REF+} remains constant (because diodes D_5 - D_6 prevent C_F to discharge), while V_{2+} discharges (helped by the current source), so logic output level of the ASK demodulator turns to “0” (stand by logic level = “1”).

It is preferable that this block presents input impedance much higher than the VM, so its influence in the VM input impedance is negligible; hence the available power for the rest of the modules is maximized. This can be managed mainly reducing the number of stages of the rectifier.

3.4 Load Modulator

Figure 11 shows the load modulator that has been implemented. The load modulator changes the input reactance of the tag, so PSK modulation is generated. This can be made with a single switch and a capacitor. The output is connected to the differential antenna.

In order to make the change softer and reduce spurious in the neighbors bands, transistors M_1 and M_2 are big enough; thus, the transition’s slope in V_G is not so

steep. Taking into account the input impedance of the tag and C_1, C_2 values, the input reactance is changed (maximizing the phase difference) without varying the real part, so only PSK modulation is generated but not ASK.

4 Digital Core

In a sensor RFID tag, the digital core gives the logical intelligence to communicate with the reader making use of the analog front end. For that, the communication protocol defined in the standard has to be implemented on the design. At the same time, the digital core must be able to control the EEPROM and the sensor.

The digital core architecture implemented in this work is shown in Fig. 12. It follows the basic architecture presented in the literature [15–18], but some additional blocks have been included. The incoming symbols are detected in the symbol detector. The command decoder obtains the operation code and the arguments of the instruction. In order to avoid a big input buffer, the variable length arguments, such as the selection mask, are processed in the command decoder. The cryptographic functions are also performed in the command decoder making all arguments transparent to the rest of modules. The control module controls the system with a finite-state machine and a register bank. The collision arbitration, session management, and memory lock are contained in this module. It does the necessary operations accessing the memory and the register bank. Finally, TX, the transmitter, encodes the answer with the required format. For this purpose, a backward link frequency synthesizer is included in the transmitter. The accesses to the tag memory (EEPROM or sensor) are handled by the memory access module.

Additionally to these basic modules, one more module has been included in the design: the PM module. The PM module is the power management unit which controls the activity of the rest of the modules reducing the dynamic power consumption.




Fig. 12 Architecture of the digital core

Table 1 Working states of the digital core

	STRTP	STDBY	RX	CNTRL	TX
PM	On	On	On	On	On
Symbol detector	Off	On	On	Off	Off
Command decoder	Off	Off	On	Off	Off
Control	On	Off	Off	On	Off
Memory access	On	Off	On	On	On
TX	Off	Off	Off	Off	On

4.1 Power Management

As the digital core has been designed as a synchronous design, the circuit consumes power every clock cycle. However, not all the operations that the core performs are always useful. Stopping the activity of the modules when it is useless reduces the total power consumption. A module called PM has been integrated in the architecture. This is the power management module, which controls the activity of every other module.

For power management, five different working states have been defined. Those are shown in Table 1. Every working state is optimized to perform a specific operation during the communication. The functionality of the tag is divided in five operations: initialize, wait, receive instruction, execute operation, and answer. Each working state is optimized to perform one of those operations. In each case, only the required modules are enabled and the rest of the modules are disabled. The combination of all the working states enables the system to work properly minimizing the activity of its circuitry.

The STRTP state is the initial working state. The tag checks the kill bit, and if it is not killed, the system turns to STDBY state, waiting for a signal to arrive. In this state only the symbol detector is active. When an input signal is detected, the command decoder is activated and the working state turns to RX. After receiving the whole message, the working states change to CTRL deactivating the command decoder and the symbol detector and activating the control. Finally, in the TX state the response is sent to the reader and the working state returns to STDBY.

Table 1 shows the relationship of the working states with the modules. Each working state activates the necessary modules to fulfill its functionality and deactivates all the other modules.

In order to reduce the power consumption, the unnecessary activity is eliminated using the clock-gating technique. A clock-gating cell is inserted in the clock net of every module, so that the activity of the clock signal is controlled by the enabling signals generated by PM. Only the clock signal in the modules of the current working state is activated. As a result, the switching activity of the resultant clock in each module is lower than the original one. The module consumes the expected total power when enabled but only the leakage power when disabled. As PM generates the enabling signals of the different modules in each working state, this module has always to be active.




Fig. 13 Power consumption profile of the digital core

Given that the active modules during each working state differ from one state to another, the power consumption of the digital core cannot be considered to be constant. Taking into account the time spent by the digital core in each state, a power consumption profile can be considered. Figure 13 shows a typical power profile of an RFID C1G2 tag digital core with power management. In this case the power profile during the execution of the *Req_RN* command is shown. This command performs the typical operations making use of all the working states. Thus, its power distribution can be used as the generic power consumption of the digital core. This profile has to be taken into account when analyzing the power and energy limitations of the tag for the design optimization.

4.2 Sensor Integration

The EPC C1G2 is specifically designed for identification applications and does not include support for sensors. Nevertheless there are some possible workarounds to integrate sensors in the C1G2 standards.

The most straightforward method is to use user commands for this purpose. A new user command can be defined to access the sensor. The new command may contain configuration parameters for the sensor, so that the behavior of the sensor can be controlled externally. The answer from the tag to this user command may contain the value of the sensor. The standard allows the use of user commands which makes such an RFID sensor C1G2 compliant. Nevertheless, compliance with the standard may not be good enough. In order to integrate these RFID sensors in any commercial C1G2 network, compatibility between all the elements in the

network must be ensured. Including user commands limits the compatibility, as every item in the network needs to support the new command. If one or more items in the communication chain do not support this user command, the end user will not be able to retrieve the information of the sensor. Therefore, in order to solve this incompatibility problem, mandatory commands of the standard shall be used.

A first approach is to use the EPC substitution solution [19]. In this case, the EPC is completely or partially replaced by the value of the sensor. The reader may just perform the typical identification algorithm and a higher-level layer could extract the sensorial information from the identification code. As product identification is the main goal of the C1G2 standard, all commands needed to retrieve the EPC are mandatory for every item in the EPC network. Substituting part of the product code with the value of the sensor guarantees the access to the sensor value through any C1G2-compliant reader. However, even this solution may cause some kind of incompatibility. The solution on its own works properly at identification layer, but it may affect to more advanced features of the EPC network. The EPC, from its definition, is a unique product identifier. The structure of the EPC is predefined. The data disposition inside the EPC is not casual, and upper layers of the EPC network work with it. Integrating the sensor value inside the EPC makes the unique identifier of the tag change with the measured magnitude. As the EPC is expected to be constant, this solution may cause unwanted behavior in higher levels of the EPC network.

In order to solve this compatibility issue and obtain a fully EPC C1G2-compatible tag, the memory-mapping method has been used. This method consists on redirecting the user memory bank (defined in the standard) to the sensor value [20, 21]. This way, the reader only needs to request the value contained in a specific memory location. If the requested address contains the direction assigned to the sensor, the answer of the tag will contain the value of the measurement. The main advantage of this solution is that any C1G2 network does support it, and as the EPC code is not altered, none of the advanced features of the EPC networks are damaged.

5 Implementation

The present section describes the prototype of a wireless passive sensor combining an ultra-low-power digital commercial sensor, a module-based digital core (FPGA and digital core replica power consumption module), and the integrated analog front end.

The analog front end has been designed with three output voltages: one unregulated, but limited around 3.0 V, and two regulated at 1.4 V and 2.1 V. A low-drop out (LDO) voltage regulator is connected to the 3.0 V output voltage. This way, many commercial sensors are suitable for the desired prototype.

The digital part has been implemented in an FPGA including the power management, the EPC C1G2 standard, and the needed digital interface in order to communicate with the sensor. The use of an FPGA makes the prototype to be versatile for any kind of sensor.




Fig. 14 RFID analog front-end layout

5.1 Analog Front End


A low-cost 2P4M 0.35 μm CMOS standard process has been used. Schottky diodes, high resistive poly, and EEPROM are also available in the process. Figure 14 shows the FE layout. The FE dimensions are 1777 × 1262 μm². Several PADS have been added in order to test the different blocks (voltage regulators, clock, POR) individually.

The correct performance of the presented analog front end has been tested on wafer offering a current consumption of 7.4 μA.

5.2 Digital Core

The digital core design has been described in VHDL. This way, the design can be easily targeted either to ASIC or to FPGA implementation. A first version of the digital core has been implemented on ASIC for a 0.35 μm process (Fig. 15). This design consists on an EPC C1G2-compatible core and a communication interface towards the sensor. In the implemented version, a proprietary communication

Fig. 15 EPC C1G2 digital core layout



protocol has been used to communicate with a digital sensor. In order to allow standard serial communication with commercial sensors, additional hardware is required. Unfortunately this implementation cannot be reused with a new sensor as new physical hardware and pads are required for this purpose.

As a solution to integrate a sensor in the tag allowing the communication with the digital core, the new code has been implemented in an FPGA. The most common digital interface protocols have been added to the VHDL code of the digital core. This way, there is no need to fabricate the digital core again for a fast and low-cost prototype.

5.3 Digital Commercial Sensors

Sensing products and technologies are rewriting future designs in an extensive range of industrial, medical, consumer, communications, and automotive applications.

The current sensor market has a wide portfolio of low and ultra-low-power devices. A complete range of commercial sensors offers accuracy and low cost with digital interfaces.

A deep analysis of low power sensors has been completed; sensors that are able to work at low voltage operation (from 1.4 V up to 5.5 V) and low quiescent current (from 10 μ A up to 300 μ A) can be listed. This list includes sensors of temperature, pressure, acceleration, and humidity from the most important manufacturers such as Texas Instruments, Analog Devices, ST Microelectronics, Sensirion, VTI Technologies, Kionix, and Bosch-Sensortec.




Fig. 16 Front-end characterization setup

The operation of the sensors is quite different depending on the manufacturer, magnitude to be measured, and the final application of the sensor. So, different modes to acquire and forward the data are found. As the sampling rate is directly correlated to the average power consumption, using slow data acquisition and forward rates allows achieving low power consumption.

Most of the sensors have high current peaks (some hundreds of μA) during the acquisition and a few μA during the standby time. For these conditions the point of view has to be changed from a constant quiescent current to an average current, regarding the maximum current peak during measurement (or conversion time). The peak current shakes directly the voltage supply and external capacitors must be added to reduce the drop voltage. In that case the list of sensors can be increased if the acquisition rate can be controlled.

As a conclusion, it is recommended, and mandatory in some cases, to use of an external capacitor to store enough energy to supply the needed power depending on the sensor and the application. The value of C_{SUPPLY} has to be high enough to support the discharge produced by the sensor measurement without falling under the minimum operation voltage.

At this point, it must be considered that the energy has to be stored previously through the RF signal. In some cases—depending on the transmitted power, the distance, and the capacitor value—the duration of the energy storage can be very high and dummy commands (EPC C1G2) shall be used [20].

5.4 Front-End Setup with External Regulator

In order to characterize the front-end capabilities to supply a sensor the measurement setup of Fig. 16 has been used. An ultra-low-quiescent-current LDO voltage regulator has been chosen. The selected LDO has been the TPS78225 from Texas Instruments with a quiescent current $I_Q = 1 \mu\text{A}$.

The measurement has been done under the following conditions: the reader is emulated by a continuous wave power amplifier and a flat panel antenna whose whole radiated output power is 2 W EIRP established by European regulations [22].

For the front-end measurement, a dipole antenna (Fig. 17) for test purpose has been designed. The antenna is not optimized considering size, polarization, and




Fig. 17 Front end with PCB dipole antenna




Fig. 18 Block diagram of the implemented prototype

gain; the main consideration has been the impedance matching with front-end input impedance, which is designed for a low input power and a high communication distance.

The evaluation has been done on the V_{SUPPLY} , which is limited internally at 3 V by the voltage limiter. The LDO is connected between V_{SUPPLY} and a variable load (R_{LOAD}) which is adjusted to measure the maximum current available for different distances.


5.5 Prototype

Once both the digital module and the analog front end have been individually tested, the performance of the whole system has to be checked. Not only the operation of the RFID tag will be validated but also the viability of the integration of a sensor comprised in the tag will be tested. Figure 18 shows the setup of the prototype, where some previous considerations thereof should be taken into account.

5.5.1 Digital Core Power Consumption

As the digital core for the prototype has been implemented on an FPGA, it has to be supplied externally. In order to obtain accurate results, the power consumption of the digital core has been emulated by means of the variable load shown in Fig. 19. This load is also controlled by the FPGA with additional control logic that can be

Fig. 19 Programmable variable load



suppressed in the final version for ASIC. With this programmable load, the power consumption profile shown in Fig. 13 is emulated. The values of the programmable load have been selected taking into account the measured power consumption of the previous ASIC version of the digital core, which has an average power consumption of around $5\mu\text{W}$.

5.5.2 Supply Capacitor

As it has been aforementioned, the supply capacitor is in charge of storing the necessary energy for supplying the rest of the circuits of the tag. In this case, it has been integrated in the analog front end. Due to area limitations, a 1.4 nF supply capacitor has been integrated. In order to check whether the size of the supply capacitor is a real limitation or not, a set of discrete capacitors is added in parallel to V_{SUPPLY} .

5.5.3 Sensor

As standard communication protocols have been implemented on the digital core, commercial sensors can be attached to the prototype. This way, the same system can be analyzed focusing on different sensors and applications.

5.5.4 Wireless Sensor Prototype

The final implementation of the prototype is shown in Fig. 20. The system is composed by three PCBs:

1. Front-end PCB. It is composed by an antenna, a front-end IC, a 2.5 V LDO, and an external capacitor of $100\mu\text{F}$.
2. FPGA PCB. It is the PCB that includes the digital part of the RFID and the digital interface for the sensor.
3. Interface PCB. It is composed of the connections between the front end and the FPGA, the socket sensor, an auxiliary 1.2 V and 2.5 V LDOs to supply the level shifters between the FPGA and the FE with the sensor and the emulator of the power consumption of the digital core.

Fig. 20 Implemented wireless sensor prototype scheme

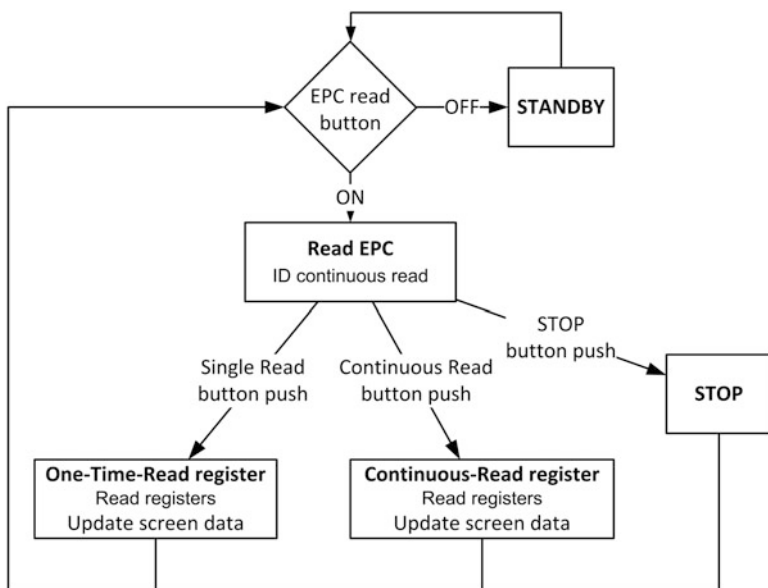



Fig. 21 Software flow chart

5.6 Reader Software Implementation

The reader used in the evaluation is the MOTOROLA handheld MC9090 at 30 dBm. Considering the output power level of the reader there is a limitation of 3 dB if compared with the previous analysis. European regulation allows transmitting up to 2 W EIRP. Additionally, given the characteristics of the APIs provided by the vendor, there is no possibility to have a continuous RF emission without sending commands. These restrictions limit the performance of the prototype and the achieved communication distance.

Figure 21 shows the flow chart diagram of the software. The program starts in a standby state. With a button push the software enters in the EPC read state and the

reader interrogates the tags continuously asking for the ID. This is the closest state to a constant RF emission. This state is used to keep the tags alive harvesting the energy from the RF wave.

When the single read button is pressed, the reader executes the necessary commands to obtain a unique and current value of the sensor. Once the measurement has been done, the software returns to the read EPC in order to keep powering the tags.

Depending on the application, a continuous sensor interrogation can be required. This can be achieved entering the continuous read state. In this state the reader interrogates the tags continuously asking for the sensor data. Pushing the stop button the reader returns to the EPC reads; if EPC read state is disabled, the reader returns to the initial standby state. This is the end of the communication, and as no more power is radiated the tags will shut down.

6 Results

In this section the results regarding the feasibility of the complete sensory system are presented. The results are shown using the previously mentioned setups and the described prototype.

6.1 *Front End and Digital Core Results*

The total power consumption of the analog front end is $7.4\mu\text{A}$. The voltage signals presented in Fig. 22 show the analog front-end performance. Any block is not activated until POR_1 circuits enable them. The reset circuit enables the digital core $175\mu\text{s}$ after POR_2 activation. It has to be taken into account that this measurement has been performed on wafer, without the matching network. That is the main reason why the charge slope is slow once the bandgap is enabled.

In order to test the communication between reader and tag, the digital core, implemented in an FPGA, has been connected to the front end and tested. Figure 23 shows the ASK signal (top), extracted from the demodulator, and the response of the digital module, LM signal (bottom). These measurements have been made taking into account a matched antenna with an impedance of $7.8 + 47.7j\Omega$. The load modulator is able to provide phase differences of 85° when answering the received command from the reader.

In order to study the maximum distance at which the front end can afford different output currents (I_{OUT}) or output power consumptions ($P_{\text{OUT}@2.5\text{V}}$), the setup detailed in Sect. 5.5 has been used. In order to characterize the system, the sensor has been replaced by a load that produces a constant current consumption.




Fig. 22 Start-up measurement of the analog front end




Fig. 23 ASK (top) and LM (bottom) signals

The results that have been obtained are shown in Fig. 24. An average output current of $15\mu\text{A}@2.5\text{V}$ ($P_{\text{OUT}} = 37.5\mu\text{W}$) can be managed at 2.4 m from the reader using a 2 W EIRP output power reader. This available power at 2.4 m is enough to supply low power sensors [13]; therefore sensor networks using long-range battery-less RFID sensor nodes are possible.




Fig. 24 Harvested current and power as a function of range

6.2 System Results

This section provides the obtained results of the RFID sensor system working as wireless temperature sensor and a wireless accelerometer. In order to demonstrate the feasibility of the RFID sensor, once the front end and the digital core have been measured and analyzed, the analog front end is connected to the FPGA, which has been programmed with the digital code, the digital interface needed to interface with the sensor, and also the controls of the resistors array to emulate the power consumption of the digital core. The implemented prototype is shown in Fig. 25.

Figure 25 shows the prototype operation; in the oscilloscope the ASK demodulated signal (up) and modulation signal (down) from the load modulator are shown. The modulation signal that carries the answer from the tag to the reader includes two different sequences: the first related to the tag identification and the second one related the requested sensor measurement.

The RFID sensor operates from wireless power at a distance up to 2 m offering a strong communication. Temperature and acceleration applications are presented. From the reader part some screen shots (Fig. 26a, b) are shown during the operation of the system. The first one (Fig. 26a) shows a screen shot of the reader in continuous mode, monitoring the temperature of the tag up to a distance of 2 m. The second picture (Fig. 26b) shows the monitoring of the 3-axis accelerometer of the tag up to a distance around 1 m.

The reduced output power of the commercial readers (30 dBm instead of 33 dBm, maximum allowed by the European regulation) and the noncontinuous emissions limit the operation distance of the wireless sensors. With fixed readers,




Fig. 25 RFID sensor prototype working with full communication

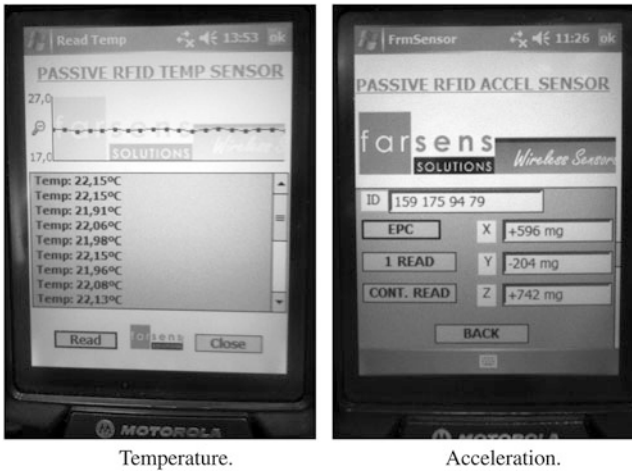


Fig. 26 Screenshot of the reader software working with commercial sensors. (a) Temperature. (b) Acceleration

the communication range of the wireless sensors can be increased. Figure 27 shows the environment and the complete setup of measurements.

Figure 27 shows the prototype of the RFID sensor in the course of the operation of a temperature sensor. The tests have been done acquiring sensor data with a commercial reader verifying the compatibility with the standard EPC C1G2.




Fig. 27 Complete setup and environment of the wireless sensor

7 Conclusions

This contribution has presented the design and experimental results of a battery-less RFID sensor device compatible with the EPC C1G2 protocol. The main theoretical limitations involving maximum communication distance between the tag and the reader are discussed obtaining useful system design guidelines. Using these guidelines a low power analog front end is designed and fabricated in a low-cost $0.35\mu\text{m}$ CMOS process. In order to maximize the communication range of a passive RFID sensor, the digital core of the tag has been optimized with a deep analysis of the dynamic power consumption. A complete wireless sensory system is implemented assembling the analog front-end chip to a matched dipole antenna, to an ultra-low-power commercial sensor, and to a module-based digital core (field-programmable gate array and digital core replica power consumption module). Measured results show a successful wireless communication up to 2.4 m from a 2 W EIRP output power reader to a digital module plus low power sensor (temperature, pressure, humidity, etc.) with average power consumption lower than $37.5\mu\text{W}$. Temperature and acceleration prototypes have been built showing communication ranges of 2 m and 1 m, respectively, using a commercial reader. These characteristics allow the use of the proposed sensory system in a battery-less wireless sensor network.

Acknowledgments This research work has been funded by CEIT, TECNUN, and FARSENS S.L., whose support is gratefully acknowledged. The authors would also like to thank the research teams

from previous institutions for supporting this work. Special thanks to: Andrés García-Alonso, Juan Francisco Sevillano, Igone Vélez, Ainhoa Cortés, Daniel Pardo, Alexander Vaz, Aritz Ubarretxena, Héctor Solar, Iñigo Gutiérrez, Ainara Jiménez, Marcos Losada, Josean Gómez, and Iker Mayordomo for their many valuable discussions and support.

This research work has been sponsored by the Spanish Ministry of Education and Science with the Torres Quevedo Grants no. PTQ0803-08788 and no. PTQ0901-00527 and the project TEC2011-29148-C02-02.

References

1. IDTechEx, "RFID Market Projections 2008 to 2018," *Technical Representative*, January 2007.
2. "New Low-Cost Temperature Sensor," *RFID Journal*, July 2002 [Online]. Available: <http://www.rfidjournal.com/article/view/28/1/1>.
3. K. Opasjumruskit et al., "Self-powered wireless temperature sensors exploit RFID technology," *IEEE Pervasive Computing*, vol. 5, no. 1, pp. 54–61, Jan.-Mar. 2006.
4. V.K. Chan and E. Mejia, Interrogation Device and Method for Scanning, US Patent#7,432,825, October 2008.
5. H. Shen, L. Li, Y. Zhou, "Fully integrated passive UHF RFID tag with temperature sensor for environment monitoring," *7th International Conference on ASIC*, pp. 360–363, October 2007.
6. N. Cho et al., "A 5.1- μ W, UHF RFID tag chip integrated with sensors for wireless environmental monitoring," *ESSCIRC*, pp. 279–282, 2005.
7. D. Yeager, F. Zhang, A. Zarrasvand and B. P. Otis, "A 9.2 μ A Gen 2 Compatible UHF RFID Sensing Tag with -12dBm Sensitivity and 1.25 μ Vrms Input- Referred Noise Floor," *IEEE International Solid-State Circuits Conference (ISSCC) - Digest of Technical Papers*, pp. 24–26, February 2010.
8. EPCGlobal, EPCTM Radio-Frequency Identity Protocols Class-1 Generation-2 UHF RFID - Protocol for Communications at 860MHz - 960MHz. V1.2.0, October 2008.
9. ISO/IEC 18000-6 – Information technology – Radio frequency identification for item management – Part 6: Parameters for air interface communications at 860 MHz to 960 MHz International Organization for Standardization April, 2011.
10. G. De Vita and G. Iannaccone, "Design criteria for the RF section of UHF and microwave passive RFID transponders," *IEEE Transactions Microwave Theory and Techniques*, vol. 53, no. 9, pp. 2985–2989, 2005.
11. J.F. Dickson, "On-chip high-voltage generation in NMOS integrated circuits using an improved voltage multiplier technique," *IEEE Journal Solid-State Circuits*, vol. 11, no. 3, pp. 374–378, June 1976.
12. U. Karthaus and M. Fischer, "Fully integrated passive UHF RFID transponder IC with 16.7 μ m minimum RF input power," *IEEE Journal Solid-State Circuits*, vol. 38, no. 10, pp. 1602–1608, October 2003.
13. A. Vaz et al., "Long range, low power UHF RFID analog front-end suitable for battery wireless sensors," *IEEE MTT-S International Microwave Symposium*, 2010.
14. D. Pardo et al., "Design Criteria for Full Passive Long Range UHF RFID Sensor for Human Body Temperature Monitoring," *IEEE International Conference on RFID*, 2007.
15. R. Barnett, G. Balachandran, S. Lazar, B. Kramer, G. Konnail, S. Rajasekhar and V. Drobny, "A passive UHF RFID transponder for EPC gen 2 with -14dBm sensitivity in 0.13 μ m CMOS," *IEEE International Solid-State Circuits Conference (ISSCC) - Digest of Technical Papers*, pp. 582–583, Feb. 2007.
16. A. Ricci, M. Grisanti, I. De Munari and P. Ciapolini, "Design of a 2 μ W RFID Baseband Processor Featuring an AES Cryptography Primitive," *Proceedings of the 15th IEEE International Conference on Electronics Circuits and Systems (ICECS)*, pp. 376–379, 2008.

17. V. Roostaie, V. Najafi, S. Mohammadi and A. Fotowat-Ahmady, "A low power baseband processor for a dual mode UHF EPC Gen 2 RFID tag," *International Conference on Design and Technology of Integrated Systems in Nanoscale Era*, 2008.
18. S. Wanggen, Z. Yiqi, L. Xiaoming, W. Xianghua, J. Zhao and W. Dan, "Design of an ultra-low-power digital processor for passive UHF RFID tags," *Journal of Semiconductors*, vol. 30, no. 4, 2009 [Online]. Available: <http://iopscience.iop.org/1674-4926/30/4/045004>.
19. P. Alanson, Sample et al., "Design of a passively-powered, programmable sensing platform for UHF RFID systems," *IEEE International Conference on RFID*, 2007.
20. I. Zalbide, I. Vález et al., "Power and energy optimization of the digital core of a GEN2 long range full passive RFID sensor tag," *IEEE International Conference on RFID*, 2008.
21. D.J. Yeager, P.S. Powledge, R. Prasad, D. Wetherall and J.R. Smith, "Wirelessly-charged UHF tags for sensor data collection," *IEEE International Conference on RFID*, pp. 320–327, 2008.
22. European Telecommunications Standards Institute (ETSI), EN 302 208: Electromagnetic compatibility and radio spectrum matters (ERM) Radio-frequency identification equipment operating in the band 865 MHz to 868 MHz with power levels up to 2 W.

Part III
Communication and Tools

Passive RFID-Based Wake-Up Radios for Wireless Sensor Networks

He Ba, Jeff Parvin, Luis Soto, Ilker Demirkol, and Wendi Heinzelman

Energy efficiency is one of the most important criteria in the design of a wireless sensor network (WSN). Sensor nodes are usually battery-powered and thus have very limited lifetime if no energy conservation method is applied. Radio transmission and reception are the two major sources of energy drain, and when a sensor node is active and waiting to receive data, it wastes energy on idle listening. To extend the lifetime of a sensor node, its radio can be turned off and its microcontroller can be set into a deep-sleep mode when it is idle, and the radio can be woken up when there are transmissions destined to the sensor node. To wake up a sensor node from its sleep mode, there are generally two approaches: (1) *duty cycling*, where the sensor node sets a timer and wakes up with the interrupt of this timer, and (2) *wake-up radio*, where the sensor node wakes up with an external interrupt.

In duty cycling, nodes are periodically set into the sleep mode and therefore save a significant amount of energy at the expense of data latency. With lower duty cycle, nodes will consume less energy, but this will introduce higher latency for data delivery. Besides the energy-performance trade-off, it is important to consider that the node loses its functionalities while sleeping. Thus, nodes will have no information about what has happened in the network prior to the time that they wake up and listen to the channel. This introduces complexities and other overhead to the MAC protocols.

H. Ba (✉) • L. Soto • W. Heinzelman
University of Rochester, Rochester, NY 14611, USA
e-mail: ba@ece.rochester.edu; luis.soto@rochester.edu; wendi.heinzelman@rochester.edu

I. Demirkol
Universitat Politècnica de Catalunya, Barcelona, Catalunya, 08034, Spain
e-mail: ilker.demirkol@entel.upc.edu

J. Parvin
University of Pittsburgh, Pittsburgh, PA, USA
e-mail: j3ff.parvin@gmail.com

Using radio wake-up techniques, such complexities and overhead can be reduced. A wake-up signal triggers a node to wake up from the deep-sleep mode and start transmission activities. Normally, the wake-up signal is sent or received by a secondary radio transceiver. In order to improve energy efficiency, the energy consumption of this extra wake-up radio receiver should be extremely low or, ideally, zero. The energy benefit of using radio wake-up in comparison with duty cycling is that nodes do not waste energy on idle listening, since they are only awakened by neighboring nodes when there is a request for transmission. This on-demand communication scheme can also help to decrease the packet latency, which is especially important for event-triggered network applications. In addition, using a wake-up signal reduces the overhead in control traffic, e.g., for synchronization or collision avoidance (RTS/CTS packet exchange).

Recently, several studies have focused on the design of wake-up radio receivers, both on the hardware design and on the protocol implementation aspects. In this chapter, we first review the state of the art in both active and passive wake-up radio receivers, followed by a discussion of potential applications of passive wake-up receivers, i.e., the type of wake-up receiver that is the focus of this chapter. One possibility for the wake-up radio is to use passive RFID as the wake-up technology, as there are off-the-shelf passive RFID tags and readers readily available. We present the design of a passive RFID wake-up device. A downside of using passive RFID tags for the wake-up functionality is their limited communication range, which results in limited wake-up range. To characterize the performance of our passive RFID wake-up device, we conduct field tests to determine the energy consumption and the wake-up probability as a function of distance in different environments. Finally, potential future work on this topic is discussed.

1 The State of the Art in Wake-Up Radio Receivers

Wake-up radios can be categorized as active wake-up radios and passive wake-up radios, depending on their energy sources [1]. An active wake-up radio has a better wake-up range than a passive wake-up radio; however, the active wake-up radio requires continuous power supply. On the other hand, a passive wake-up radio operates within a relatively smaller range, but it does not require an external power supply. A passive wake-up receiver harvests energy to power itself from the wake-up signal transmitted by the sender. Based on the way the wake-up signal destinations are defined, the wake-up radios can be classified as broadcast-based wake-up, where all nodes that receive the wake-up signal are awakened, and ID-based wake-up, where only targeted nodes are awakened. In this section, we review the state of the art of both active and passive wake-up radio receivers. Table 1 provides a summary of these different wake-up receiver design techniques and a comparison of the energy consumption, sensitivity, and implementation status of these wake-up radio receivers.

Table 1 Proposed wake-up receiver comparisons

Active wake-up receivers	Main technique	Frequency	Power consumption	Sensitivity	Implementation
Otis et al. [2]	Super-regenerative with BAW matching network	1.9 GHz	400 μ W	-100.5 dBm	Yes
Pletcher et al. [3]	BAW matching network	1.9 GHz	65 μ W	-50dBm	Yes
Le-Huy et al. [4]	Zero-bias Schottky voltage doubler	2.4 GHz	19 μ W	-50dBm	No
Von der Mark et.al [5]	Three-stage wake-up, zero-bias Schottky diode used at first stage	N/A	nWs for first stage and μ Ws for second stage	N/A	No
Ansari et al. [6]	Five-stage charge pump	N/A	2.628 μ W	N/A	Yes
Van der Doorn et al. [7]	Commercial filter and amplifier	868 MHz	171 μ W	-51 dBm	Yes
Austria Microsystems [8]	N/A	15-150 kHz	8.1 μ W	-37dBm	Yes
Passive wake-up receivers	Main technique	Frequency	Power consumption	Sensitivity	Implementation
Gu et al. [9]	Zero-bias Schottky diode	433 MHz	0 or 3.69 μ W	N/A	No
Ba et al. [23]	Passive RFID (WISP)	902-928 MHz	0	-10dBm	Yes

N/A: Data not available

1.1 Active Wake-Up Radio Receivers

Several different low-power active wake-up receivers have been proposed [2–8]. In [2], B. Otis et al. propose the use of a super-regenerative architecture with a 1.9 GHz bulk acoustic wave (BAW) resonator to reduce the power consumption of the wake-up radio. The power consumption of this radio is 400 μ W for the receiver and 1.6 mW for the transmitter. This approach is further optimized to create a 65 μ W wake-up receiver [3], using a 1.9 GHz BAW resonator matching network for RF signal filtering. This wake-up receiver can provide a sensitivity of -50 dBm at 40 kbps and -48 dBm at a maximum data rate of 100 kbps.

A different approach is developed by Le-Huy and Roy [4] and Von der Mark et al. [5], where zero-bias Schottky diodes are used because they have no bias current through the diodes. The low-power 2.4 GHz wake-up receiver proposed in [4] is designed to work with a directional antenna and pulse-width modulation in order to reduce energy dissipation. Simulation results show that the receiver can reach -50 dBm sensitivity with only 19 μ W power consumption. A three-stage wake-up scheme is introduced in [5]. In this approach, a very low power (on the order of nW) always-on stage is used to trigger an intermediate higher power (on the order of μ W) stage for wake-up signal verification. Only if the wake-up signal is confirmed is the main transceiver activated.

Other approaches for active wake-up radios are described in [6, 7]. Junaid et al. propose a wake-up receiver including a five-stage charge pump used to increase the received signal voltage [6]. The only active parts of the wake-up circuit are the digital comparator and the voltage divider, which consume 350 nA and 526 nA, respectively. In [7], Van der Doom et al. implement a wake-up receiver using only commercial components to reduce the extra hardware costs. Their design consumes 171 μ W with -51 dBm sensitivity.

Although there are several hardware proposals for active wake-up radios, not many physical implementations or commercialized products are available today. Recently, Austria Microsystems announced their latest 3-channel low-frequency wake-up receiver working at 15–150 kHz [8]. This product consumes 8.1 μ W and can reach a sensitivity of about -37 dBm (as calculated from the provided specifications).

1.2 Passive Wake-Up Radio Receivers

Compared with active wake-up receivers, passive wake-up receivers do not require energy from a physically connected power supply; instead, they harvest energy from the transmitted wake-up signal. While this makes passive wake-up radios energy efficient, the wake-up range for passive wake-up radios is relatively shorter.

Currently, there are limited studies on passive radio wake-up receivers. L. Gu et al. propose a passive radio wake-up circuit that theoretically could operate at

a range of 10 feet with 5 ms latency, according to SPICE simulation results [9]. If a comparator and an amplifier are added, which respectively consume negligible currents of 350 and 880 nA, the radio could theoretically reach up to 100 feet with 55 ms latency.

A performance study on the use of passive RFID wake-up radios is given by Jurdak et al. [10, 11]. In their work, an RFID wake-up mechanism is proposed, namely RFIDImpulse, which assumes a commercial RFID reader and a passive RFID tag are attached to each sensor node, providing radio wake-up capability. The performance of the proposed mechanism is investigated through MATLAB simulations and compared with BMAC and the IEEE 802.15.4 standard. Their results show that RFIDImpulse outperforms both other methods in terms of energy efficiency and transmission rate for low and medium traffic scenarios. However, the analysis is based on an important assumption that all nodes have the capability to wake up their neighbors, which is not feasible currently, due to the considerable amount of energy required by the RFID reader, which sends out the wake-up signal, and its large size. In addition, their energy consumption analysis does not include the energy consumed by the nodes to wake up. In reality, the wake-up energy consumption includes the energy used for MCU boot-up and for radio initiation, which could be comparable to the energy consumed for radio transmission. Later in this chapter, a passive RFID-based wake-up device named WISP-Mote will be introduced. To the best of our knowledge, the WISP-Mote is the first reported complete implementation of a passive radio wake-up device.

2 Applications of Passive Wake-Up Radios for WSNs

The applications of passive wake-up radio sensor networks are determined by the characteristics of the network in terms of transmission range, asymmetric energy consumption values, and the equipment cost of the receiver and the transmitter. The power consumption of the wake-up transmitter is the major obstacle to realizing multi-hop passive radio wake-ups, as typically the battery-powered sensors cannot afford to transmit a high-power wake-up signal. Hence, considering the hardware cost and energy constraints, it is expected that there will only be a few powerful nodes in a sensor network that have the capability to wake up other nodes, in realistic applications. Even if all the sensors are assumed to have enough energy, the wake-up range is relatively short due to path loss and the low efficiency of power harvesting at the receiver. As a result, to cover a large area of sensor nodes, either the special wake-up signal transmitter has to be mobile to wake up the sensor nodes and collect data, or the sensor nodes have to be mobile to move to the base station to deliver data. Based on these features and constraints, we present several potential real-world applications that can benefit from passive wake-up sensor networks.

Nowadays, sensor networks are widely used for environmental observation and habitat monitoring. Such applications usually require long-term operations and therefore the longevity of the sensor nodes is important. Wildlife monitoring

is one of the potential applications for passive radio wake-up sensor networks. A branch of zoology research investigates the behavior of a species or interactions between species. It is important to gather information on individual animals such as their location or physiological data, as well as environmental information such as temperature and humidity, to understand the effects of the environment and influences from other species. The cost for collecting data by equipping animals with sensor nodes is much lower than the cost of other approaches, e.g., volunteers. However, it is usually difficult for scientists to retrieve the sensor nodes from wild animals for battery replacement or recharging individually. Hence, energy-efficient operation of sensor nodes is crucial, which is the main motivation for using the radio wake-up technique. Scientists may put sensor nodes on the animals, e.g., using a sensor collar, as done by researchers in the ZebraNet project [12], and place data collectors (equipment including a wake-up signal transceiver, data transceiver, and large energy supply) at places where the animals are expected to congregate, such as ponds and rivers, or where the animals are expected to roam. When the animals get close to the data collector, the radio on them will be awakened and start transmitting the gathered sensor data to the data collector. Since the sensor nodes may last for years, the battery replacement cost can be saved.

Another potential application is patient monitoring in long-term care facilities, e.g., sanatorium or assisted living locations. In these places, patients are cared for over a relatively long period of time and require nonemergency but long-term health monitoring (e.g., the data from the patients only need to be collected by nurses several times per day). Instead of using traditional health monitors, which connect to the patient's body with wires limiting his/her daily activities, each patient can be equipped with various wireless sensor nodes for collecting different physiological information, e.g., heart rate, blood pressure, blood glucose, and insulin level. With all the equipment (the wake-up transmitter, the data receiver, the computer, and the power source) carried in a trolley, the nurse can easily gather all the sensor data from patients while periodically traveling throughout the patients' rooms. This scenario is similar to a Data MULE scenario [13] where a MULE plays the role of a data collector. In this scenario, the ID-based wake-up will be beneficial, since there are multiple nodes on a patient's body and possibly multiple patients in a room. If a broadcast-based wake-up is used, all of the nodes of all patients will be awakened at the same time, which would cause congestion. Therefore, waking up the nodes one by one based on their type and/or patient ID will greatly reduce the chances of collisions and hence extend the sensor lifetime.

Intelligent transportation system (ITS) is a very promising application that aims to improve traffic safety while providing other benefits such as reducing fuel consumption and latency. Using transportation facilities to collect sensor data throughout the city, i.e., *urban monitoring* [14], can be part of the ITS. The sensor data may include traffic conditions and weather information that will be helpful for congestion avoidance and improving traffic safety. Battery-powered sensor nodes can be easily attached to traffic lights, bus stations, toll booths, and traffic signs

and may work for years due to the energy efficiency improved by using wake-up radios. Buses, trains, and garbage trucks are good candidates for performing data collecting duties as they can cover a fairly wide area [15, 16] and they have no energy limitations so that transmitting a powerful wake-up signal is not a problem.

Energy waste is becoming a global concern nowadays. Wake-up radios have the potential to improve the energy efficiency of other devices as well. For example, in a public place with a wireless access point, there may not be Wi-Fi users all the time, but the router is constantly consuming energy even when it is in low-power mode. Using a wake-up radio receiver as a remote switch to the access point and attaching the wake-up transmitter to the user's device such as a laptop computer or a smart phone, the Wi-Fi user will have the capability to wake up the access point only when they want to start a connection; otherwise the access point will be set in deep-sleep mode to save energy. This approach will improve the energy efficiency of the device without sacrificing the throughput or response delay.

3 Design and Characterization of a Passive RFID Wake-Up Sensor Device

The potential applications presented previously inspire us to design a passive wake-up sensor device, with which the sensor lifetime can be greatly extended. In this section, we introduce a physical implementation of a passive RFID wake-up device, which we call a WISP-Mote. The WISP-Mote is created by integrating a WISP tag, an RFID tag developed by Intel research [17], and a Tmote Sky sensor node [18]. We characterize the WISP-Mote's performance by measuring its energy consumption values for different operations and testing its wake-up range in different environments.

3.1 Design and Implementation of the WISP-Mote

To wake up a sensor node, a trigger signal must be sent to the microcontroller (MCU) of the sensor node. The wake-up receiver is used to receive the wake-up radio signal and then send a trigger wave to the MCU. We employ Intel WISP (wireless identification and sensing platform) as an external wake-up signal receiver for the Tmote Sky mote (see Fig. 1). A WISP is an RFID tag with sensing and computing capabilities, developed by Intel research. Using energy harvesting, a WISP can be powered remotely by a UHF RFID reader. In our implementation, we use a UHF Gen2 Speedway RFID reader from Impinj, as shown in the testbed setup in Fig. 2. The RFID reader sends a continuous wave along with commands according to the C1G2 protocol [19] to the tag. The tag sends data back by




Fig. 1 A WISP-Mote



Fig. 2 Test setup

modulating the reflection coefficient of the backscattered signal (via changing the antenna impedance) [20]. The tag also collects the energy sent by the reader and regulates it to a proper DC value as its power supply. Therefore, using a WISP as a wake-up receiver does not consume extra power from the sensor node battery. The Tmote Sky mote is a wireless sensor node developed by UC Berkeley. Both devices' specifications are shown in Table 2. By combining an Intel WISP with a Tmote Sky mote, we create our passive wake-up mote.

Table 2 Device specifications

	WISP 4.1DL [21]	Tmote Sky [22]
Microcontroller	TI MSP430 F2132 (512B RAM, 8K+256B Flash)	TI MSP430 F1611 (10k RAM, 48k Flash)
Transceiver features	900 MHz dipole antenna	250 kbps 2.4 GHz IEEE 802.15.4 Chipcon wireless transceiver
Sensors	Accelerometer, temp	Humidity, temperature, light
Wake-up time		Fast wake-up from sleep (<6 μ s, typically 292 ns)

We use the WISP to generate an interrupt signal to the mote. In broadcast-based wake-up, whenever the WISP harvests enough energy using the reader's continuous wave signal, it sends a pulse to wake up the mote from the sleep state. Thus, any WISP-Mote within range of the reader will be awakened. By default, the operations that require energy on an RFID tag are the modulation and demodulation of the signals sent and received, and the processing of the data. Since in our RFID wake-up system, there is no need for tag-to-reader transmission, we disable the energy consuming operations on the WISP that are used to respond to commands sent from the reader. As a result, the only responsibility of the WISP is to harvest energy from the reader and to send the interrupt signal to the mote. By minimizing the energy required for the WISP, we maximize the wake-up range of the WISP-Mote.

In dense networks, the broadcast-based wake-up results in a high number of collisions in the communication channel, since all nodes within the wake-up range are awakened by the reader and attempt to send their data. Certain applications define the wake-up of a specific node or a class of nodes, which reduces unnecessary wake-ups and the number of collisions, compared to broadcast-based wake-ups. For such applications, we programmed the WISP to generate an interrupt for the mote only after receiving a packet that includes its ID or class number. This functionality requires the demodulation of the received signal and additional computations by the MCU, both of which require extra energy to be harvested by the WISP. Therefore, the wake-up range for the ID-based scheme is expected to be smaller than that for the broadcast-based scheme. We performed field tests of both broadcast-based and ID-based wake-up in various environments to measure the corresponding wake-up ranges.

For both the broadcast-based and ID-based wake-up schemes, higher energy efficiency can be achieved if the motes do not wake up when they have no data to send, e.g., in the wildlife monitoring scenario, when the animal is close to the data collector at the pond and has already delivered all the buffered data. To achieve this, we set up the mote such that it disables the interrupt from the input port of the WISP when it has no data to send. As a result, the node will wake up only when it has buffered data and receives a wake-up signal. Otherwise, the node will remain in the low-power sleeping state. This method reduces unnecessary energy waste and improves the energy efficiency significantly.

Table 3 Current consumption measurements of a Tmote Sky node

Operation	Average current consumption (mA)	Average power consumption (mW)	Duration (ms)
Wake-up	10.4	31.2	5
Transmit 12 byte packet	18.2	54.6	30
Receive and idle listening	20.2	60.6	
Sleep	0.2	0.6	

3.2 WISP-Mote Energy Consumption Characteristics

Unlike active wake-up radios that constantly consume power, passive RFID wake-up radios do not consume any energy from the sensor node. Hence, the only energy consumption of the WISP-Mote is caused by the mote components. The Tmote Sky datasheet [22] provides the current consumptions in typical operating conditions for the mote, but it lacks information about the current consumption and the duration of the booting process, which is essential for the energy consumption analysis of RFID wake-up sensor networks. In addition to current consumptions in transmission and reception, we measured the current and time consumed in booting and radio initiation for the wake-up operation. The energy consumption characteristics of a Tmote Sky mote are extracted by measuring the current consumptions of different operations along with the duration of these operations. The measurement results are presented in Table 3 and are consistent with those from the Tmote Sky datasheet. A battery voltage of 3 V is considered in the calculations, since motes use 2 AA batteries. The results show that besides radio transmission and reception, the mote's booting also consumes energy that cannot be ignored. These measurements support the need for an accurate energy analysis for sensor networks.

3.3 Field Testing

There are two main factors that have a very large effect on the wake-up probability, i.e., on the probability that a node can be awakened in a real-world implementation. The first factor is the distance from the RFID reader that transmits the wake-up signal. As the range between the reader and the tag (WISP) increases, the energy that the WISP can harvest decreases due to path loss. The second factor is the environment that the system is in. The environment plays a critical role, because reflections can have large effects on the wake-up probabilities at different locations within the three-dimensional space around the reader. In certain locations, there are destructive effects, causing dead zones where the node is unable to be awakened. There will also be constructive effects of the reflections, resulting in locations with much higher wake-up probabilities than locations that are closer to the reader.

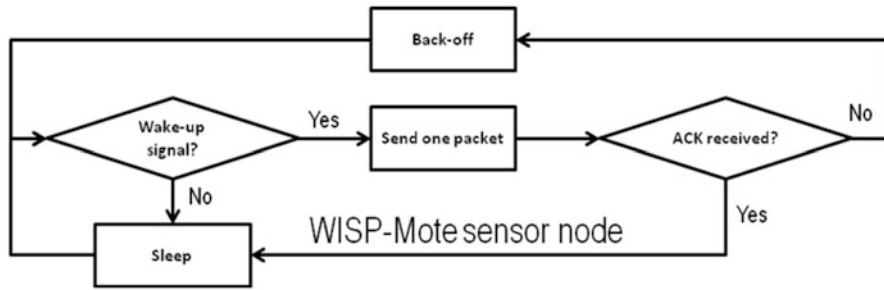


Fig. 3 WISP-Mote wake-up flow chart

Because the ID-based wake-up coding scheme has higher power consumption than the broadcast-based one, theoretically it should have reduced wake-up probabilities in all environments. The main goal of the tests we present in this section is to explore the wake-up probabilities observed for both schemes. We investigated three different environments with both ID-based and broadcast-based wake-up approaches to demonstrate the effects of environment, distance, and the wake-up scheme. We measure the wake-up probability as a function of distance to the reader and height from the reader in a hallway and an open-air environment, and as a function of distance in an office environment.

Experiment Setup

To ensure the succinctness of our experiments and the accuracy of the results, we set up our experiments as follows. The reader is raised off of the ground to a height of 84 cm (33 inches) to reduce reflective effects due to the ground close to the reader. Due to the necessity to vary both distance from the reader and height, the WISP-Mote was attached to a small tripod, which allows both *x*-axis movement, toward and away from the reader, and *y*-axis movement, up and down with respect to the reader.

For the broadcast-based wake-up scheme, the data collector broadcasts a generic wake-up signal that causes every WISP-Mote receiving this signal to accumulate energy to wake up. After waking up, the WISP-Mote contends to send a packet to the data collector, which also includes a wake-up count. After the delivery, the WISP-Mote promptly returns to sleep. In real applications, acknowledgements and back-offs may be introduced for stable communications, as illustrated in Fig. 3. In our single WISP-Mote field tests, no ACKs or back-offs are implemented, since the backward link quality is good (mote-to-data collector) and no contention exists.

To calculate the wake-up probability, we set up the system as follows. Due to the difficulties to control the commercial RFID reader to send the wake-up signal periodically, we use the wake-up counts from the nodes to calculate the wake-up probability. To count the number of times that a WISP-Mote is able to be awakened,

we periodically enable the interrupt from the wake-up signal input port once every 0.5 s and disable it after the mote is awakened from the WISP. As a result, WISP-Motes can wake up as long as the wake-up receiver is harvesting enough power and the interrupt is enabled. The reader is then run for a fixed amount of time, in this case 100 s, and the final wake-up count sent by the WISP-Mote is recorded. We can then calculate the wake-up probability of the investigated scenario by dividing the observed number of wake-ups to the total possible wake-ups, i.e., 200 wake-ups. The ID-based wake-up experiments were conducted in an identical manner, except that instead of simply broadcasting a generic wake-up signal, the data collector transmits an ID, against which the WISP compares its own ID before deciding whether to wake up the mote or not. This additional demodulation and computation require an additional amount of power to be harvested.

Wake-Up Probabilities for Different Environments

1. Open Environment

The purpose of the open environment is to test a location with as few reflecting surfaces as possible. This environment could be a large indoor area, or an outdoor area, but for testing purposes a large gymnasium was used.

Figure 4 shows the test results for the broadcast-based wake-up in the open environment.¹ The y -axis of the graph shows the vertical distance between the WISP-Mote and the data collector, and the x -axis shows the horizontal distance between the WISP-Mote and the data collector. The data collector is located at point (0, 0) on the graph. The colors at each point represent the wake-up probability at that particular point. As can be seen, the WISP-Mote has almost 100% wake-up probability for all points within 2.5 m, and after that point, reflections from the ground start to have a significant effect on the wake-up probability. These reflections appear to match the two-ray ground model, and there are clearly areas where there are effects of destructive interference that are creating dead zones, as well as areas of constructive interference that enable 100% wake-up probability far from those dead zones.

Figure 5 shows the test results for the same environment, but this time with the ID-based wake-up. There has been a decrease of the wake-up range from the 3 m of the broadcast-based scheme down to 2.5 m for the ID-based one at the height of 0.2 m above the data collector, as expected. However, generally the ID-based scheme has a similar performance to the broadcast-based scheme in this environment. The two-ray ground interference is also evident in the ID-based wake-up scheme.

¹ The detailed field measurement data can be found at: <http://www.ece.rochester.edu/research/wcng/code.html>.




Fig. 4 Broadcast-based wake-up probabilities in an open environment




Fig. 5 ID-based wake-up probabilities in an open environment

2. Closed Environment

When such pronounced reflective effects are witnessed in an environment with only one reflective surface, it can be hypothesized that even more significant and far more unpredictable reflections will occur in a closed environment. The tests for a closed environment were done in a long hallway, meaning that the number of reflective surfaces was increased to four surfaces.




Fig. 6 Broadcast-based wake-up probabilities in a closed environment




Fig. 7 ID-based wake-up probabilities in a closed environment

As can be seen in Figs. 6, and 7, a closed environment vastly increases the number of reflections and therefore the amount of interference among the multiple copies of the signal sent. For the broadcast-based wake-up, 100% wake-up probability for all heights measured is only observed for very small distances, and for ID-based wake-up, there is no distance that offers 100% wake-up probability for all heights.




Fig. 8 Broadcast-based and ID-based wake-up probabilities in a cluttered environment

3. Cluttered Environment

In real-life applications, it is more likely that the sensor network will be deployed in a cluttered environment such as a warehouse, a grocery store, or a conservatory. For the cluttered environment, it was infeasible to set up the in-depth tests as for the previous two environments, due to physical obstacles. Instead, average wake-up probability was tested at different distances from the data collector. For each distance on the plots, ten measurements were taken in various spots within an office and then averaged.

Although not shown in the same detail as the previous figures, Fig. 8 does illustrate that even more interference is occurring in the cluttered environment than in the previous two environments for both broadcast-based and ID-based wake-up. The wake-up probability drops very quickly between 3 and 4 m, and no wake-up is recorded at or outside of the 4 m mark, anywhere in the room.

As the data show, the WISP-Mote has a relatively stable and long wake-up range in an open environment, a volatile but similar long range in a closed environment, and an unpredictable and short wake-up range in a cluttered environment.

Broadcast-Based Wake-Up Versus ID-Based Wake-Up

There are quite a few meaningful conclusions we can draw from the measurements about the performance difference between the ID and broadcast wake-up schemes,

and also about the effectiveness of the system as a whole. In every environment, comparing the average wake-up probabilities for distance from the reader, we can see that broadcast-based scheme results in slightly higher average wake-up probabilities. The advantage of ID-based wake-up is that it functions as broadcast wake-up, but with the capability of waking up a particular class of nodes or a specific node. If there are only specific nodes that need to be woken up, then ID-based wake-up can lead to significant energy savings, as it does not force all nodes to wake up and contend for the medium, and then process all the redundant data. In a scenario where all nodes serve the same purpose, the usefulness of having ID-based wake-up may be reduced, but in the case where different nodes have different sensors and information, it may lead to a very large savings in energy. In most scenarios, the very small sacrifice made in wake-up range is less crucial than the increased functionality and possible energy savings.

4 Conclusions and Future Work

The radio wake-up technique has several very promising features: very low energy consumption and the potential to reduce protocol complexities. However, on the downside, the short wake-up range limits its application. The directional antenna and beam forming technique mentioned in [4] provide a possible solution to extend the wake-up range without increasing the wake-up transmitter power. As another approach, Omni-ID [24] has developed a patented technology called “plasmatic structure” that captures incoming RF waves and creates a region of highly concentrated energy around the RFID tag. This technology makes the RFID tag extremely energy efficient and greatly extends its operational range. By applying this technique to the wake-up radio receiver design discussed in this chapter, the range problem may be reduced. On the other hand, using other energy-harvesting approaches can provide additional energy, such as using sources like kinetic, thermal, and light. These extra energy sources could be used to support the energy consumed by an active wake-up radio receiver to achieve a larger wake-up range.

Researchers have already proposed many partial or complete implementations of wake-up radio receivers. Although simulations have shown the benefits and feasibility of applying a wake-up radio to sensor network applications, as presented in [11, 23], no network level test results based on a real-life application are provided. On the other hand, protocols designed for wake-up radio sensor network implementations will be needed. Since multi-hop wake-up is not applicable yet due to the large energy consumption for the wake-up transmitter, MAC protocols that aim to improve link throughput, collision avoidance, and delay reduction will be needed. With improvements in wake-up range and with decreasing hardware costs, we can expect wake-up radio-based sensor network systems to be deployed in the near future.

References

1. Demirkol I, Ersoy C, Onur E (2009) Wake-up receivers for wireless sensor networks: benefits and challenges. *IEEE Wireless Communications* 16:88–96. doi: 10.1109/MWC.2009.5281260
2. Otis B, Chee YH, Rabaey J (2005) A 400 μ W-RX, 1.6mW-TX super-regenerative transceiver for wireless sensor networks. *IEEE ISSCC* 1:396–397. doi: 10.1109/ISSCC.2005.1494036
3. Pletcher N, Gambini S, Rabaey J (2007) A 65 μ W, 1.9 GHz RF to digital baseband wakeup receiver for wireless sensor nodes. *IEEE CICC* 539–542. doi: 10.1109/CICC.2007.4405789
4. Le-Huy P, Roy S (2008) Low-power 2.4 GHz wake-up radio for wireless sensor networks. *IEEE WIMOB* 13–18. doi: 10.1109/WiMob.2008.54
5. Von der Mark S, Kamp R, Huber M, Boeck G (2005) Three stage wakeup scheme for sensor networks. *SBMO/IEEE MTT-S* 205–208. doi: 10.1109/IMOC.2005.1579978
6. Ansari J, Pankin D, Mähönen P (2008) Radio-triggered wake-ups with addressing capabilities for extremely low power sensor network applications. *PIMRC* 1–5. doi: 10.1109/PIMRC.2008.4699501
7. Van der Doorn B, Kavelaars W, Langendoen K (2009) A prototype low-cost wakeup radio for the 868 mhz band. *IJSNET* 5:22–32. doi: 10.1504/IJSNET.2009.023313
8. Austria Microsystems (2010) AS3933 3-D low frequency RF wake-up receiver. <http://www.austriamicrosystems.com/Wake-up-receiver/AS3933>. Accessed 12 December 2010
9. Gu L, Stankovic JA (2005) Radio-triggered wake-up for wireless sensor networks. *Real-Time Syst* 29:157–182. doi: 10.1007/s11241-005-6883-z
10. Ruzzelli AG, Jurdak R, O’Hare GMP (2007) On the RFID wake-up impulse for multi-hop sensor network. *ACM SenSys*
11. Jurdak R, Ruzzelli AG, O’Hare GMP (2008) Multi-hop RFID wake-up radio: design, evaluation and energy tradeoffs. *ICCCN’08* 1–8. doi: 10.1109/ICCCN.2008.ECP.124
12. Juang P, Oki H, Wang Y (2002) Energy-efficient computing for wildlife tracking: design tradeoffs and early experiences with zebranet. *ASPLOS* 96–107. doi: 10.1145/605397.605408
13. Shah RC, Roy S, Jain S, Brunette W (2003) Data MULEs: modeling a three-tier architecture for sparse sensor networks. *IEEE SNPA* 30–41. doi: 10.1109/SNPA.2003.1203354
14. Murty RN, Mainland G, Rose I, Chowdhury AR, Gosain A, Bers J, Welsh M (2008) CitySense: An Urban-scale wireless sensor network and testbed. *IEEE THS* 583–588. doi: 10.1109/THS.2008.4534518
15. Lee U, Magistretti E, Gerla M, Bellavista P, Corradi A (2009) Dissemination and harvesting of urban data using vehicular sensing platforms. *IEEE TVT* 58:882–901. doi: 10.1109/TVT.2008.928899
16. Gil-Castineira F, Gonzalez-Castano FJ, Duro RJ, Lopez-Pena F (2008) Urban pollution monitoring through opportunistic mobile sensor networks based on public transport. *IEEE CIMSA* 70–74. doi: 10.1109/CIMSA.2008.4595835
17. Sample AP, Yeager DJ, Powledge PS, Smith JR (2008) Design of an RFID-based battery-free programmable sensing platform. *IEEE TIM* 57:2608–2615. doi: 10.1109/TIM.2008.925019
18. Polastre J, Szewczyk R, Culler D (2005) Telos: enabling ultra-low power wireless research. *IPSN* 364–369. doi: 10.1109/IPSN.2005.1440950
19. http://www.epcglobalinc.org/standards/uhfclg2/uhfclg2_1.2.0-standard-20080511.pdf. Accessed 12 December 2010
20. Want R (2006) An introduction to RFID technology. *IEEE Pervasive Computing* 5:25–33. doi: 10.1109/MPRV.2006.2
21. <http://wisp.wikispaces.com/>. Accessed 12 December 2010
22. Tmote Sky, <http://sentilla.com/files/pdf/eol/tmote-sky-datasheet.pdf>. Accessed 12 December 2010
23. Ba H, Demirkol I, Heinzelman W (2010) Feasibility and benefits of passive RFID wake-up radios for wireless sensor networks. *IEEE Globecom*
24. <http://www.omni-id.com/>. Accessed 12 December 2010

BAT: Backscatter Anything-to-Tag Communication

Andrés Molina–Markham, Shane S. Clark, Benjamin Ransford, and Kevin Fu

1 Introduction

Research has extended RFID tags into general-purpose batteryless computing devices and explored their computational capabilities, identifying the potential to create RFID-scale sensor networks [2, 3, 14, 20, 25]. The increasing capabilities of these tags enable them to implement applications that incorporate requests for data or services from other computer systems. For example, tags may need to check for software updates or obtain information such as weather or environmental conditions for some applications. Tags may also need to communicate with one another to accomplish collaborative tasks such as environmental monitoring or data aggregation.

Supply-chain RFID technologies were developed with narrower design goals in mind, primarily inventorying or data collection. The EPC Gen 2 protocol—the de facto standard for UHF RFID communication [9]—reflects these design goals by offering a limited set of commands not well suited for bulk data transfers [12]. The assumption that tags will not implement any functionality locally and will instead act as simple static identifiers in most cases also imposes other restrictions. Tags must conform to a rigid state machine allowing no time to sense or process data. They are instead required to respond to reader commands rapidly whenever in range.

Without protocols that present appropriate abstractions for richer applications, computational RFIDs (CRFIDs) must use limited resources to shoehorn these new applications into a suboptimal paradigm [14]. Unlike supply-chain RFID tags, CRFIDs have their own microcontroller units: they are capable of running their own application logic and managing their own memory and communication links. They are therefore able to participate in radio protocols that are more flexible than Gen 2.

A. Molina–Markham (✉) • S.S. Clark • B. Ransford • K. Fu
University of Massachusetts, Amherst, MA, USA
e-mail: amolina@cs.umass.edu; kevinfu@cs.umass.edu

The BAT protocol and software stack provides fast and flexible backscatter anything-to-tag communication. The key insight is that BAT separates tag applications from the networking stack by ensuring that the networking layer does not impose unnecessary constraints on abstractions confined within the tag. For instance, BAT separates memory management abstractions from the networking stack such that tag applications can efficiently store data in arbitrary locations without needing several interactions with the network stack. Further, we show that this logical detachment between the networking stack and applications does not impede secure anything-to-tag communication even using current CRFID prototypes.

2 Anything-to-Tag Communication

The prevailing protocol for supply-chain RFID systems, EPC Class 1 Gen 2 [9], is designed around the abstraction of RFID tags as remotely addressable memory. The Gen 2 commands a reader may use fall into two categories. Readers use *singulation* commands to search a tag population for a single tag; then they may issue *memory access* commands to read or write its protocol-defined memory banks.

The tags-as-memory abstraction is entirely appropriate for supply-chain applications, but it imposes several hindrances on applications running on CRFIDs. First, it restricts communications to those that can be formulated in terms of reader-to-tag memory access commands and simple tag responses. It is possible to reuse fields of existing commands to carry information—for instance, the Gen 2 *write* command takes an arbitrarily long *WordPtr* parameter intended to specify where to start writing—but reader support for embedding arbitrary data in these fields varies. Second, the tags-as-memory abstraction as implemented in Gen 2 readers lacks the notion of a shared device-address space, so it cannot express useful mainstream networking concepts such as multicast, anycast, or peer-to-peer (tag-to-tag) communication. Finally, it offers no facility for rich tag messages. Tags respond to most commands with short fixed-length status messages and to *read* commands with the requested data. Tags cannot address messages to other tags.

In contrast to Gen 2, BAT comprises a set of abstractions designed to enable *anything-to-tag* communication, in which multiple tags can communicate with one another and with external entities such as Internet resources. Figure 1 depicts BAT at a block level.

BAT's increased flexibility and efficiency in comparison to Gen 2 result from several key distinctions. In terms of flexibility, tags are *addressable network nodes* rather than simple memory stores. Each tag and each relay has an address in a shared address space. In addition, tags and relays formulate their communications as explicitly addressed messages (packets) rather than implicitly addressed responses. Finally, BAT enables confidential tag-to-tag communication via untrusted *relays*.

In terms of efficiency, BAT conceptually replaces supply-chain RFID readers with BAT *relays* that maintain message queues. As in an IP network, interactions




Fig. 1 Overview: Relays collect packetized messages from tags and forward them through other relays, which deliver them. Packets are split into frames locally to maximize throughput

between relays and tags consist mainly of packet-exchange commands. Also, tags and relays negotiate message sizes to adapt to lower- or higher-quality links. Messages from tags and relays have the same message format; a relay can forward messages to other tags without altering them.

BAT’s abstractions are more closely matched to conventional networking abstractions than those of Gen 2, and they consequently enable a variety of functions. For example, multiple tags that need to collaborate toward a larger goal, such as collective time synchronization, can exchange messages with relays and tags to reach consensus using Paxos [16]. As another example, tag-to-tag messaging enables tags to solve optimization problems collectively with *ant algorithms*, biologically inspired algorithms that harness the power of collective behavior in solving complex problems using simple agents [17]. Ants are known to collectively compute the shortest path between their nest and a source of food by each depositing a certain amount of pheromone on their paths while walking, attracting other ants to follow the same path. Similar mechanisms can be generalized to allow swarm agents to solve problems collectively, such as the traveling salesman problem [7]. The coordinated action of multiple independent agents has also been studied as *amorphous computing* [1], inspired by metaphors in biology and physics.

3 BAT Design Overview

BAT’s basic model consists of relays that move into the vicinity of tags and offer to provide power and communication. Messages follow a path conceptually similar to mail in the US postal system: items to be delivered are picked up at their sources (BAT tags) by mail carriers (BAT relays), routed through the postal system (powered networks of BAT relays), and delivered by local carriers (BAT relays) to their destinations (BAT tags).

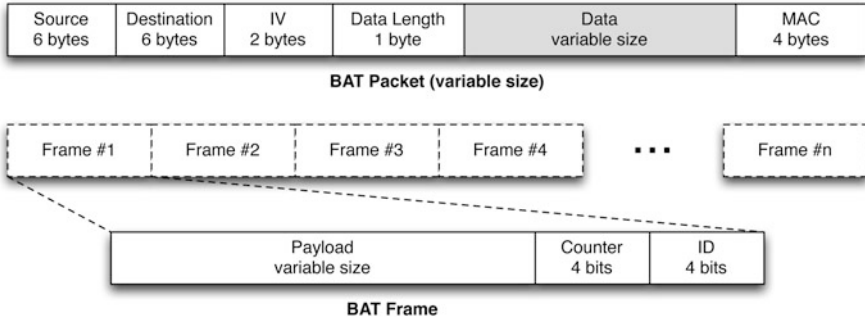


Fig. 2 The *framed* packet format accommodates up to $2^8 = 256$ bytes of data payload per packet by breaking packets into one or more frames of variable size

A BAT system comprises a set of CRFID tags, all of which are programmed initially by a trusted tag programmer, and a set of interchangeable relays that understand common frame and packet formats. A relay is a powered device akin to the RFID readers used in supply-chain applications. It powers tags by transmitting RF energy and can exchange messages with them via BAT's link layer. Multiple relays may be connected to a centralized application controller that knows the network's topology and can coordinate message routing; they may also independently query their surroundings for tags for which they have messages. The organization and operation of relays is application-dependent. Figure 3 summarizes a BAT interaction between a relay and a tag.

Our experimental results demonstrate the benefits of BAT's packet-framing and frame-size-negotiation mechanisms. As explained in Sect. 4, when there is a high-quality link between relays and tags, larger frame sizes may result in higher throughput. However, when the link is lossy or noisy, frame sizes must be smaller in order to achieve sustained communication.

The packet format depends on the higher network layers that are used above BAT. In our prototype implementation, we chose to follow a format similar to that proposed by Karlof et al. for sensor networks [15]. Another option could be 6LoWPAN [18], for example. A packet is split into frames as illustrated in Fig. 2. A frame includes a group ID, a frame-counter, and the payload. All frames with the same group ID are concatenated sequentially. The frame-counter gives the position of the frame within the packet. Once all frames are concatenated, the initialization vector (IV) together with the source and destination fields provide a unique packet identifier. The source and destination fields consist of an *ID*, a *group*, and a specified *domain* in order to facilitate routing: when a tag *A* has a message for a tag *B*, *A* sends the message addressed to *B*'s ID at the specified domain. Then, *B* retrieves messages addressed to it opportunistically. When a tag *A* wants to send a message to a particular group of tags in the local domain, it can simply *post* a message to

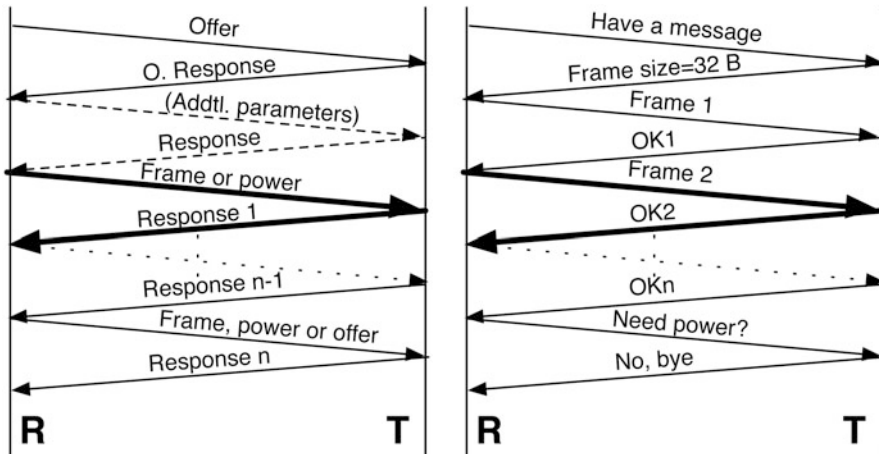


Fig. 3 BAT messages. A tag may request a relay to *deliver a message*, *provide power*, or *check for messages addressed to the tag*. In some cases, singulation may be necessary, but not always

the given group at the local domain. BAT implements the retrieval and collection of messages in a given domain, and the network layer routes and delivers messages between domains.

In order to facilitate the use of untrusted readers, BAT encrypts message payloads. A per-packet message authentication code (MAC) allows the recipient to detect tampering or transmission errors upon receipt of a packet. Section 5 provides the details of the security mechanisms BAT supports.

3.1 Prototype Implementation

Our relay implementation uses the GNU Radio software-defined radio platform to drive universal software radio peripheral (USRP) hardware [10, 13]. This hardware and software combination allows complete specification of message structure in both directions. We build upon Buettner’s RFID reader implementation [4] for our own relay software. We implement tags using the UMass Moo [26], see Fig. 5. This prototype has an MSP430F2618 with 8 KB of RAM and a maximum operating frequency of 4 MHz under harvested power. Our prototype fixes some transmission options to specific values for simplicity of implementation; in particular, it uses phase-reversal amplitude shift keying (PR-ASK) modulation with pulse interval encoding (PIE) from relay to tag, and phase-shift keying (PSK) modulation with Miller-4 encoding from tag to relay. Tags are implemented so that facilitation of power requests and rate adaptation are transparent to applications; individual applications on tags do not need to be aware of physical conditions or power budgets.

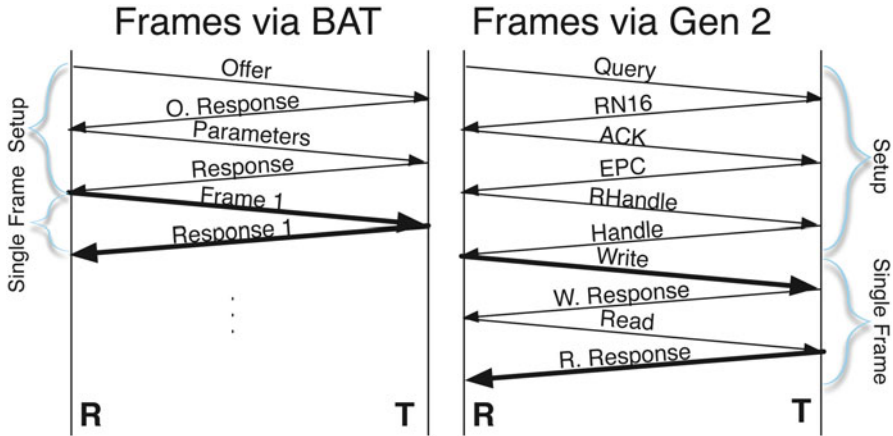


Fig. 4 Gen 2 requires messages to achieve singulation before Read/Write commands are issued. In order to implement custom round-trip messages as in BAT ($R \rightarrow T$, $T \rightarrow R$), a write command would have to be followed by a read command

4 BAT Evaluation

BAT's networking abstractions do not add significant communication overhead in comparison to supply-chain RFID networking. BAT's mechanisms to enable the negotiation of optimal frame size generally result in an increased bidirectional throughput.

BAT negotiates optimal frame size at the beginning of an interaction to maximize throughput. When a tag needs to receive more than a few bytes of data on a regular basis, using a fixed, conservative frame size adds significant overhead. The optimal frame size may vary considerably depending on operating conditions, the tag's internal memory speed, or the particular combination of reader and tag. In our experiments the throughput increased linearly with the frame size to achieve the highest throughput of 18 Kbps with a frame size of 112 bytes. With larger frame sizes, the error rate increases and therefore the goodput decreases.

Because there is not yet a wide variety of CRFID prototypes, we use high-capacity Gen 2 tags and readers to further illustrate the need for variable frame sizes. In this experiment, the payload of the Gen 2 BlockWrite command carries data from a reader to a tag. We vary the size of the payload and record the throughput for two commercial tags: the Xerafy Sky-ID and the Ramtron MaxArias, which can store up to 8 KB and 2 KB of user data, respectively. The readers in the experiment are the ThingMagic reader M5e and a Ramtron MaxReader development kit.

When using the *User* memory bank of the Xerafy Sky-ID, the number of BlockWrite commands necessary to fully write the payload value increases linearly with frame size, except when the tag is placed less than 3 mm from the reader's

Fig. 5 Hardware used to implement BAT. The relay (*foreground*) is a USRP with RFX900 daughterboards driven by GNU Radio. The UMass Moo (*background*) is a CRFID tag derived from the DL WISP 4.1 [23]



antenna. In order to achieve a high successful-write rate, a BlockWrite command should carry at most 4 bytes of data. MaxArias tags would achieve the best throughput when using BlockWrite commands with 56 bytes of data at a time. In order to implement BAT-style communication atop Gen 2, a BlockWrite command would have to be followed by a Read command (with ~ 4 bytes of data) in order to allow a tag to reply with a non-Gen 2 protocol response as illustrated in Fig. 4.

BAT's abstractions do not add a communication overhead. Its frame-size negotiation happens only once per transmission and its overhead is approximately 50ms, roughly the time it takes to perform a Gen 2 read. Gen 2 does not negotiate optimal frame size, but there is still some overhead associated with transmission, as illustrated in Fig. 4. For example, Gen 2 requires tag singulation because it is meant to be used in tag-dense environments. However, BAT may or may not need this singulation. Gen 2's overhead typically involves three initiating messages before each command is sent to a tag. In addition, Gen 2 Read and BlockWrite commands have fields that are unnecessary from the standpoint of a protocol designed to facilitate anything-to-tag communication.

We expect that the throughput of BAT will scale with improvements in prototypes and the physical layer. For example, one of the parameters of the physical layer that we use, the Tari (the time interval to encode a bit in the $R \rightarrow T$ direction), significantly impacts throughput. In our experiments, we found no errors when

the Tari was greater or equal to $13\mu\text{s}$. When Tari was equal to $11\mu\text{s}$, only 2 of 100 128-bit frames contained no errors. In this case, errors consisted of missed bits and not bit flips. When Tari was less than $11\mu\text{s}$, all frames had missed bits and flipped bits. The limiting factor hindering communication with Tari values smaller than $13\mu\text{s}$ is the MCU in the UMass Moo. The MSP430F2618 in the UMass Moo can process an interrupt no more than every $7\mu\text{s}$ regardless of the operating clock speed, and the time required to process the input is around $5\mu\text{s}$. The demodulator, however, supports higher rates. In contrast, other non-CRFID tags such as the MaxArias use a $6\mu\text{s}$ Tari.

5 Using Untrusted Relays

CRFIDs are well suited for network applications that require long-term deployments because of their low maintenance requirements. As discussed in Sect. 2, in many cases it may be important to provide message secrecy and integrity for such deployments. The ability to use untrusted relays for message routing minimizes the size of the trusted computing base and simplifies key management problems. Untrusted relays are only required to understand how to route packets in the network.

In order to support the use of untrusted relays, BAT encrypts packet payloads using AES, a well-studied cipher. To ensure packet integrity, each packet includes a 32-bit MAC to prevent tampering. Karlof et al. [15] provide arguments for the sufficiency of this MAC size for sensor network applications. Our preliminary results show that CRFID prototypes are capable of performing state-of-the-art cryptographic operations to support encrypted communication at the link layer. The most expensive operations correspond to the public-key cryptography necessary when shared keys are not distributed in advance. Once these keys have been established, providing secrecy and integrity is relatively inexpensive, as shown in Table 1. Opportunistic encryption allows for 61 Kbps of throughput, higher than the maximum 18 Kbps link speed achieved in our experiments. Once a shared key has been constructed, tags must encrypt the packet payload using a symmetric block cipher (AES-128) and then calculate a MAC (we use CMAC [8]) over an IV and ciphertext. This sort of computation is also possible because BAT allows tags to request additional power to encrypt/decrypt as needed.

5.1 Performance on Future CRFID Prototypes

The ratio of the computational performance of cryptographic operations over the overall throughput of BAT's communication will continue to decrease with improved microcontroller capabilities, such as power-efficiency and higher clock speeds. Our implementation was based on a current CRFID prototype with an

Table 1 The cycle count for each security operation was measured using a hardware debugger and a UMass Moo

Operation	Cycles	Time (ms) (calculated)
AES-128 setup	144,677	36.1
AES-128 enc (per block)	7,795	1.9
AES-128 dec (per block)	8,003	2.0
CMAC (AES-128 80 B)	39,669	9.9

The reported time is calculated assuming a typical sending clock speed of 4 MHz

Table 2 Computation times for cryptographic operations on the following MCUs: (1) MSP430F5310 @ 8 MHz; (2) ARM Cortex-M0+ @ 30 MHz; (3) ARM Cortex-M0 @ 50 MHz

Primitive	MCU	Clock freq. (MHz)	Time (calculated)
AES setup	MSP43F5310	8	18 ms
AES setup	Cortex-M0+	30	1.2 ms
AES setup	Cortex-M0	50	0.7 ms
AES enc	MSP43F5310	8	978 μ s
AES enc	Cortex-M0+	30	74 μ s
AES enc	Cortex-M0	50	44 μ s
AES dec	MSP43F5310	8	1,004 μ s
AES dec	Cortex-M0+	30	80 μ s
AES dec	Cortex-M0	50	48 μ s

MSP430F2618 microcontroller, the UMass Moo [26]. In the near future we may see other CRFID prototypes with ARM microcontrollers or other MSP430 microcontrollers with higher clock speeds, which will enhance the performance of BAT’s cryptographic operations. There are MCUs in the MSP430F5xx series that would allow a CRFID to operate at 8 MHz using harvested energy, or at a maximum of 25 MHz using an alternative source of energy. There are also several MCUs in the ARM Cortex-M0 and ARM Cortex-M0+ families that are even more power-efficient than MCUs in the MSP430 family, while providing higher clock speeds. For example, the ARM Cortex-M0+ may consume as little as 11.2 μ W/MHz compared to 803 μ W/MHz consumed by the MSP430F2618. An ARM Cortex-M0 may consume 16 μ W/MHz in its lowest-power setting. Table 2 forecasts BAT’s performance on these MCUs.

5.2 Shared-Key Generation

Shared-key generation is the most expensive security operation required by BAT. However, our micro-benchmarks show that identity-based key agreements are feasible on CRFID tags. The Smart–Chen–Kudla (SCK) key agreement [6] is less computationally expensive than a traditional Diffie–Hellman (DH) approach. However, the non-interactive Sakai, Ohgishi and Kasahara (SOK) construction [21] is only marginally more expensive than DH.

Table 3 Computation times for key agreements on the MSP430F2618 @ 4MHz

Key agreement	Security	Time (s) (calculated)
Diffie–Hellman	2,048-bit	208
Diffie–Hellman	1,024-bit	50
Diffie–Hellman	640-bit	21
ECC Diffie–Hellman	~DL 640-bit	27
SOK with type-1	~DL 1,024-bit	53
SOK with type-1	~DL 640-bit	30

All multi-precision arithmetic, ECC arithmetic, and pairings are implemented by Miracl [5]

The DH key exchange based on the ECDLP uses the NIST curve P-192

The Type-1 pairings use supersingular curves $E(\mathbb{F}_p)$ for a 512-bit prime and $E(\mathbb{F}_{2^{379}})$

Identity-based key-agreement schemes allow for the creation of private–public key pairs, such that the public key is an arbitrary string and only a trusted entity—a private-key generator (PKG)—can compute the corresponding private key. Thus, for example, it would be easy for a tag to encrypt a message so that only a tag with serial number x could obtain the corresponding decryption key from the PKG. Additionally, the sending tag can include an expiration time in its public key. SOK key exchange is also desirable because DH is vulnerable to man-in-the-middle attacks unless a third party authenticates protocol participants. This limitation is usually addressed by the addition of a certificate authority (CA), but key management is difficult in CA-based systems [11]. Table 3 lists the times necessary to compute a shared key using the approaches described above and the elliptic curve pairing implementations by the Miracl Crypto SDK [5].

6 Related Work

Networks of RFID tags [2] or CRFIDs have been of interest in the last few years, and other authors have proposed applications including building instrumentation [25] and user-activity inference [3]. This past work assumes that tags are not capable of communicating with one another and that the network is actually composed of readers that report data to a central database where all computations occur. Other work has focused on enhancing supply-chain RFID-style backscattering. For example, Wang et al. [24] describe how to reduce singulation overhead by allowing multiple tags to communicate simultaneously in spite of collisions.

Reynolds has suggested the use of *semi-passive* tags capable of intercommunication, but these tags abandon backscatter in favor of traditional radio technology, so the network model is much different from our own [20]. Application-specific communication between backscattering RFID tags has been explored previously by Juels [14]. Nikitin et al. [19] propose tag-to-tag communication by allowing

passive tags to listen to one another's transmissions. Their approach is restricted by requiring close proximity. BAT could be implemented atop this physical layer to improve throughput. CCCP [22] supports tag-to-tag communication of a sort, but it is limited to storage and retrieval of data from a single tag to an untrusted reader. It does not give tags the ability to exchange data.

7 Conclusion

BAT allows CRFIDs to communicate with computer systems as well as other tags using relays that can be easily replaced and do not need to be aware of application information. Our preliminary implementation demonstrates that BAT's abstractions are better suited for RFID-scale sensor networks than communication protocols designed for supply-chain tags. Our micro-benchmarks show that offering a secure network layer that uses untrusted relays to exchange information does not add a significant overhead, which is attractive for applications that require tags to be left exposed and unattended.

Acknowledgements We thank Deepak Ganesan for his feedback on this work. This research was supported by the National Science Foundation CNS-0845874, CNS-0831244, and CNS-0923313.

References

1. H. Abelson, D. Allen, D. Coore, C. Hanson, G. Homsy, T. F. Knight, Jr., R. Nagpal, E. Rauch, G. J. Sussman, and R. Weiss. Amorphous Computing. *Communications of the ACM*, 43(5), May 2000.
2. M. Buettner, B. Greenstein, and A. Sample. Revisiting Smart Dust with RFID Sensor Networks. *Hot Topics in Networks*, Jan. 2008.
3. M. Buettner, R. Prasad, M. Philipose, and D. Wetherall. Recognizing Daily Activities with RFID-Based Sensors. In *Proceedings of the 11th International Conference on Ubiquitous Computing (UbiComp)*, Oct. 2009.
4. M. Buettner and D. Wetherall. A Software Radio-based UHF RFID Reader for PHY/MAC Experimentation. In *IEEE International Conference on RFID (RFID)*, 2011.
5. CertiVox. MIRACL Crypto SDK. <http://certivox.com/index.php/solutions/miracl-crypto-sdk/>.
6. L. Chen and C. Kudla. Identity Based Authenticated Key Agreement Protocols from Pairings. In *Proceedings of the 16th IEEE Computer Security Foundations Workshop*, July 2003.
7. M. Dorigo, G. Caro, and L. Gambardella. Ant Algorithms for Discrete Optimization. *Artificial Life*, 5(2), pp. 137–172, 1999.
8. M. Dworkin. I. T. L. N. I. of Standards, and T. C. S. Division. *Recommendation for Block Cipher Modes of Operation: The CMAC Mode for Authentication*. US Department of Commerce, National Institute of Standards and Technology, 2005.
9. EPCglobal Inc. EPCglobal Class 1 Generation 2 Air Interface. V. 1.2.0, Oct. 2008.
10. Ettus Research. Universal Software Radio Peripheral. <http://ettus.com/>.
11. N. Ferguson, B. Schneier, and T. Kohno. *Cryptography Engineering: Design Principles and Practical Applications*. Wiley, New York, 2010.

12. J. Gummesson, P. Zhang, and D. Ganesan. Flit: A Bulk Transmission Protocol for RFID-Scale Sensors. In *Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services*. ACM, 2012.
13. GNU Radio. <http://gnuradio.org/>.
14. A. Juels. “Yoking-Proofs” for RFID Tags. In R. Sandhu and R. Thomas, editors, *First International Workshop on Pervasive Computing and Communication Security*. IEEE Press, Mar. 2004.
15. C. Karlof, N. Sastry, and D. Wagner. TinySec: A Link Layer Security Architecture for Wireless Sensor Networks. In *Proceedings of ACM International Conference on Embedded Networked Sensor Systems*, 2004.
16. L. Lamport. The part-time parliament. *ACM Transactions on Computer Systems*, 16(2), May 1998.
17. R. Mullen, D. Monekosso, S. Barman, and P. Remagnino. A review of ant algorithms. *Expert Systems with Applications*, 36, pp. 9608–9617, 2009.
18. G. Mulligan. The 6lowpan architecture. In *Proceedings of the 4th workshop on Embedded networked sensors*, EmNets '07, pages 78–82, New York, NY, USA, 2007. ACM.
19. P. Nikitin, S. Ramamurthy, R. Martinez, and K. Rao. Passive Tag-to-Tag Communication. In *IEEE International Conference on RFID (RFID)*, April 2012.
20. M. Reynolds. Beyond RFID: Peer to Peer, Semi-Active RFID Tags. *NSF Workshop on Animal Tracking and Physiological Monitoring*, 2007. Presentation.
21. R. Sakai, K. Ohgishi, and M. Kasahara. Cryptosystems based on pairing. In *Proceedings of the Symposium on Cryptography and Information Security*, Jan. 2000.
22. M. Salajegheh, S. S. Clark, B. Ransford, K. Fu, and A. Juels. CCCP: secure remote storage for computational RFIDs. In *Proceedings of the 18th USENIX Security Symposium*, Montreal, Canada, Aug. 2009.
23. A. P. Sample, D. J. Yeager, P. S. Powlledge, A. V. Marnishev, and J. R. Smith. Design of an RFID-based battery-free programmable sensing platform. *IEEE Transactions on Instrumentation and Measurement*, Nov. 2008.
24. J. Wang, H. Hassanieh, D. Katabi, and P. Indyk. Efficient and Reliable Low-Power Backscatter Networks. In *Proceedings of the ACM SIGCOMM 2012 Conference*, 2012.
25. E. Welbourne, K. Koscher, E. Soroush, M. Balazinska, and G. Borriello. Longitudinal Study of a Building-Scale RFID Ecosystem. In *Proceedings of the International Conference on Mobile Systems, Applications, and Services*, MobiSys '09, June 2009.
26. H. Zhang, J. Gummesson, B. Ransford, and K. Fu. Moo: A Batteryless Computational RFID and Sensing Platform. Technical Report UM-CS-2011-020, Department of Computer Science, University of Massachusetts Amherst, Amherst, MA, June 2011.

Implementing the Gen 2 MAC on the Intel-UW WISP

Michael Buettner and David Wetherall

1 Introduction

The Intel-UW wireless identification and sensing platform (WISP) [10] is the first example of a new class of computational device. WISPs are small, low-power computers complete with a microprocessor, sensors, memory, and communication that are powered completely by energy harvested from the RF environment. The technology to develop these devices builds on that of passive UHF RFID tags. Passive RFID tags gather the entirety of their operating energy from RF signals transmitted by nearby RFID readers; they have no batteries or long-term energy stores. This means they can be very small, long-lived, and deeply embedded into the physical environment.

WISPs inherit these advantages and add new capabilities such as sensing and computation. This has the potential to enable a broad range of applications including cold-chain monitoring, access control, embedded monitoring of bridges and planes, gestural interfaces, activity recognition, and non-intrusive physiological monitoring [1]. That is, WISPs are of interest as an enabler for cyber-physical systems, ubiquitous computing applications, and sensor networks that approaches the vision of “smart dust” [12].

Applications built using WISPs generally require that data, such as sensor readings or the results of computation, be transmitted to a host computer for processing. By masquerading as standard RFID tags, WISPs can leverage existing RFID infrastructure to do this. Gen 2 RFID tags [5], the standard used in modern systems, transmit simple 96 bit identifiers. WISPs transmit data to Gen 2 readers by

M. Buettner (✉) • D. Wetherall
University of Washington, Seattle, WA, USA
e-mail: michael.buettner@gmail.com; djw@cs.washington.edu

using shorter identifiers and embedding their data in the remainder of the 96 bits. The reader simply reports identifiers to the host application which is responsible for extracting the data payload.

In reality, WISPs operate very differently from traditional RFID tags. Most significant is the amount of power they consume compared to the amount of power they can harvest. When standard RFID tags are receiving sufficient power to turn on, they can communicate with the RFID reader indefinitely unless the reader stops supplying power, i.e., they harvest more power than they consume during communication. WISPs, on the other hand, consume far more power when communicating than they can harvest because their added capabilities come at the cost of increased power consumption. This imposes an operational model where WISPs spend much of their time sleeping in a low-power mode, where energy is buffered in a capacitor, and only wake up to execute tasks and communicate when sufficient energy has been stored. This new operational model means that MAC protocols designed for traditional RFID tags are a poor fit for emerging devices because they may run out of energy during a message exchange.

Determining how much energy must be stored to transmit data to the RFID reader is difficult. This is because the number of messages that must be processed by the tag is nondeterministic due to the anticollisions mechanism defined by the Gen 2 standard. To work around this, the existing WISP MAC implementation does not strictly adhere to the protocol. Instead of “waiting its turn” so as to not collide with other tag transmissions, WISPs transmit at every opportunity. This approach is not only deterministic but also consumes a minimal amount of energy for each data transmission. In practice, it has worked well for most studies seen in the literature as they have generally experimented with only a single WISP [3,4,6,9,13]. However, as full-scale deployments begin to emerge [2] this simple approach to communication can become problematic. When many WISPs are present, their simple MAC behavior means that they often transmit at the same time and their transmissions collide.

In this chapter, we experiment with the Gen 2 anticollision mechanism to determine if it is appropriate for use on WISPs. To avoid collisions, Gen 2 tags randomly choose one of a reader specified number of slots and reply only in that slot. However, this means that a tag may need to process many reader commands before responding, and common wisdom has suggested that the full Gen 2 MAC is thus too energy intensive for use on the WISPs. We find this not to be the case, as the cost of Gen 2 slotting is high only when collisions are likely and the cost is worthwhile.

We present the implementation of the Gen 2 slotting mechanism on the Intel-UW WISP and evaluate its effectiveness. We were able to implement this functionality while still meeting the tight timing constraints of the protocol by carefully choosing when to execute time-intensive operations. When using the previous MAC approach, even one well-powered tag significantly degrades system performance, and with only five tags present the system becomes unusable. With Gen 2 slotting enabled, all tags in the deployment are able to communicate with the reader at generally better than ten responses per second. Lastly, we find that the additional

energy cost of Gen 2 slotting is close to that of the previous, non-slotting, MAC when collisions are low. Thus, Gen 2 slotting is desirable for a wide range of deployments.

The remainder of this chapter is organized as follows: in Sect. 2 we describe the salient aspects of the Gen 2 protocol, Sect. 3 gives an overview of the WISP with a focus on the communication subsystem, Sect. 4 describes our implementation of Gen 2 slotting on the WISP, Sect. 5 presents the evaluation of our system, and we conclude in Sect. 6.

2 Gen 2 RFID Background

In this section, we present relevant background on the Gen 2 PHY and MAC layers. The Class-1 Generation 2 (Gen 2) RFID specification was formalized by EPCglobal in 2004 and is the standard for UHF RFID. The WISP was designed to take advantage of widely deployed RFID infrastructure by harvesting power from and communicating with commercial Gen 2 readers. Two constraints drove the design of the specification. First, it had to be implementable on very low cost RFID tags (each costing a few cents) that are wirelessly powered and computationally weak. Second, it had to operate in a regime where readers can communicate with tags and vice versa, but tags cannot communicate with each other or even hear other tag transmissions.

2.1 Gen 2 Physical Layer

The Gen 2 physical layer has the dual purpose of maximizing harvestable power at the tags and facilitating downlink and uplink communication. While a reader is communicating with tags it must transmit a continuous RF wave (CW); tags power themselves by harvesting the energy in this RF signal. Downlink communication uses on/off keying where bit boundaries are indicated by brief pulses of zero amplitude in the CW, and pulse interval encoding (PIE) where the time between zero amplitude pulses differentiates a zero or a one. Depending on PIE durations, downlink rates range from 27 to 128 kbps.

Uplink data rates are between 5 and 640 kbps, and communication is achieved via “backscatter” transmission. Passive RFID tags do not technically transmit any energy. Instead, they manipulate how well they reflect (backscatter) the incident CW. This reflected signal, though weak, is received by the reader which decodes the modulated data. Uplink modulation is determined by two parameters specified by the reader; uplink frequency and data encoding. Gen 2 specifies four encodings schemes that differ in the number of cycles per symbol (varying from one cycle to eight). For a given link frequency, an encoding using more cycles will have a lower data rate, but will be more robust to noise.

Fig. 1 Reader message and tag response

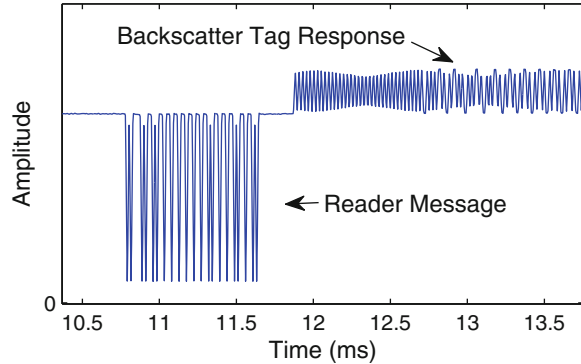


Figure 1 shows communication between a reader and tag as captured by a software radio. The CW results in the DC offset of the received waveform, with the series of low amplitude pulses being a reader transmission. The backscattered tag response can be clearly seen as a combination of the incident CW and the reflected CW from the tag. It should be noted that the signal seen in the figure was captured with the software radio’s antenna placed inches from the RFID tag. This accounts for the prevalent backscatter signal.


To limit interference with other devices in the 902–928 MHz ISM band, FCC regulations stipulate that UHF RFID systems frequency hop across fifty 500 kHz channels, with dwell times no longer than 400 ms. However, because tags do not “tune” to different channels, when multiple readers are active in an area their transmissions will collide at tags even if the readers are on different channels. Similarly, if multiple tags respond to a reader command their transmissions will collide.

2.2 Gen 2 MAC Layer

Gen 2 tags decode reader transmissions using a simple edge detector that detects the conspicuous zero amplitude pulses in the CW. However, tags are unable to decode, or even detect, the backscattered signals of other tags. This is in contrast to Ethernet or 802.11 where nodes can hear each other, at least when they are not transmitting. Consequently, the Gen 2 MAC protocol is based on Framed Slotted Aloha [8] which was designed to operate in a context where transmitting nodes cannot hear each other.

The general model of RFID is that readers are continuously “inventorying”, i.e., looking for tags that are in range. An inventory round begins with the reader transmitting a *Query* command that indicates the number of slots in the frame. The number of slots in a frame is a power of two in the range [1, 32768]. After

Fig. 2 Message exchange for a Query round




receiving a *Query*, tags randomly choose a slot in which to reply and transmit a 16-bit random number (*RN16*) in that slot. If the reader receives an *RN16* in a slot, it *ACKs* the *RN16* and the tag that transmitted it replies with its 96-bit identifier. Tags that collide in a slot are not *ACKed* and respond again after the next *Query*. This protocol provides collision avoidance, as tags randomly choose their slots, and collision resolution as collided tags respond again in the next round.

Figure 2 shows a message exchange with only one tag present. In the example, the reader powers up and transmits a *Query* message that specifies four slots in the frame. The first slot immediately follows the *Query* command and is empty in this example. The second reader command is a *QueryRepeat* which indicates the beginning of a new slot. In this case, the tag had chosen the second slot and so it transmits its *RN16* in response to the *QueryRepeat*. It is then *ACKed* by the reader, and transmits its *ID*. Because the reader specified four slots in the *Query*, it sends two additional *QueryRepeats* looking for any remaining tags. In this case, there is only a single tag so the last two slots are empty. The reader then powers down.

3 Wireless Identification and Sensing Platform

In this section, we give relevant background on WISP operation and communication subsystem, with a focus on the current MAC implementation. The Intel-UW WISP [10], shown in Fig. 3, is a battery-free, wirelessly powered platform for sensing and computation that combines many of the strengths of passive RFID and wireless sensor motes. WISPs are powered by and communicate with EPC Gen 2 RFID readers at a range of nearly 4m. Computation is provided by a fully programmable ultra-low-power 16-bit MSP430 microcontroller with an analog to digital converter. This WISP includes 8K of flash program space, a 3D accelerometer, temperature sensor, and 8K serial flash. Small header pins expose microcontroller ports for expansion daughter boards, external sensors, and peripherals. Application software is written in C, as is the transmit and receive functionality.

Fig. 3 Gen 2 tag, Intel-UW WISP, Telos Mote



3.1 Basic Operation

The WISP is designed to function at 1.8 V, and when the CPU is in active mode it has a clock speed around 3.5 MHz and draws approximately 600 μ A. This is far more power than can usually be harvested from the RFID reader. In contrast, when the WISP is in a low-power state it draws close to 1 μ A and harvested energy is banked in a 10 μ F capacitor. This imposes an operating model in which the WISP spends much of its time sleeping and charging the capacitor and then waking up briefly to execute tasks. Though designed to operate at 1.8 V, as long as the WISP is supplied with at least 1.5 V, it can operate and the contents of RAM can be maintained. If the supplied voltage drops below this threshold the WISP will “black out” and lose all state.

WISP circuitry harvests energy whenever there is a strong enough reader signal. When the WISP charges to 1.5 V it powers up and immediately goes to sleep to store more energy; if it started execution right away it would drop below 1.5 V and black out. To wake up when there is sufficient energy to do useful work, WISPs are equipped with a voltage supervisor that generates an interrupt when the voltage of the capacitor rises above a fixed threshold of 2 V. This wakes the WISP from sleep mode and it begins to execute its program, possibly communicating with the reader. After the WISP completes a task, it puts itself to sleep and charges back to 2 V.

3.2 Physical Layer Implementation

The WISP can decode reader messages and transmit messages to the reader. Both receive and transmit functionality are implemented in software and consist of tuned assembly code to decode and modulate signals at hundreds of Hz. The WISPs hardware demodulator generates interrupts at the CPU for all falling edges in the

reader signal. By measuring the time between these interrupts the WISP decodes reader commands. To transmit messages, the WISP carefully modulates a transistor that shorts the antenna thus generating a backscattered signal. This entails that the WISP must be awake while receiving messages from the reader or transmitting responses, which consumes energy stored in the capacitor. However, the WISP sleeps during the periods between reader messages; the first edge of a reader command triggers an interrupt that wakes the WISP.

3.3 *MAC Layer Implementation*

Because the WISP must be awake to receive and transmit messages, the WISP may run out of energy before transmitting its data if it has to process too many reader commands. Consequently, existing WISP firmware does not implement the full Gen 2 MAC protocol because processing many *QueryRepeats* is too energy intensive. Instead of randomly choosing a slot, and responding only once per frame, the WISP transmits data whenever it has sufficient energy. More specifically, when the WISP has data to send to the reader it transmits an *RN16* after every *Query* or *QueryRepeat* that it decodes. This approach removes the collision avoidance aspect of the Gen 2 protocol in the interest of lower energy consumption. It works on the assumption that, in the common case, only one WISP will be powered (and have data to send) during a given slot.

3.4 *Drawbacks to WISP MAC Implementation*

The benefit of the WISP MAC implementation is that the 2 V wake-up threshold is guaranteed to be sufficient to complete the message exchange; in fact, the 2 V threshold was chosen specifically for this reason. As long as no other tag responds in the slot, the WISP will not run out of energy before transmitting its *Identifier*. In contrast, if the Gen 2 slotting MAC were used the WISP may choose a slot late in the frame and, because processing the *QueryRepeats* consumes energy, it could run out of energy before reaching its slot. To accommodate a large number of slots, the WISP would need to store more energy either by using a larger capacitor or waking up at some threshold above 2 V. This would result in longer charge times and lower response rates.

The above approach works well as long as tag collisions are rare. However, this is not the case when there are many tags in the deployment, or if even a single tag is very well powered. In that case, tags will consistently collide, possibly resulting in congestion collapse where no tags can communicate with the reader and the deployment becomes unusable. Hence, there is a trade-off between consuming more energy per response and wasting energy by colliding with other tags.

4 Implementing the Gen 2 MAC Protocol

To determine if the Gen 2 anticollision protocol is appropriate for the WISP, we integrated slotting into existing WISP firmware. The firmware is written in a mix of C and assembly for timing sensitive operations. The communication subsystem is particularly timing sensitive due to the strict timing of the Gen 2 protocol. Once a WISP begins timing sensitive with the reader, it must respond to reader commands and be ready to process new commands quickly, generally on the order of tens of microseconds. If the deadlines imposed by the protocol are violated, the reader will ignore tag transmissions. Meeting these timing requirements was the main challenge for implementing Gen 2 slotting on the WISP. We describe our implementation in this section.

4.1 *Random Number Generation*

The Gen 2 MAC protocol requires that tags choose slots randomly. As a source of randomness, the WISP can use the mismatch between the internal voltage oscillator and the crystal oscillator or the noise inherent in ADC readings [11]. Both mechanisms are too time-consuming, require too much energy, and provide more randomness than we need.

Instead, we sample the voltage in the capacitor once immediately when the WISP first powers up and use this value as a seed for a pseudorandom number generator. The variance in this voltage sample, due to input power and noise in the ADC, gives us sufficient randomness. Alternatively, we could have used SRAM state as a random source, with similar efficiency [7].

4.2 *Choosing a Slot*

After receiving the *Query* command the WISP should generate a new random number and take its modulus to determine what slot to respond in. However, we found that incrementing the random number generator after decoding the command took too long and caused the WISP to send the *RN16* too late if it chose the first slot. To work around this, we generate a random number during each slot that the WISP does *not* respond in. This assures that there is no tight deadline for the WISP to meet (the WISP just needs to be ready to decode the next reader message), and at each *Query* a new random number is used.

When the *Query* is received, the WISP must scale the 16-bit random number to the frame size. We found that the modulus operation was too time intensive and caused the WISP to again miss the deadline for its *RN16*. Because frame sizes are

powers of two, we instead scale the random number via bit shifting to the right. The number of positions to shift is based on the frame size. This approach is sufficiently fast to meet the timing deadline of the protocol.

4.3 *Modifying the State Machine*

Modifications to the state machine were straight forward. When a *Query* is received, the frame size is extracted, the slot is randomly chosen, and the slot counter is set to zero. If the first slot is chosen the WISP responds immediately to the reader. Otherwise, at every *QueryRepeat* the slot counter is decremented and the *RN16* is transmitted when the counter reaches zero. This is in contrast to WISP communication without slotting, where the WISP transmits an *RN16* at every *Query* or *QueryRepeat*.


5 Evaluation

In this section, we compare the performance of a group of WISPs when they implement Gen 2 slotting and when they use the approach of responding to the reader whenever possible. Though slotting may require more energy per response, as more messages must be processed by the tag, we find that slotting significantly improves WISP performance when more than one tag is present and does not appreciably degrade performance in the case of a single tag.

5.1 *Experimental Setup*

The data presented in this section was gathered using an Impinj Speedway reader and up to 5 Intel WISPs. We evaluate performance in terms of response rate for WISPs that use slotting and WISPs that do not. Response rate is the number of time per second that a WISPs identifier is received by the reader, and the results of three trials were averaged for each experiment. For the WISPs that do not use slotting, the tag attempts to communicate with the reader at every opportunity, transmitting an *RN16* whenever it decodes a *Query* or a *QueryRepeat* message. For the WISPs that use slotting, we use the implementation described in the previous section. For this study, there is no workload aside from communication; the WISPs do not take sensor readings or perform any computational tasks. The same 5 WISPs were used for all experiments and were reprogrammed between experiments. Care was taken to replace each WISP at its original position after reprogramming.

Fig. 4 Multiple tags without slotting



5.2 Tags Near the Reader

When WISPs are close to a reader, they have sufficient energy to respond very often. This results in high response rates for a single tag, but many collisions and poor performance when more than one tag is present. This can be a significant problem for deployments such as [2] where many tagged objects are on a table and a reader antenna is nearby.


To illustrate this problem, we experimented with WISPs that did not use slotting. We placed the WISPs on the poster board one meter from the reader antenna, set the reader transmit power to 30 dBm, and measured the average response rate over one minute. This experiment was repeated as WISPs were added one at a time until 5 WISPs were present, and then the WISPs were removed one at a time in the order that they were added.

Figure 4 shows the response rates for the WISPs, with the vertical dashed lines separating the different sets, e.g., the far left shows only WISP-1, the set in the middle of the figure shows all five WISPs, and the far right shows the case when only WISP-5 is left on the poster board.

When only WISP-1 is present it responds approximately 50 times per second. When a second tag is added, the response rate of WISP-1 decreases slightly, but the second WISP responds less than once per second. This is because of the capture effect, where a stronger signal is decoded while the weaker signal is ignored. Though both tags respond at every opportunity, the signal of WISP-1 is stronger than WISP-2, so WISP-1 is read every time. Similarly, when WISP-3 and WISP-4 are added, their response rates are close to zero while WISP-1 continues to perform quite well.

When WISP-5 is added, all WISPs perform very poorly; none respond even once per second. This is because the combined signals of all WISPs mean that none can be read well. As we remove tags, no tag responses are sufficiently strong to be decoded until only WISP-4 and WISP-5 are present; at this point WISP-5 is stronger and

Fig. 5 Multiple tags with slotting



it responds nearly 50 times per second. One thing to note is that if even *one* tag can respond to most reader messages the system can experience congestion collapse and no other tags will be read. For example, there may be many WISPs deployed in a room or across a workspace, but if a single tag is placed near the antenna the deployment may become unusable.

If tags use Gen 2 slotting, their responses are randomized across the slots in the frame and collisions are less frequent. We repeated the above experiment with Gen 2 slotting enabled, and the results are shown in Fig. 5. When WISP-1 and WISP-2 are present, the response rate of WISP-1 increases slightly and WISP-2 responds more than 20 times per second. This is in contrast to WISP-2 responding less than once per second when not using slotting. With all five WISPs present, the minimum response rate is above 6 per second, most are above 10 per second, and WISP-1 performs as well as it did when it was the only tag present.

Surprisingly, we found that the response rate of a given tag sometimes *increases* when tags are added to the deployment; this can be seen when WISPs 1–4 were present and WISP-1 achieves better than 90 responses per second. This is because the duration of an inventory round is greater when more tags are present, but the time that the reader powers down between rounds is fixed at 2 ms. Consequently, when many tags are present the reader is powered down a smaller fraction of the time, so there is more energy available to the WISP and it charges faster and responds more often.

These results show that our Gen 2 slotting implementation effectively avoids congestion collapse and assures that all WISPs get a chance to transmit their data. However, slotting does not mean that all tags will respond equally as often. This can be caused by different tags receiving different amounts of energy, so they are not all powered up and attempting to communicate at the same rate. Additionally, slotting does not entirely eliminate collisions so tags with stronger signals may still be decoded even if a collision occurs.

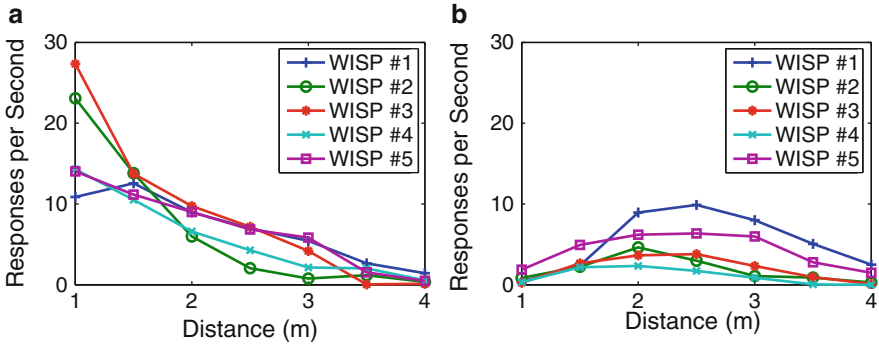


Fig. 6 Response rates at range for the two schemes (a) Slotting enabled (b) No slotting

5.3 Tags at Range

The Gen 2 slotting mechanism may require tags to process more reader messages before responding. For example, if a tag chooses the fifth slot for its *RN16* it must process at least the *Query* command and four *QueryRepeat* commands before responding. This consumes more energy than simply waking up and responding at the next opportunity. A higher energy cost per response means tags can respond less often, as they must sleep longer to recharge their energy store. Our previous experiment showed that the increase in energy consumption is worthwhile at close range. However, even with tags further from the reader and receiving less power, we do not want the increased energy consumption to degrade performance.

To evaluate WISPs at increasing distance, we placed the 5 WISPs on the poster board and measured their response rates while variably attenuating the transmit power of the reader. Seven settings were used to approximate ranges from 1 m to 4 m based on free-space propagation. Varying the transmit power of the reader allows us to avoid the changes in multi-path environment that would result from physically moving the tags between experiments. Figure 6a, b shows the results of this experiment when using Gen 2 slotting and when the WISP responds at every opportunity, respectively.

With slotting enabled, the response rates for the WISPs generally reduce monotonically with distance; beyond 4 m no tags responded. This is the expected behavior because the tags receive less power from the reader so recharging the capacitor takes longer. When slots are not used the case is quite different, as shown in Fig. 6b. Here, response rates are very low at close range, increase at middle distances, and then reduce as distance increases past 2.5 m. At close range, collisions degrade performance as discussed previously, and beyond 2.5 m response rate is reduced due to lack of power just as in Fig. 6a. At middle distances, the odds of two tags waking up in exactly the same slot and colliding is low. In effect, the natural

duty cycling that results from tags receiving different amounts of power sufficiently avoids collisions. When this is the case, tag response rate is largely independent of the total number of tags.

Initially, we expected slotting to degrade performance when only one tag was active in a slot, because tags would pay the cost of processing more messages even though the probability of collisions, and the expected benefit from avoiding them, was low. In practice, this was not the case; the median response rates were approximately equal for the two schemes beyond 2 m. The reason performance does not noticeably degrade is that commercial RFID readers adapt the frame size based on the number of tags that responded in the last frame. If many tags are being detected, larger frames are used to reduce the chance of collisions. If zero or one tags respond, the reader uses frames with only one slot. In our scenario, there is effectively only one tag active at a time beyond 2 m, as tags spend most of their time asleep harvesting energy. Because there is only one slot per frame, WISPs that implement slotting do not consume much more energy than WISPs that do not implement slotting. In effect, the cost of Gen 2 slotting is only significant when collisions are likely, i.e., precisely when the benefits outweigh the cost.

6 Conclusion

In this chapter, we explore using the Gen 2 anticollision protocol on the Intel-UW WISP. This is motivated by the fact that the previous WISP approach to medium access, where tags transmit at every opportunity, can result in congestion collapse when many WISPs are present. Though conventional wisdom has suggested that the full Gen 2 MAC would be too energy intensive for use on the WISP, we find this not to be the case, as the cost of Gen 2 slotting is high only when collisions are likely and the cost is worthwhile.

By carefully choosing when time-intensive operations are executed, we are able to implement the full Gen 2 MAC on the WISP hardware while still meeting the timing requirements of the protocols. We show through experimentation with five WISPs that slotting significantly improves the response rates of the deployed tags. When deployed close to the reader, WISPs that use slotting are generally able to respond more than 10 times per second, whereas without slotting congestion collapse occurs and none of the tags respond more than once per second. Additionally, we show that when tags are further from the reader and collisions are rare, the added energy cost of slotting is minimal and does not negatively impact performance. Consequently, WISPs should implement the full Gen 2 MAC protocol as it enables tags to communicate effectively in deployments consisting of one or more tags.

References

1. M. Buettner et al. Revisiting smart dust with RFID sensor networks. In *HotNets*, 2008.
2. M. Buettner et al. Recognizing daily activities with RFID-based sensors. In *Ubicomp*, 2009.
3. H. J. Chae et al. Maximalist cryptography and computation on the WISP UHF RFID tag. In *Proc. RFID Security*, 2007.
4. A. Czeskis et al. RFIDs and secret handshakes: Defending against ghost-and-leech attacks and unauthorized reads with context-aware communications. In *CCS*, 2008.
5. EPCglobal. EPC radio-frequency identity protocols class-1 generation-2 UHF RFID protocol for communications at 860 mhz–960 mhz version 1.0.9. 2005.
6. D. Halperin et al. Pacemakers and implantable cardiac defibrillators: Software radio attacks and zero-power defenses. In *IEEE Symposium on Security and Privacy*, 2008.
7. D. E. Holcomb et al. Initial sram state as a fingerprint and source of true random numbers for RFID tags. In *Proc. RFID Security*, 2007.
8. L. G. Roberts. Aloha packet system with and without slots and capture. *SIGCOMM Comput. Commun. Rev.*, 5(2):28–42, 1975.
9. M. Salajegheh et al. CCCP: Secure remote storage for computational RFIDs. In *Proc. of USENIX Security*, 2009.
10. A. P. Sample et al. Design of an RFID-based battery-free programmable sensing platform. In *IEEE Transactions on Instrumentation and Measurement*, 2008.
11. Texas instruments application notes: Random number generation using the msp430. 2006.
12. B. Warneke, M. Last, B. Liebowitz, and K. S. Pister. Smart dust: Communicating with a cubic-millimeter computer. *Computer*, 34:44–51, 2001.
13. D. Yeager, R. Prasad, D. Wetherall, P. Powledge, and J. Smith. Wirelessly-charged UHF tags for sensor data collection. In *Proc. IEEE RFID*, 2008.

WISP Monitoring and Debugging

Richa Prasad, Michael Buettner, Ben Greenstein, and David Wetherall

1 Introduction

WISPs [11] often run out of energy while doing work, even when there would seem to be sufficient energy to complete the task at hand. Working out what went wrong (and characterizing what went right) is the subject of this chapter.

WISPs must work with small amounts of energy that are unpredictably available to carry out tasks that vary widely in complexity. WISPs receive all of their operating energy from readers. A commercial reader such as the Impinj Speedway [5] radiates 1 W, but only a small fraction of that power (6 mW at 1 m according to the Friis transmission equation [10]) is received by WISPs due to path loss associated with RF propagation. More distant WISPs, for example those beyond a meter, receive an even smaller share. Moreover, a reader frequently powers down as it follows the EPCglobal Class 1 Gen 2 protocol [1] and hops frequencies to conform to FCC and international spectrum regulations. Different amounts of energy are available to be harvested on different frequencies due to frequency selective fading, and, of course, no energy is available for harvesting when the reader is off. In this environment of severely limited availability of energy, the key to understanding WISP behavior is gaining visibility in WISPs' accumulation and use of energy (Fig. 1).

In this chapter, we present a tool we designed to provide visibility in WISP behavior, which consists of a daughter card for the WISP and its associated

R. Prasad
Microsoft, Redmond, WA, USA
e-mail: rich_pl@hotmail.com

M. Buettner • D. Wetherall
University of Washington, Seattle, WA, USA
e-mail: michael.buettner@gmail.com; djw@cs.washington.edu

B. Greenstein (✉)
Google, Seattle, WA, USA
e-mail: bengreenstein@gmail.com




Fig. 1 The C1G2 inventory sequence. From the EPCglobal C1G2 specification

firmware. The tool logs an accurate trace of WISP state, including the amount of energy it has available to it. It can monitor the WISP's demodulator, modulator, unregulated voltage, supervisor and demodulator enable lines, as well as parse any state information sent explicitly by having the WISP toggle debug pins. Information collected by the board can be sent over a serial connection to a PC for analysis. We describe the design considerations and implementation of the tool. We find that this tool creates only 18 cycles of overhead in WISP behavior and can catch interrupts from the WISP that are only $5.36\ \mu\text{s}$ apart.

The rest of this chapter is organized as follows. In Sect. 2, we motivate the need for our tool by describing the limitations of existing techniques to monitoring WISP behavior. In Sect. 3 we explain our approach to gaining visibility into WISP energy usage and describe the implementation of our tool. Section 4 evaluates the accuracy and performance of our tool. We present applications of our tool in Sect. 5 and conclude in Sect. 6.

2 Limitations of Existing Approaches

WISP behavior is highly sensitive to received power from the reader, and thus an accurate characterization requires a tool that can monitor the details of WISP activity and available energy without impacting its energy consumption. Additionally, to be practical, such a tool should be relatively inexpensive and unobtrusive, so that one can operate in realistic deployments and scale to monitoring several WISPs simultaneously. This section discusses several monitoring and debugging approaches and highlights their limitations.

2.1 Debugger

The joint test action group (JTAG) interface [4] that programs the WISP's microcontroller [8] can be used as a hardware debugger. By setting breakpoints in code,

WISP operation can be paused, and the microcontroller registers can be inspected to understand the WISP's execution state. Emerging profiling tools such as [3] can be used in conjunction with JTAG to optimize power consumption.

There are several limitations to this approach. First, JTAG supplies power to the WISP and is thus useless for understanding how the WISP harvests and uses power. Second, JTAG cannot be used to monitor state changes while the WISP is executing naturally. The WISP's microcontroller must be paused to view WISP registers, which disrupts normal program timing, since WISPs are largely driven by external events. Third, this method does not scale to several WISPs.

2.2 Oscilloscope and Data Acquisitions

WISP code can be modified to toggle unused pins in order to convey its state, which can be probed by an oscilloscope along with probing of lines used for communication and power management. The first limitation of this approach is that there are a limited number of oscilloscope probes, thus requiring the WISP developer to guess where a problem might be occurring and change the code at that point to toggle unused pins. Second, the traces captured on the oscilloscope are segmented and thus do not tell a contiguous story of WISP operation. Third, an oscilloscope is cumbersome, expensive, and does not scale to several WISPs.

Likewise, a data acquisitions (DAQ) device can sample the analog and digital lines of WISP pins to log a contiguous state trace on a host computer. DAQs however are large in size, expensive, have limited monitoring channels, and do not scale well to many WISPs.

2.3 Messaging

Information about a running WISP can be gathered by programming the WISP to transmit it in backscattered messages; this is a commonly used technique in the related field of sensor networks [7, 9]. The easiest ways to do this would be to overload the tag identifier to carry this additional information; EPC Class 1 Gen 2 conforming WISPs return a 96-bit identifier, so there is ample room to carry additional state. Alternatively, a reader-side program can be written to ask for state information explicitly, by using Gen 2's READ command.

This approach suffers from several limitations. First, it is difficult to get temporally fine-grained data due to processing overheads and limited backscatter bandwidth. Second, the natural behavior of WISP is changed since logic must be added to WISP firmware to overload identifiers with current state information or execution behavior must change to handle READ commands. Third, this approach fails in the common case of the WISP losing power between the recording of state information in memory and the eventual transmission of that information. Finally,

this approach has limited use in debugging. The program being debugged (the WISP firmware) is the same program responsible for transmitting messages; if it crashes, it won't transmit.

3 The WISP Monitor

Given the limitations of other approaches, we developed a monitor board that attaches to the WISP, reads its energy and state, and conveys this information via a serial back to a PC for analysis.

3.1 *Design Considerations*

We discuss the design considerations of this monitor board. First, the monitor should record information relevant to debugging and energy profiling. This includes monitoring the unregulated voltage and demodulator lines as well as receiving state information from the WISP. Examples of useful state information include determination of whether a parsed packet matches a Class 1 Gen 2 command, failures in reception or transmission, and checks for power sufficiency.

Second, the monitor should interoperate with the WISP without changing the WISP's behavior. The operating voltage of the WISP is 1.8 V. Attaching a monitor board to the WISP that runs at a higher voltage could lead to problems. If some pins on the monitor board are set as output, then the monitor will behave as an on-board power supply, thus changing the energy available to the WISP. Moreover, if the monitor board's pins are set as input, then the voltage threshold at which the monitor can detect interrupts could be higher than the voltage achievable by the WISP. Given these problems, it would seem to make sense to run the monitor board at the same voltage as the WISP.

Third, the times of WISP events must be recorded accurately. Accurate timestamps are key to logging good state traces. Since clock synchronization between a host computer and external device can suffer from unpredictable latencies, the monitor should have the capability to use its own clock for timestamping WISP events.

Finally, the monitor should incur minimal overhead in WISP software. We need rich state information without changing the WISP software significantly. There are three kinds of information that are central to the WISP state: communication, power management, and software decision information. The first two pieces of information can be accurately gathered by eavesdropping on the WISP input/output lines. A crucial benefit of this method is that it requires no change to the WISP software. Software decisions however occur internally and thus to gather this information we need to add "hooks" into the WISP software. These "hooks" are code segments that toggle a small number of unused microcontroller pins to indicate the decisions made by the WISP.




Fig. 2 The monitor board

3.2 Monitor Board Hardware

Here we describe the implementation of our monitor board (Fig. 2) and explain how it fulfills each of the design considerations stated in Sect. 3.1. Furthermore, we evaluate the performance of our monitor board on a set of microbenchmarks. Figure 3 shows the hardware architecture of the monitor board. The *USB-UART* block translates serial communication from the microcontroller into the USB protocol for communication with a host computer and vice versa. The microcontroller and *USB-UART* block are fed a 1.8 V supply from the regulator which uses the 5.0 V power supply from the USB as its input. This means that the monitor board has a constant source of power supply unlike the WISP. The microcontroller pins on the monitor board are connected to headers, to which the WISP can be directly affixed. This allows the board to monitor all exposed WISP pins. The WISP unregulated voltage pin is routed through the *voltage sampling* circuitry before being fed into

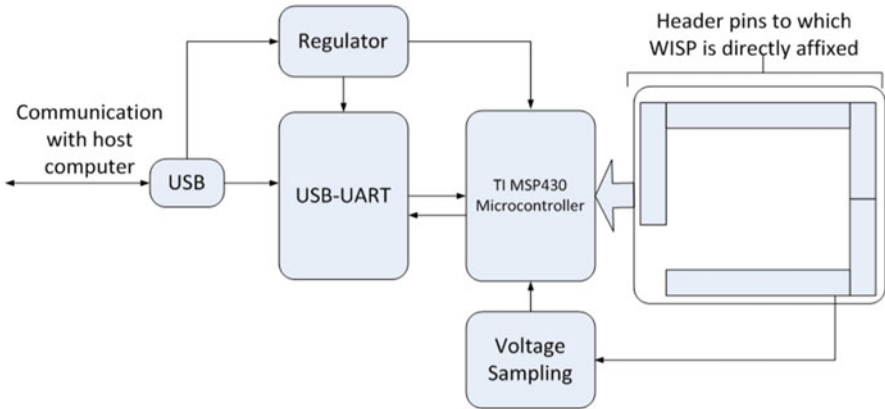


Fig. 3 Block diagram of monitor board hardware architecture

the monitor board's microcontroller. This circuitry ensures that the unregulated voltage is within microcontroller tolerable levels ($\leq 1.8\text{ V}$) and also ensures that the microcontroller is not drawing current from the WISP.

3.2.1 USB-UART Block

The *USB-UART* block consists of a FTDI FT232RL chip [2]. This chip connects to the USCI port of the microcontroller which performs UART communication. The chip translates the UART frames into USB protocol frames for transmission to the host computer and vice versa. The chip and microcontroller are powered by a 1.8 V supply from the regulator. If the chip was functioning at a higher voltage, then the voltage threshold of its pins to detect an interrupt would be greater than the voltage achievable by the microcontroller. In order to overcome this problem, we ensure that both modules function at the same operating voltage: 1.8 V.

3.2.2 Voltage Sampling Circuitry

If the monitor board's microcontroller were to be directly connected to the WISP unregulated voltage pin, the microcontroller would clamp down the WISP voltage to its operating voltage (1.8 V). This would change the WISP behavior dramatically. In order to prevent this, we add a high-impedance circuit in the form of two large resistors (20 k Ω and 1 k Ω). This divides the voltage by a factor of 3, thus mapping the upperbound of WISP unregulated voltage, 6.0 V, to 2.0 V. The microcontroller ADC operates by allowing the current to charge a capacitor for a software-defined period of time and then reporting the capacitor voltage. The factor of 3 division causes a nA current to be output from the high-impedance circuit, which would take

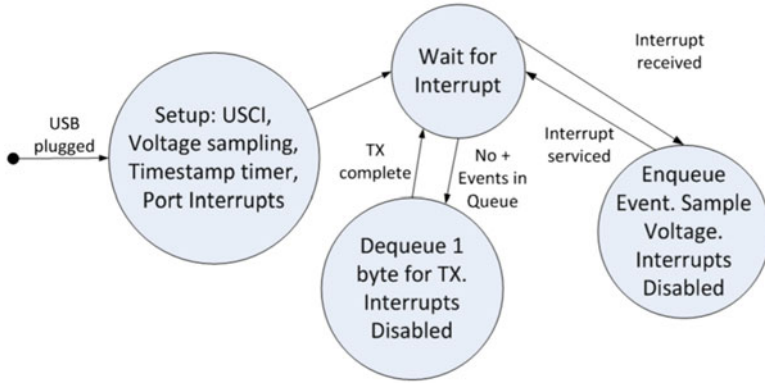


Fig. 4 Block diagram of monitor board software flow

a very long time to charge the capacitor. This would cause massive delays in the microcontroller code. Thus, we feed the nA current into an op-Amp which amplifies it to a few μA . This magnitude of current charges the capacitor quickly and allows the voltage sampling to be quick. In fact, we can attain accurate and stable results by waiting for only $10\mu\text{s}$ before sampling the capacitor voltage.

3.2.3 Monitor Board to WISP Pin Mapping

We use the same microcontroller as the WISP microcontroller, the MSP430F2132, on the monitor board. The USCI port of the microcontroller is connected to the FTDI FT232RL chip while its USI port is connected to the WISP USCI port. This enables us to perform direct real-time communication with the WISP if needed, although we do not use this feature. Since the microcontroller we use has interrupt capability only on ports 1 and 2, we attempted to map the WISP demodulator, modulator, supervisor, and sensor pins to these ports. We mapped the WISP unregulated voltage pin and three unused WISP pins to port 3.

3.3 Software

Figure 4 shows the software flow of the microcontroller code. The goal of the monitor board software is to log all wanted WISP events and to successfully transfer the log of events to the host computer. In order to accomplish these goals, we need to minimize the time spent in code segments where interrupts are turned off and ensure that the event rate is not greater than the transmission rate to the host computer. Here we describe the current software implementation and explain how it attempts to fulfill these goals.

3.3.1 Monitored Pins

The monitor board has the capability to monitor the WISP demodulator, modulator, unregulated voltage, supervisor and demodulator enable lines as well as parse any state information sent by the WISP through toggling of its unused pins (called *debug pins* henceforth). In practice, we monitor only a subset of these traces. For example, when interested in knowing when the WISP sends a packet, we do not need to sample the WISP modulator line, because packet sends can be indicated through the toggling of the WISP debug pins. Sampling a subset of the traces also reduces the rate at which events are generated, thus reducing the likelihood that the monitor board's event buffer overflows.

The WISP shares its state through the toggling of debug pins and can do so whenever it parses a valid or invalid command, completes transmission of a packet, and when it goes into a low-power state (e.g., the MSP430's *LPM4*) to build power. In theory, debug pins can be toggled at any point during WISP program execution, but in practice, setting debug pins should be avoided during packet reception and transmission to meet the very tight timing constraints of those operations.

The monitor board microcontroller samples the debug pins and also eavesdrops on the WISP demodulator line to detect the beginning of a packet. Upon detecting a positive edge on the demodulator line, the default operation of the monitor board firmware is to stop eavesdropping on the demodulator. Continued eavesdropping would make it possible to parse and collect every bit received by the WISP, but more commonly, we find it more helpful and concise to have the WISP toggle pins to indicate when entire commands have been received. When an *ID* is transmitted or an invalid command is received by the WISP, the monitor board re-enables eavesdropping on the WISP demodulator line. The monitor board by default also samples the WISP unregulated voltage line when a *Query* command is received, an *ID* packet is transmitted, an invalid command is received, and when the WISP goes into a low-power state to build power.

3.3.2 High-Level Description

We run the monitor board's microcontroller at a frequency of 5.6 MHz (approximately the fastest frequency at a 1.8 V operating voltage), which is at least 2.0 MHz faster than the fastest WISP frequency. This enables us to complete tasks quickly and prepare for the next WISP interrupt. When the monitor board microcontroller first starts executing code, it enables the USCI port for communication with the *USB-UART* block, sets up the ADC for WISP voltage sampling, enables a timer for timestamping the captured WISP events, and enables interrupt capability on the pins that map to wanted WISP events. In the default firmware, this corresponds to the WISP demodulator line and the debug pins. If there is a port interrupt, we enter the corresponding port interrupt service routine, log the event, and sample and log WISP voltage if appropriate. If there is no port interrupt, we check if there are events to be transmitted to the host computer. If this is true, a byte is transmitted. We loop

Fig. 5 Since the WISP can use only five debug pins to indicate its state, we map each reader command and other WISP states to a 5-bit integer value. The states we defined are presented; values 17–31 are undefined

5-bit Integer Value	Interpreted WISP State
0	Ready State
1	Arbitrate State
2	Reply State
3	Acknowledged State
4	Open State
5	Read Sensor State
6	Parse Packet Failure Type 1
7	Parse Packet Failure Type 2
8	Sleep
9	Query command
10	ACK command
11	Request RN command
12	Query Repeat command
13	Query Adjust command
14	Select command
15	NAK command
16	Read command

indefinitely between waiting for port interrupts and transmitting bytes to the host computer. Port interrupts are ignored during logging of an event, byte transmission, voltage sampling, and servicing of an interrupt. In order to minimize the time spent performing these activities, the monitor board transmits only one byte at a time, voltage sampling is set to its fastest rate of 10 μs, and servicing of an interrupt and logging of events are reduced by monitoring a subset of WISP events.

3.3.3 WISP State Information Collection

The WISP uses five debug pins to indicate its state. A high pin line indicates a binary 1, while a low pin line indicates a binary 0. The five debug pins together form a 5-bit number. Figure 5 shows the mapping of the debug pins configuration to the WISP state. We use the WISP debug pin that maps to the monitor board port 1 pin as the trigger for sampling the other debug pin lines. We call this pin the *trigger pin*. The other four pins are called the *LSB3*, *LSB2*, *LSB1*, and *LSB0* pins as per their position in the 5-bit number. Thus, for example, if the WISP wants to communicate to the monitor board that it just received a *Query* command, it would set the *LSB0* and *LSB3* pins high, ensure that the *LSB1* and *LSB2* pins are set low, toggle the trigger pin high to generate an interrupt, and then immediately set the trigger pin low so as to ensure that it is sampled as a binary 0 by the monitor board. Upon the receipt of this interrupt, the monitor board samples all five pins, maps their states to binary 1's and 0's, and logs the 5-bit integer value (9 in the case of the *Query* command) in RAM. The application on the host computer receives this value and maps it to the corresponding state (*Query* in this example).

3.3.4 Event Buffer Implementation

All captured WISP events are pushed into a circular queue structure. Each event consists of six bytes. The first four bytes correspond to the timestamp of the event, while the last two bytes indicate the type of event. When the queue is full, captured events are not logged, thus resulting in dropped events.

3.3.5 Timestamp Implementation

A timer is enabled when the monitor board first starts executing code. The timer counter rollover occurs every 90 ms, and thus we maintain a wraparound variable which counts the number of times the timer counter has rolled over. A timestamp consists of four bytes. The first two bytes are the wraparound value while the last two bytes are the current timer counter value.

Each timer tick is of duration $1.3\ \mu\text{s}$. Each positive edge in the *Miller-4* encoding is $25\ \mu\text{s}$ apart from the previous positive edge. Thus, the timestamp on events is accurate. In order to calculate energy accurately, we need accurate timestamps. Furthermore, the timestamps on commands can also reveal timing issues in the WISP software.

4 Evaluation

We evaluate the performance of the monitor board based on these microbenchmarks: timestamp accuracy, voltage resolution and accuracy, overhead, interrupt capture latency, and cross correlation of gathered data across multiple WISPs.

4.1 Timestamp Accuracy

We measured the timestamp accuracy by generating five consecutive interrupts on the WISP and capturing them on the monitor board and a Tektronix TDS7254 2.5 GHz oscilloscope, which serves as the ground truth. We compared the time interval reported by the monitor board with the one reported by the oscilloscope and found that the absolute error to be $\pm 9\ \mu\text{s}$.

4.2 Voltage Resolution and Accuracy

Voltage resolution is the smallest quanta of voltage that the monitor board can measure. The monitor board has a 10-bit ADC, its operating voltage is 1.8 V, and

Fig. 6 The overhead in WISP software for indicating its state to the monitor board is 18 or 19 CPU cycles per state

Command	WISP Overhead (# CPU Cycles)
NAK	19
QUERY	18
QUERY REPEAT	19
QUERY ADJUST	19
ACK	19
SLEEP	18

there is a voltage divider that reduces WISP voltage by a factor of 3 before feeding it into the monitor board's microcontroller. This means that it has 1,024 values to read between 0 and 1.8 V. Thus, its resolution is $(1.8 \times 3)/1024 = 5.25$ mV/bit. The WISP consumes $529.4 \mu\text{A}$ at minimum when running in active mode. This means that the monitor board will detect a voltage change every $90 \mu\text{s}$ or lesser (Eq. 3). As per the settings in our experimental study, the minimum duration of a bit is $25 \mu\text{s}$ (a binary 0) and hence the monitor board can report the power consumption every 4 bits:

$$I = C \frac{\Delta V}{\Delta t}, \quad (1)$$

$$\Delta t = \frac{C}{I} \Delta V, \quad (2)$$

$$\Delta t = \frac{10; \mu\text{F}}{529.4 \mu\text{A}} \times 1.75 \text{ mV} = 90 \mu\text{s}. \quad (3)$$

We measure the accuracy of voltage measurements by providing the WISP with constant power through a power supply and sampling the WISP voltage on the monitor board for 10 s. We perform this experiment for three voltage values: 2.0 V, 3.0 V, and 4.0 V. We find that the absolute error is ± 10 mV.

4.3 Overhead

Overhead is the number of cycles spent by the WISP to indicate its state to the monitor board. The CCSTEP register of the microcontroller indicates the number of cycles spent per C code instruction. Figure 6 shows the number of cycles spent by the WISP to share state information. We see that the WISP overhead is 18 to 19 cycles per state information. The WISP runs at 3.5 MHz, which means that it spends 5.14 to $5.42 \mu\text{s}$ to communicate its state. The most time-sensitive code segment in which the WISP communicates with the monitor board is between receiving and transmitting packets. For our reader settings, the minimum wait time between commands is $45 \mu\text{s}$. Thus, the WISP overhead is reasonable.

Fig. 7 Time spent in activities during which all interrupts are turned off

Activity	Debug Board Latency (us)
Interrupt Service Routine	3.57 (Minimum), 82 (Maximum)
Voltage Sampling	31
Enqueue event	29.4
Dequeue event	5.71

4.4 Interrupt Capture Latency

Interrupt capture latency is the minimum duration needed between consecutive interrupts for them both to be captured successfully by the monitor board. The monitor board ignores interrupts when it is sampling voltage, transmitting a byte to the host computer, logging an event, or servicing an interrupt. Figure 7 shows the time consumed by each of these activities. We obtained these measurements by counting the number of cycles in the CCSTEP register for each activity.

When an interrupt is captured, the interrupt flag register is immediately copied and cleared, thus enabling the capture of a consecutive interrupt on the same pin regardless of the time spent completing the service of the interrupt. Thus, the minimum interrupt capture latency in this case is only $3.57\mu\text{s}$. However, if two consecutive interrupts occur on the same pin while the first interrupt is being serviced, the monitor board will capture the first of the two consecutive interrupts but miss the second one. In this case, the interrupt capture latency is $82\mu\text{s}$ at most.


5 Applications

The monitor board provides visibility into WISP behavior that is not possible or prohibitively tedious to achieve using other techniques. There are two domains where this is particularly useful: debugging firmware and gathering fine-grained data for evaluation. In the former case, debug interrupts can be triggered at key points in the WISP's execution to give the developer insight into what code is being executed and when. In the latter case, data about how often tasks, or even single operations, are being executed and how much energy they consume can be gathered in situ. We present examples from our own experience that illustrate these uses.

5.1 Debugging

The monitor board allows 32 codewords to be used to describe changes in WISP state. As triggering the debug interrupts is very lightweight, debugging output can be triggered at essentially any point in the code execution. As an example, Fig. 8 shows a trace of WISP behavior that was gathered from the monitor board. Debug points were inserted after each event completed to indicate when each reader command was received and when each WISP response was transmitted. Additionally, debug

Fig. 8 Annotated trace generated from monitor board data



points were inserted to indicate when the WISP wakes up and sleeps. We can see from the trace exactly what code is being executed on the WISP and also what the voltage of the WISP's capacitor is at each stage in the protocol.


In the example, the WISP wakes up when the voltage of the capacitor is approximately 2.1 V and it immediately takes a sensor reading from the accelerometer. This takes close to 5 ms and results in the voltage of the capacitor being reduced by around 0.3 V. The WISP then processes two Query commands sent by the reader and transmits its *RN16* (not shown), and the reader transmits an *ACK*. The WISP next sends its ID (which contains the sensor reading) and then goes to sleep. The WISP wakes up again approximately 8 ms later, and a similar exchange follows. By having such detailed data, developers can determine where in the execution path firmware may be failing and also track energy consumption for individual operations enabling them to understand where their code could be optimized.

5.2 Gathering Evaluation Data

Prior work on WISPs have generally based their evaluations on two forms of data: results based on response rate as measured by the reader, and those based on one (or a few) hand measurements taken with an oscilloscope or power meter. This is because it is tedious to gather large sample sizes for any metric other than read rate. The monitor board allows researchers to gather more fine-grained data and from devices that are deployed in situ.

One instance where we used the monitor for gathering evaluation data was when we experimented with a full implementation of the Gen 2 MAC protocol. Previous WISP MAC implementations used a simple scheme where WISPs attempt to transmit data whenever they have sufficient energy to do so. However, this results in many collided transmissions when multiple WISPs are present. The Gen 2 specification describes an anticollision mechanism, but it had been thought to be too energy intensive for use on the WISP.

Fig. 9 Voltage drop for Gen 2 communication in situ




To characterize how much energy is consumed when WISPs use a full Gen 2 implementation, we implemented the protocol and experimented with a group of WISP at varying distances from the reader. We wanted to see how much energy was used during communication and also how this varies with distance. To do this, we triggered a debug message when the tag begins communicating with the reader and another when communication is finished. This gave us a trace that showed when communication started and ended, and also what the voltage of the capacitor was at each of these points.

Figure 9 shows the results of this experiment, with each item in the figure consisting of more than 1,000 data points. When the WISP is closer to the reader, the drop in voltage is less, because stored energy is supplemented with energy being harvested during task execution. This effect was easily measured because data was gathered from a WISP in situ. The second thing we saw is that the drop in voltage varied widely from execution to execution. This is because the Gen 2 protocol is nondeterministic in the number of messages that the WISP must process to communicate. Using the monitor board, we could easily gather data from thousands of executions to gain an understanding of the distribution in energy usage.

Another place where the monitor board is useful is for gathering task completion rates for tasks that do not communicate with the reader. For example, we were experimenting with a motion detection system where the host computer wanted to know when a WISP moved. We were comparing two approaches. The first had the WISP determine if it was moving and only communicate with the reader when it detected motion. The second approach had the WISP transmit every sample to the host using the Gen 2 protocol and the motion detection was implemented there.

For our study, we wanted to compare the sampling rates of the two approaches across varying distances to give us a proxy of system responsiveness. This comparison would be difficult without the monitor board, as determining the responsiveness of the first approach would be difficult because it only responds to the reader occasionally. Figure 10 shows the results of this experiment, based on data we gathered using the monitoring board. To measure task execution rate of the first

Fig. 10 Task execution rate for sampling a sensor and for sampling and transmitting the data



task, we simply inserted a debug point after every execution. We could measure the change in performance as the WISP moved further from the reader, which showed that the second approach was infeasible beyond approximately 1.5 m.

6 Conclusion

This chapter presented an approach to monitoring and debugging WISPs.¹ The monitor tool provides a way to gain visibility into WISP state trace with negligible impact on its behavior. Since it is externally powered and is connected to WISP debug pins, it can accurately log energy usage and map it to WISP program state changes.

References

1. EPCglobal. Epc radio-frequency identity protocols class-1 generation-2 uhf rfid protocol for communications at 860 mhz–960 mhz version 1.0.9. 2005.
2. FTDI. Future technology devices international ltd. - ft232r. <http://www.ftdichip.com/Products/FT232R.htm>. Accessed 17 Dec 2012
3. IAR. What is power debugging? <http://www.iar.com/en/Products/IAR-Embedded-Workbench/Power-debugging/>. Accessed 17 Dec 2012
4. IEEE. IEEE Standard Test Access Port and Boundary - Scan Architecture. <http://standards.ieee.org/>. Accessed 17 Dec 2012. IEEE Std 1149.1, 1990. pp. 0–1
5. Impinj. Impinj’s uhf gen 2 speedway rfid reader. <http://www.impinj.com/products/rfid-reader.aspx>. Accessed 17 Dec 2012
6. R. Prasad. Energy debugging for rfid sensor networks. In *University of Washington Department of Computer Science and Engineering Masters Thesis*, May 2009.

¹This chapter is based on work first presented in Richa Prasad’s masters thesis [6].

7. N. Ramanathan, K. Chang, R. Kapur, L. Girod, E. Kohler, and D. Estrin. Sympathy for the sensor network debugger. In *SenSys '05: Proceedings of the 3rd international conference on Embedded networked sensor systems*, pp. 255–267, New York, 2005. ACM.
8. TI. Getting started with msp430 from texas instruments. http://www.ti.com/lscds/ti/microcontroller/16-bit_msp430/getting_started.page. Accessed 17 Dec 2012
9. K. Whitehouse, G. Tolle, J. Taneja, C. Sharp, S. Kim, J. Jeong, J. Hui, P. Dutta, and D. Culler. Marionette: using rpc for interactive development and debugging of wireless embedded networks. In *IPSN '06: Proceedings of the 5th international conference on Information processing in sensor networks*, pp. 416–423, New York, 2006. ACM.
10. Wikipedia. Friis transmission equation. http://en.wikipedia.org/wiki/Friis_transmission_equation. Accessed 17 Dec 2012
11. D.J. Yeager, A.P. Sample, and J.R. Smith. Wisp: A passively powered uhf rfid tag with sensing and computation. In *RFID Handbook: Applications, Technology, Security and Privacy*. CRC Press, Boca Raton, 2008.

Part IV
Cryptography and Security
for Computational RFID

Maximalist Cryptography and Computation on the WISP UHF RFID Tag

Hee-Jin Chae, Mastooreh Salajegheh, Daniel J. Yeager, Joshua R. Smith, and Kevin Fu

1 Introduction

Because of computational constraints on many RFID tags, classical cryptographic primitives such as block ciphers and asymmetric cryptography were thought to be unrealistic on a low-resource tag [8]. To this end, many lightweight cryptographic protocols have been proposed [3, 5, 14–16, 18, 19]. However, many such protocols have serious vulnerabilities [6, 10, 11]. Moreover, the lack of a development platform makes it difficult to determine the feasibility of proposed cryptographic schemes. Thus, a popular approach is to minimize cryptographic operations to ensure feasibility on an RFID tag. This minimalist approach [8] can leave spare computational resources unused. An open question is then how to best maximize the security on an RFID tag to fully utilize available computational resources.

Our approach to maximizing security relies on low-power microcontrollers. Continuous improvements in efficiency of microelectronics (i.e., required energy per instruction) now enable wirelessly powered, general-purpose microcontrollers—infesible at any reasonable range a few years ago. Such microcontrollers make traditional cryptographic methods more feasible on RFID tags.

H.-J. Chae • M. Salajegheh (✉) • K. Fu
University of Massachusetts, Amherst, MA, USA
e-mail: chae@cs.umass.edu; negin@cs.umass.edu; kevinfu@cs.umass.edu

D.J. Yeager
Department of Electrical Engineering, University of Washington, Seattle, WA, USA
e-mail: yeagerd@gmail.com

J.R. Smith
Department of Computer Science and Engineering, Department of Electrical Engineering,
University of Washington, Seattle, WA, USA
e-mail: jrs@cs.washington.edu

Using a maximalist approach to cryptography, our results show that an RF-powered UHF tag can perform strong encryption. We provide preliminary experimental results of implementing RC5-32/18/16 [23] on WISP (Wireless Identification and Sensing Platform) [26]—a battery-free platform powered and read by a standards compliant UHF RFID reader running the EPC Class 1 Gen 2 protocol. Our contributions include:

1. We provide preliminary experimental data on how much computation is available on a microcontroller-based RFID tag.
2. We show that symmetric cryptography is feasible on an RF-powered, general-purpose RFID tag. To the best of our knowledge, this is the first implementation of conventional cryptography on an RF-powered UHF RFID tag.

This book chapter is based on a paper [4] presented in RFIDsec 2007 Conference. The paper gives measurements based on an earlier version of the WISP which uses MSP430F1232. Having only 256 bytes of RAM available, the older WISP would allow only 12 rounds of RC5. The platform used in this book chapter is the new WISP 4.1 which has more RAM available and allows more rounds of RC5 (up to 51).

2 Background on WISP Architecture

This section provides a condensed background on the WISP, more fully described in [25, 26]. The WISP and its block diagram appear in Fig. 1. An antenna and impedance matching circuit precede the analog front end. The power harvester block rectifies incoming RF energy into DC voltage to power the system. The demodulator follows the envelope of the RF carrier wave to extract the Amplitude-Shift-Keyed (ASK) data. This extracted baseband waveform is read by the MSP430 microcontroller to receive downlink data from the reader. Uplink data is sent via the modulator circuit, which functions by changing the antenna impedance. The WISP is made of a two-layer FR4 PCB with components limited to the top side. A dipole




Fig. 1 WISP [25, 26] is a batteryless, microcontroller-based UHF tag that implements RFID protocols in software. WISP uses a TI MSP430 microcontroller and minimal analog circuitry to function as a UHF RFID tag

antenna made of 22 gauge (0.6 mm diameter) copper magnet wire is visible. Small header pins expose all ports of the microcontroller for expansion to daughter boards, external sensors, and peripherals.

2.1 Analog Front End and Tuning

Due to the relatively high power consumption of WISP, its rectifier is designed to supply more current than ordinary tags. This circuit is excited by commercial, EPC Class 1 Generation 2 compliant readers operating at 902–928 MHz with an allowable transmission power of $4W_{\text{EIRP}}$ (Effective Isotropic Radiated Power).

Efficient conversion of the incoming RF energy to DC power for the tag maximizes range. A matching network provides maximum power transfer from the antenna to the rectifier, and a five-stage voltage doubling circuit converts the incoming power to voltage. Low threshold RF Schottky diodes maximize the voltage output of the rectifier. Finally, this rectified DC voltage is stored in a large capacitor and supplied to a 1.8 V regulator to power the WISP.

2.2 Demodulation and Modulation

To encode reader-to-tag data, the reader amplitude-modulates the 915 MHz RF carrier wave it emits. Normally the carrier waveform remains at a constant amplitude; when bits are transmitted, the amplitude of the carrier drops to approximately ten percent of its normal value. The duration of the low “break” indicates a logical “one” or a “zero.” A short break (1.78 μs) indicates a “zero,” and a long break (5.34 μs) indicates a “one.” To decode this data, the RF signal is fed through a small (two-stage) voltage doubling rectifier in parallel with the main (five-stage) harvester. We call the two-stage harvester the “mini-harvester.” The time constant of the mini-harvester is much smaller than that of the main harvester, allowing it to track the dynamic range of the incoming bits. The first two voltage doubling stages of the mini-harvester, in conjunction with a lowpass filter, effectively demodulate the 915 MHz carrier, and leave a baseband data signal on the order of 70 kHz. A final “extra” diode performs an additional rectification step, removing the 70 kHz data signal and leaving a slowly varying average power level (i.e., just fast enough so that it can change on the timescale that the tag moves in space, say 10 Hz) that provides a dynamic reference for bit detection.

The 70 kHz data signal is fed through a Schmitt trigger inverter that thresholds this waveform to remove noise and glitches. Finally, a level shifter converts the relative magnitude of the incoming data waveform into a 1.8 V logic level for the MSP430. The slowly varying average power signal serves as the power supply for the Schmitt trigger and level shifter.

RFID tags do not actively transmit radio signals. Instead they modulate the impedance of their antenna which causes a change in the amount of energy reflected back to the reader. This modulated reflection is typically called backscatter radiation. In order to change the impedance of the antenna, a transistor is placed between the two branches of the dipole antenna. When the transistor conducts current, it short circuits the two branches of the antenna together, changing the antenna impedance; in the non-conducting state, the transistor has no effect on the antenna, and thus the power harvesting and data downlink functions occur as if it were not present. This impedance modulation is currently implemented with a 5 GHz RF bipolar junction transistor which allows for effective shunting of the 915 MHz carrier wave.

2.3 Digital Section and Power Conditioning

The WISP's general-purpose computation capabilities are provided by an ultra-low-power microcontroller, the TI MSP430F2132. This 16-bit device has just over 8 KBytes of flash memory, 512 bytes of RAM, and a 10-bit, 200 kilo-samples per second Analog to Digital Converter (ADC). The WISP 4.1 runs TI MSP430 at 6 MHz with a 1.8 V supply voltage.

In active operation at 1.8 V, the microcontroller consumes approximately 690 μ A to 1.4 mA. Erasing and writing flash both require approximately 1 mA. Reading from flash requires no additional power (above ordinary active-mode operation). The flash in the MSP430's information memory must be erased in 64-byte blocks; its main memory is erased in 512-byte blocks. Bytes are individually writable. Erasing the flash memory takes more time and therefore more energy than writing or reading flash. Erasing a block requires 1.24 ms; writing a byte takes 7.7 μ s; reading a byte from flash requires no more time than accessing RAM.

The MSP430 has various low-power modes that are very useful for wirelessly powered operation. Its lowest-power mode, "RAM-retention mode," which stops all computation but maintains state, requires only 0.1 μ A at 1.6 V. The low-power consumption of this relatively new device is a critical factor in enabling use of a general-purpose microcontroller in RF-powered RFID systems.

3 Related Work

Because of the resource-constrained nature of nodes in sensor networks, many of the design criteria for security coincide with that for security of RFID tags. SPINS [17] and TinySec [9] both present experimental data on implementation of security protocols in sensor networks. Both works implemented RC5 for their block ciphers because of its small code size and high efficiency, and they show that RC5 provides a balance between security and performance for sensor networks. Although sensor

Table 1 Comparison of WISP with other RFID devices

Platform	Power	Computing	Storage	Communication	Distance
EPC Gen1 [20]	UHF RF	State machine	64-bit	UHF backscatter	3–7.5 m [28]
EPC Gen2 [20]	UHF RF	State machine	96/128-bit	UHF backscatter	3–7.5 m [28]
WISP 4.1 [25]	UHF RF	16-bit 8 MHz	8 KB+ 256-byte flash 512-byte RAM	UHF backscatter	< 4.5 m
SoCWISP [30]	UHF RF	State machine	96/128-bit	UHF backscatter	3 m
DemoTag [1]	Battery	8-bit 16 MHz	4 KB EEPROM 4 KB SRAM 128 KB flash	HF backscatter	N/A
TELOSB [13]	Battery	16-bit 8 MHz	48 KB flash 10 KB RAM	UHF/ 2400 MHz	75–100 m
Microchip MCRF202	Inductive	12-bit 400 kHz	96/128-bit	LF backscatter	1.3–10.2 cm [27]
Proxmark3 [29]	USB	32-bit 16 MHz	256 KB flash 64 KB SRAM	LF and HF backscatter	N/A
RFIDGuardian [22]	Battery	32-bit 520 MHz	16 MB flash 64 MB SRAM	HF backscatter	0.5 m

nodes have limited resources, general-purpose RFID tags are even more limited in memory, power, and computing capabilities. Sensor nodes like the Mica2 platform have their own power source onboard, and they have orders of magnitude more storage. Table 1 provides side-by-side comparison of different devices with RF interfaces. WISP is most comparable to general-purpose RFID tags such as EPC Gen1 and Gen2.

There are many HF tags capable of cryptography because of their common usage in building access control and authentication applications [12]. HF tags can afford to be more expensive in terms of manufacture cost and therefore provide stronger cryptography because of the demand for tags with built-in security and privacy mechanisms, which is usually achieved with a separate cryptographic engine. Israsena proposes the Tiny Encryption Algorithm (TEA) as a suitable encryption engine for low-cost RFID applications [7]. Three different architectures for the TEA encryption algorithm are presented that fit within the budget of a five-cent per tag cost. Aigner and Feldhofer discuss their results on an implementation of Tiny AES (TINA) as an ASIC [2]. Although both systems comply with stringent

requirements for low-cost RFID systems (average current consumption below 10 μ A), they depend on separate special-purpose circuitry. While ASICs are efficient in terms of power consumption and cost, they are inflexible and limited to a narrow set of applications. Because of this inflexibility, the cost for prototyping and development is relatively high—making it difficult and time-consuming to realize in actual applications. To the best of our knowledge, our system provides the first UHF RFID tag using general-purpose microcontroller to provide cryptographic capabilities under nontrivial computational constraints.

Despite the extreme resource limitations of UHF tags, mementos [21] make energy and/or time-intensive computations possible on programmable RFID tags by dividing the workload among life cycles. Mementos system opens the door for conventional cryptography algorithms with no requirement for software modifications. On the other hand, systems like CCCP [24] show that employing the already developed security schemes in WISP is possible if the capacity of the tags is fully exploited and the cryptography algorithms are chosen carefully.

4 Measurements and Estimates of Computation and Power Consumption

In this section, we demonstrate that symmetric cryptography is feasible on microcontroller-based RFID tags through an empirical study of RC5-32/18/16 on the WISP. Experimental data shows how much computation is available on WISP with varying workloads. Although WISP is relatively power inefficient compared to ASIC designs, the logic gate count of WISP is comparable to an EPC UHF tag—enabling a reasonable measurement of computational capabilities and limitations of an RFID tag.

4.1 Measurements of Computation and Symmetric Cryptography

To show the feasibility of classical cryptography on a general-purpose RFID tag, we present experimental results of implementing RC5 on WISP. We have chosen RC5 because of its simplicity and relatively small memory requirements. RC5 is implemented with 32-bit word, 18 rounds, and 16 bytes of secret key. The 16-byte secret key is stored in flash. Even with careful choices of block cipher parameters, there exist practical challenges in implementing RC5 on such a resource-constrained platform. Because of the extremely limited RAM memory




Fig. 2 WISP lifecycle

(512 bytes), minimizing stack size is crucial. For instance, our first implementation resulted in an immediate stack overflow upon running the key scheduler. Since RC5 requires expanded key table of size $2(r + 1)$ words, where r is the number of rounds, careful memory planning is required to reduce any unnecessary memory writes. The current implementation fully utilizes all of its RAM memory so that expanded key table $(2(r + 1) * 4 \text{ bytes}) = 104 \text{ bytes}$ is kept in RAM along with just enough room for the stack. One of the approaches to remedy possible stack overflow would be to use a precomputed expanded key table. It can be precomputed and stored in ROM or flash as long as the secret key remains the same without hurting the performance since reading from flash or ROM takes the same amount of power and cycles as reading from RAM. In our implementation, we compute the extended key table once in every hardware reset. In other words, the key table is computed during the first active cycle and kept in RAM unless WISP reaches a brownout voltage. This decision is based on the assumption that it is unlikely for the secret key to be changed while an RFID reader is polling.

All three major functions of RC5—`setupKey()`, `encrypt()`, and `decrypt()`—have been implemented and measured on WISP. Our data comes from running WISP at 3 MHz at 1.8 V. `setupKey()` is executed once after hardware reset, and the 64-bit ID value is encrypted or decrypted once every duty cycle (e.g., during “Generate Packet” stage in Fig. 2).

The duty cycling is enabled by a hardware voltage supervisor. This supervisor creates software interrupts by toggling a microcontroller input high to indicate sufficient voltage for operation. When the supervisor detects a voltage of 2.0 V or greater, the microcontroller is enabled and a computation is initiated. At the end of the computation, if the voltage is between 1.8 V and 2.0 V the microcontroller enters very low-power (RAM-retention only) mode. While in this low-power mode, the harvester provides power to recharge the supply capacitor back to 2.0 V for the next computation. Below 1.8 V, a hardware reset is generated by an insufficient voltage indicator output on the voltage regulator.

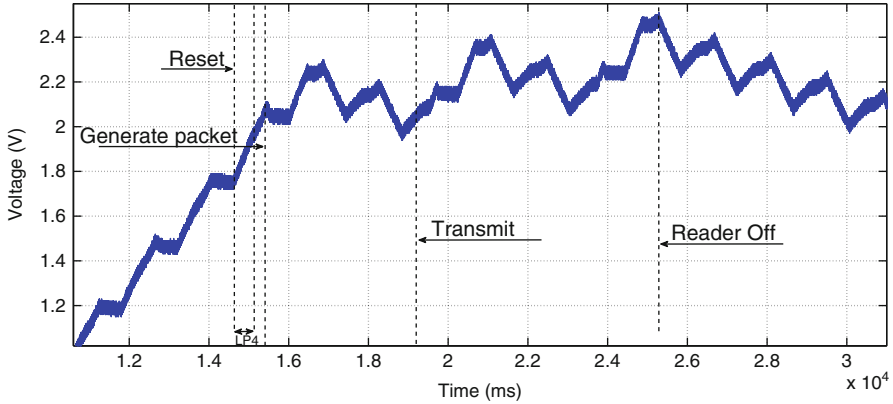


Fig. 3 An annotated scope trace of WISP execution and power consumption. The *solid line* measures the supply voltage V_{OUT} in 2 V increments on the y-axis. While the reader is ON (i.e., sending queries), the WISP’s voltage level stairs up. The voltage supervisor wakes up the WISP from LP4 when the voltage level exceeds 2.0 V, and WISP begins its computation (“Generate Packet”). When the reader receives the WISP’s response, the reader stops sending queries, therefore cutting off RF power to the WISP. This is observed as a gradual decline of voltage at the *right side* of the figure. The WISP first enters LP4 and then resets as the voltage level falls below the minimum operating voltage. The total latency from 0 V until the end of RF response transmission in this case is approximately two seconds

Table 2 Execution time for operations in RC5 with 32-bit words, 18 rounds, and a 16-byte secret key on a WISP tag

RC5-32/18/16 function	Execution time (ms)	Energy consumption (μ J)
setupKey ()	27.84	25.75
encrypt ()	5.40	6.92
decrypt ()	5.39	5.56

The WISP received uninterrupted power
Encryption happens on 64-bit messages

Figure 3 depicts a typical lifecycle of WISP with RC5 encryption enabled and voltage level patterns throughout the WISP duty cycle after a hardware reset, captured by an oscilloscope. WISP is in active mode during “Generate Packet” and “Transmit and Receive” stages.

Table 2 presents the execution time and the energy consumption for three operations of RC5 on WISP tags operating at a distance of 1 foot (0.3 m). The time measurements are the mean of five separate trials of execution after a hardware reset. An external power supply charges the WISP to 4.5 V and then let the WISP run each function using the energy stored in the capacitor.




Fig. 4 This estimate shows the computation available in one active cycle, where the power harvester charges a 10 uF capacitor to its peak value and then the microcontroller enters active mode, discharging the capacitor down to the minimum operating voltage. The Instructions Available line ends abruptly at the distance where the harvester is no longer able to supply the microcontroller’s minimum operating voltage

4.2 Estimates of Maximal Computation and Measurements of the Effect of Flash Writes on Computation

Figure 4 provides an estimate of available microcontroller instructions versus wireless distance. The model used to generate this plot is based on experimental data of WISP performance, on the Friis transmission equation, and on published microcontroller power consumption specifications.

The Friis transmission equation provides an estimate of power received as a function of distance from the transmitter, transmit power, antenna gains, and wavelength:

$$P_R = P_T - 20 \log \left(\frac{4\pi d}{\lambda} \right) + G_T + G_R. \tag{1}$$

The RFID reader power $P_T = 30dBm$; the receive and transmit antenna gains are given by $G_R = 2dBi$ and $G_T = 6dBi$, respectively. The wavelength $\lambda = 0.33m$ at 915 MHz. Using the Friis equation with these parameters, we can find the expected power available as a function of distance.

In [25], experimental data of output voltage as a function of input power for the WISP power harvester is given. Thus, this data, together with the Friis equation, allows us to predict the WISP’s output voltage V_{rec} as a function of distance d . This data captures the finite efficiency of the WISP harvester. An ideal, lossless harvester could produce any desired voltage if it were able to accumulate energy long enough. Actual WISP devices reach a steady state voltage, at which point power lost and power harvested balance.




Fig. 5 Preliminary measurements of the number of cycles available after writing to flash—computed after a minimal set of operations during the “Generate Packet” stage and without tag response over RF. Each data point represents a mean of ten samples

The TI-MSP430 data sheet specifies its current consumption for given voltage and operating frequencies. This allows us to calculate its energy per instruction: 281 pJ per instruction at $V_{dd} = 1.8$ V or 900 pJ per instruction at $V_{dd} = 3.0$ V.

Of the energy stored in the WISP’s storage capacitor, only the voltage above the microcontroller’s minimum operating voltage is usable. If the capacitor is initially charged to V_{rec} and its minimum operating voltage is V_{dd} , then $\frac{1}{2}CV_{rec}^2 - \frac{1}{2}CV_{dd}^2$ Joules are available to run the computation. Using the previously calculated dependence of V_{rec} on distance, we find energy available as a function of distance. Dividing by energy per instruction provides number of instructions as a function of distance.

Our actual measurements in Fig. 5 show the relationship between the reader-to-tag distance and number of cycles available in one duty cycle with varying workloads of flash writes to the MSP430 information memory. While Fig. 4 provides the estimated upper bound on how much computation we can afford on WISP, Fig. 5 presents how many cycles are available after flash writes of different sizes. For consistency, the experiment involves two steps, initialization and measurement. During the initialization, WISP’s flash memory is checked if its state is erased or not. The flash is wiped out (writing ‘1’ in all bits) before each measurement is taken. During the measurement phase, WISP consumes power writing to flash then raises a bit in a loop until WISP completely runs out of power (brownout voltage of 1.5 V). These bit raises appear as a pulse on the oscilloscope. Since raising a bit requires four cycles, the number of bits is counted to compute available number of cycles after writing different number of bytes to the flash. Note that a minimal set of operations are done during the “Generate Packet” stage in this experiment (e.g., preloading the 64-bit ID into an array) and there is no RC5 computation or CRC computation.

We observed that the number of cycles available does not deviate significantly from each other at short read ranges. At the medium range of 0.5 m, the power consumption of a flash write causes a significant decrease in available computation. Beyond the distance of 0.75 m, WISP is not able to reach the minimum operating voltage of 2.7 V for flash writes; therefore the plot lines for flash writes end abruptly. These measurements are preliminary, and further research is necessary to determine a more precise relationship between flash writes, erase segment sizes, and computation. However, one conclusion is that unnecessary writes to flash at midrange distances will significantly reduce available computation.

5 Conclusions

Our preliminary experimental data shows that UHF RFID tags with cryptographic capabilities are no longer infeasible. Although our experimental platform exceeds the current EPC UHF Class 1 tags in terms of computing power and storage, we believe that the device is comparable and a good representation of the future of UHF RFID tags. We believe that the trend in microelectronics will continue to bring power-efficient and cost-effective microcontrollers capable of more sophisticated computation. We hope that our work will encourage further research to determine the feasibility of maximal-strength cryptography with actual power measurements of RF-powered UHF RFID tags—beyond algorithmic estimates of space and running times.

Acknowledgments We thank Dan Holcomb and Thomas Heydt-Benjamin for their assistance in debugging experimental setups and for giving critical feedback; Salma Mirza for her help in conducting experiments with WISP; and Peter Desnoyers and Gaurav Mathur for their advice on the MSP430. This material is based upon work supported by the National Science Foundation under Grant No. 0627529.

References

1. Aigner, M.: DemoTag (2006) Last viewed May 16, 2007 http://www.iaik.tugraz.at/research/vlsi/02_products/05_rfid_demotag/.
2. Aigner, M., Feldhofer, M.: Secure symmetric authentication for RFID tags. In: Telecommunication and Mobile Computing – TCMC 2005, Graz, Austria (2005)
3. Calmels, B., Canard, S., Girault, M., Sibert, H.: Low-cost cryptography for privacy in RFID systems. In Domingo-Ferrer, J., Posegga, J., Schreckling, D., eds.: International Conference on Smart Card Research and Advanced Applications – CARDIS. LNCS, Tarragona, Spain, IFIP, Springer-Verlag (2006)
4. Chae, H.J., Yeager, D.J., Smith, J.R., Fu, K.: Maximalist cryptography and computation on the WISP UHF RFID tag. In: Proceedings of the Conference on RFID Security. (2007)
5. Cui, Y., Kobara, K., Matsuura, K., Imai, H.: Lightweight asymmetric privacy-preserving authentication protocols secure against active attack. In: International Workshop on Pervasive Computing and Communication Security – PerSec 2007, New York, USA, IEEE (2007) 223–228

6. Defend, B., Fu, K., Juels, A.: Cryptanalysis of two lightweight RFID authentication schemes. In: International Workshop on Pervasive Computing and Communication Security – PerSec 2007, New York, USA, IEEE (2007) 211–216
7. Israsena, P.: Securing ubiquitous and low-cost RFID using tiny encryption algorithm. In: International Symposium on Wireless Pervasive Computing, Phuket, Thailand, IEEE (2006)
8. Juels, A.: Minimalist cryptography for low-cost RFID tags. In Blundo, C., Cimato, S., eds.: The Fourth International Conference on Security in Communication Networks – SCN 2004. Volume 3352 of LNCS., Springer-Verlag (2004) 149–164
9. Karlof, C., Sastry, N., Wagner, D.: TinySec: A link layer security architecture for wireless sensor networks. In: Second ACM Conference on Embedded Networked Sensor Systems (SenSys 2004). (2004) 162–175
10. Kwon, D., Han, D., Lee, J., Yeom, Y.: Vulnerability of an RFID authentication protocol proposed at SecUbiq 2005. In: International Workshop on Security in Ubiquitous Computing Systems – Secubiq 2006. LNCS, Seoul, Korea, Springer-Verlag (2006)
11. Li, T., Wang, G.: Security analysis of two ultra-lightweight RFID authentication protocols. In: IFIP SEC 2007, Sandton, Gauteng, South Africa, IFIP (2007)
12. Microelectronic, E.: EM 4035 datasheet (2006)
13. Moteiv Corporation: Telos (rev b): Datasheet (2004)
14. Peris-Lopez, P., Hernandez-Castro, J.C., Estevez-Tapiador, J., Ribagorda, A.: LMAP: A real lightweight mutual authentication protocol for low-cost RFID tags. In: Workshop on RFID Security 2006(RFIDSec 06), Graz, Austria, Ecrypt (2006)
15. Peris-Lopez, P., Hernandez-Castro, J.C., Estevez-Tapiador, J., Ribagorda, A.: M2AP: A minimalist mutual-authentication protocol for low-cost RFID tags. In: International Conference on Ubiquitous Intelligence and Computing – UIC06. Volume 4159 of LNCS., Springer-Verlag (2006) 912–923
16. Peris-Lopez, P., Hernandez-Castro, J.C., Estevez-Tapiador, J.M., Ribagorda, A.: EMAP: An efficient mutual authentication protocol for low-cost RFID tags. In: OTM Federated Conferences and Workshop: IS Workshop – IS’06. Volume 4277 of LNCS., Springer-Verlag (2006) 352–361
17. Perrig, A., Szewczyk, R., Wen, V., Culler, D.E., Tygar, J.D.: SPINS: security protocols for sensor networks. In: Mobile Computing and Networking (2001) 189–199
18. Piramuthu, S.: HB and related lightweight authentication protocols for secure RFID tag/reader authentication. In: Collaborative Electronic Commerce Technology and Research – COLLECTeR 2006, Basel, Switzerland (2006)
19. Poschmann, A., Leander, G., Schramm, K., Paar, C.: DESL: An efficient block cipher for lightweight cryptosystems. In: Workshop on RFID Security 2006(RFIDSec 06), Graz, Austria, Ecrypt (2006)
20. Ranasinghe, D.C., Lim, D., Cole, P.H., Devadas, S.: White paper: A low cost solution to authentication in passive RFID systems. Technical Report WP-HARDWARE-029, Auto-ID Labs, The University of Adelaide, Adelaide, Australia (2006)
21. Ransford, B., Clark, S., Salajegheh, M., Fu, K.: Getting things done on computational RFIDs with energy-aware checkpointing and voltage-aware scheduling. In: USENIX Workshop on Power Aware Computing and Systems (HotPower). (2008)
22. Rieback, M., Gaydadjiev, G., Crispo, B., Hofman, R., Tanenbaum, A.: A platform for rfid security and privacy administration. In: USENIX/SAGE Large Installation System Administration conference – LISA’06, Washington DC, USA (2006)
23. Rivest, R.: The RC5 encryption algorithm. In Preneel, B., ed.: FSE. Volume 1008 of LNCS., Springer (1995) <http://theory.lcs.mit.edu/~rivest/Rivest-rc5rev.pdf>.
24. Salajegheh, M., Clark, S., Ransford, B., Fu, K., Juels, A.: CCCP: Secure remote storage for computational RFIDs. In: Proceedings of the 18th USENIX Security Symposium, Montreal, Canada (2009)
25. Sample, A.P., Yeager, D.J., Powledge, P.S., Smith, J.R.: Design of a passively-powered, programmable platform for UHF RFID systems. In: IEEE International Conference on RFID 2007 (2007)

26. Smith, J.R., Sample, A.P., Powledge, P.S., Roy, S., Mamishev, A.: A wirelessly-powered platform for sensing and computation. In: 8th International Conference on Ubiquitous Computing (UbiComp 2006), Orange Country, CA, USA (2006) 495–506
27. Spotlight, R.: 5 things must know for RFID starters. WWW (2006) Blog entry, Last viewed May 15, 2007 <http://www.innovez-one.com/blogs/>.
28. ThomasNet: Short range UHF EPC tag is designed for item-level tagging. WWW (2007) Last viewed May 12, 2007 <http://news.thomasnet.com/fullstory/482149>.
29. Westhues, J.: Proxmark3. WWW (2007) Last viewed May 16, 2007 <http://cq.cx/proxmark3.pl>.
30. Yeager, D., Zhang, F., Zarrasvand, A., Otis, B.P.: A 9.2ua Gen 2 compatible UHF RFID sensing tag with -12dBm sensitivity and 1.25uVrms input-referred noise floor. In: IEEE International Solid-State Circuits Conference (ISSCC). (2010)

Security Enhanced WISPs: Implementation Challenges

Alexander Szekely, Michael Höfler, Robert Stögbuchner,
and Manfred Aigner

1 Introduction

The wireless identification and sensing platform (WISP) is an interesting platform for a variety of new applications. The combination of sensors, a programmable microcontroller, a contact-less communication interface, and energy harvesting from the reader field makes it much more than a “contact-less sensor.” The tag communicates via a standardized RFID protocol with off-the-shelf UHF RFID readers; it harvests the energy for its operation from the electromagnetic field that is supplied by the reader. The sensor is a maintenance-free device, since it does not require an onboard power source that can run out nor can interfaces or cabling get worn due to mechanical burden. The attached microcontroller allows for on-tag processing, storage or filtering of the measured data before it is transferred to the reader. The RF channel does not require a direct line of sight, supports continuous access to multiple tags in the field, and operates over a range of more than one meter.

Beside the major advantages described above, the contact-less interface has also drawbacks. Eavesdropping of the communication between the sensor and the reader is much simpler than the interception of wire-based communication. Depending on the application, the measured data can be interesting for competitors or other attackers. In case that a WISP measures data related to persons, e.g., their body temperature, serious end-user privacy concerns may arise.

The operator of the measurement system should be aware that the maximal eavesdropping distance may be much higher than the operation range of a WISP. An attacker can eavesdrop communication between WISP and reader from further away, since the attacker’s receiver does not need to power the tag with its field.

A. Szekely (✉) • M. Höfler • R. Stögbuchner • M. Aigner
Institute for Applied Information Processing and Communications (IAIK),
Graz University of Technology, Inffeldgasse 16a, 8010 Graz, Austria
e-mail: Alexander.Szekely@iaik.tugraz.at; Michael.Hoefler@student.tugraz.at;
Stoegbuchner@student.tugraz.at; Manfred.Aigner@iaik.tugraz.at

Beside the obvious scenario of an eavesdropper intercepting the communication between tag and original reader, one should consider that an attacker may connect to a WISP when it is not operated by its legitimate owner. An attacker may pass by with her own reading device whenever possible. Depending on the application, such cases of rouge measurements may be unpleasant to highly tenuous.

As long as the WISPs measure and operate in a closed environment, these obvious security threats can be solved by protecting the area from unwanted access by potential attackers. The area may be shielded so that susceptible information does not leak. In scenarios where the sensors are placed outside protected areas, access control and shielding are impossible. This is for instance the case when WISPs are used to monitor the temperature profile of perishable goods in the supply chain. Protection of the measured and communicated data on basis of cryptographic mechanisms is necessary in such applications. One can even consider cases where non-trusted readers may transfer collected sensor data from remote sensors to an application server. In those scenarios, the data originating from the sensor has to be protected against illicit modification on its way via the untrusted reader and the public network to the application server.

When cryptographic mechanisms are applied, it is good practice to rely on publicly open algorithms following Kerckhoffs' principle [6]. For protection of the communication between WISPs and readers, we suggest to use the advanced encryption standard (AES) [10] which was standardized by NIST in 2001 after an extensive public evaluation. AES is a modern block cipher that allows efficient implementation in software as well as in hardware. Feldhofer et al. have shown in 2005 that AES can be implemented as hardware module fulfilling the demanding requirements of passive RFID tags [5]. Dominikus et al. suggest how the cryptographic functionality of the tag can be integrated into the tag-to-reader communication protocol [1]. Their approach considers that passive low-cost tags do not provide a programmable microprocessor, that would allow implementation of cryptographic algorithms in software. In difference to low-cost RFID tags, WISPs feature a programmable 16-bit microprocessor.

This chapter presents first results how AES encryption can be integrated into the WISP firmware so that measured sensor data can be transferred encrypted to the reader. We show that it is possible to integrate an optimized AES implementation into the firmware without major reduction of the operating distance. We think that the approach is meaningful for many future applications that will require private data transmission. Our results show that communication between WISPs and readers can be protected with state-of-the-art cryptographic mechanisms running on the low-performance controller of the sensors.

Whenever cryptographic devices can get into the hands of potential attackers, their vulnerability to so called implementation attacks has to be considered. Attackers may measure power, electromagnetic (EM) radiation, or similar to reveal the secrets stored on the devices, for instance, the cryptographic key. A WISP operating in an untrusted environment is certainly susceptible to this threat. We present first investigations in this direction. We evaluated the feasibility of a side-channel analysis (SCA) attacks [9] on an AES-equipped WISP. By measuring the

EM radiation of a WISP during multiple encryption operations, we were able to reveal its secret key. The platform is therefore susceptible to this kind of attacks and special countermeasures need to be considered for secure applications.

This chapter is organized as follows. In Sect. 2, we shortly present the WISP platform as used in our work. After that, we provide some background on AES in Sect. 3 and present the details about how AES was integrated in Sect. 4. Additionally, we present the resulting performance figures for the AES enhanced WISP firmware. In Sect. 5, we give a short introduction into SCA attacks and show our results of the first tests on the vulnerability against SCA attacks. In Sect. 6 we present the conclusions drawn from our results.

2 The WISP Platform

The WISP is a programmable passive RFID tag including sensing and computing devices. Powered by the UHF reader, the WISP harvests the energy from the emitted radio signals. The included sensing devices are a light sensor, a temperature sensor, and a 3D acceleration sensor. The onboard microcontroller, a Texas Instruments MSP430F2132 [12] ultra-low-power 16-bit microcontroller, is used for all computing tasks, sampling the sensors, and controlling the communication to the reader. The controller provides 8 kB of flash memory, 512 Bytes of RAM, and a 10-bit analog to digital converter.

The WISP does not actively transmit data, but it communicates with the UHF reader via backscatter modulation. Therefore, the tag modulates the impedance of its antenna to change the amount of reflected energy.

The firmware on the WISP is responsible for the measurement of the sensors, the power management, and also implements the EPC Class-1 Generation-2 UHF RFID protocol [3]. It is written in a mix of C and assembly language, where the assembly code mostly maintains the EPC protocol to meet the timing constraints.

Managing the power consumption of the microcontroller and peripherals on a passive RFID tag is a challenge. The energy requirements of the MSP430 and the sensors are significantly larger than those of typical passive RFID tags. If the power harvester cannot provide enough energy, the WISP enters a sleep mode where it stops all computation but retains its RAM content. In this mode the device can harvest energy for several reader cycles. As soon as the energy harvester accumulated enough energy, the external voltage supervisor issues a hardware interrupt to wake up the processor.

To transmit sensor data to the reader, the WISP supports two types of operation. It can either transmit the data as a response to the read command or it can transmit the data in the tag ID. Transmitting the data in the ID involves less communication overhead and gives better range and performance results. Figure 1 shows the standard EPC Class-1 Generation-2 ID reply. In the *transmit-data-in-ID* mode, the sensor data is encoded as shown in Fig. 2. The 96-bit ID field is divided into four logical units. The first field, the 8-bit tag type, indicates the sensor type which



Fig. 1 Structure of a 96-bit ID tag reply according to the EPC Class-1 Generation-2 UHF RFID protocol



Fig. 2 Structure of the 96-bit ID used to transmit sensor data within the ID

Table 1 WISP tag types associated with the sensor devices

Tag type	Sensor
0x00	Static ID
0x0C	Null sensor
0x0B	Acceleration (quick)
0x0D	Acceleration (standard)
0x0E	Temperature (external sensor)
0x0F	Temperature (internal sensor)

generated the data. A list of possible tag types can be found in Table 1. The next eight bytes are used to transmit the actual sensor data. Depending on the active sensor type, two to six bytes of the data field are used to encode the measured data. The remaining two bytes are used as a running counter. The data field is followed by the 8-bit WISP hardware version and a 16-bit WISP serial number. As the ID changes with the measured data, one WISP appears as several different tags to a standard EPC reader.

3 Advanced Encryption Standard

The AES [10] is a symmetric-key encryption algorithm. It supports key lengths of 128 bits, 192 bits, or 256 bits and operates on fixed-sized blocks of 128 bits. The versions of the algorithms with different key lengths differ only in the number of rounds and the key schedule. On the WISP, we implemented the encryption with a key size of 128 bits (AES-128).

In AES, the data, also called the state, and the key are represented in two rectangular arrays arranged as shown in Fig. 3. Each square represents eight bits of plain text or key data.

The AES algorithm encrypts each data block by applying the same round function multiple times to the state. In case of AES-128, the round transformation is iterated ten times. The *KeySchedule* function generates a new round key for each of the rounds.

Fig. 3 Arrangement of the AES state and key bytes

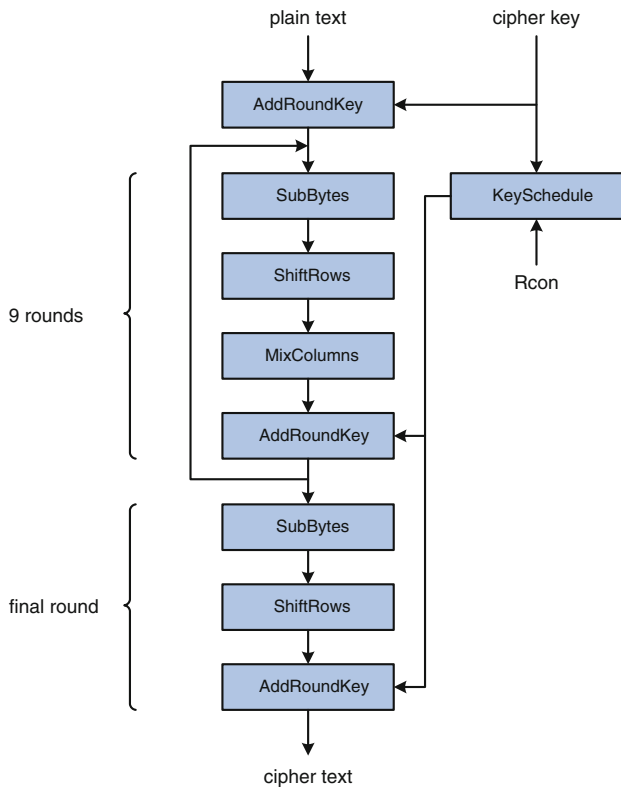



Fig. 4 Structure of the AES encryption rounds

As shown in Fig. 4, the round function consists of the four basic steps *SubBytes*, *ShiftRows*, *MixColumns*, and *AddRoundKey*. Each of these operations modifies the state array.

The *SubBytes* operation is the only nonlinear step in the round transformation. Each byte in the state is replaced with the associated value from the S-box. Although

it is possible to calculate the values of the S-box from its definition, it is usually implemented with the help of a lookup table.

ShiftRows rotates the rows of the state to the left. The first row is left unchanged, the second row is rotated by one position, the third row by two positions, and the last row rotates three positions to the left.

MixColumns operates on the columns of the state. It is defined as a matrix multiplication by a constant polynomial in a finite field modulo $x^4 + 1$.

The *AddRoundKey* operation adds the 128-bit round key to the 128-bit state *mod* 2. This is equivalent to a bitwise XOR operation of the round key and the state.

The *KeySchedule* function generates the 128-bit round keys from the cipher key. It is based on the S-box function and the round constants *Rcon*. Note that the key for the initial key addition is the unmodified cipher key.

4 Enhancing the WISP with Cryptography

In this section, we show how state-of-the-art cryptography can be added to the WISP. We extended the *transmit-data-in-ID* mode so that all sensor data is encrypted with AES. To provide security and privacy to the WISP, changes to the protocol, the firmware, and the reader application were necessary. However, all our changes are backward-compatible to non-modified version of the WISP protocol, so that security-enhanced and original WISPs can coexist.

The original WISP firmware implements the EPC Class-1 Generation-2 protocol. In the *transmit-data-in-ID* mode, the whole communication between the WISP and the reader is handled in an inventory round. After receiving the *Select* and *Query* commands, the tag answers with a 16-bit random number (RN16). Once the tag receives the *ACK* command from the reader, it backscatters the measured sensor data in the EPC ID. The firmware of the WISP is implemented in a state machine, resembling the EPC Gen2 tag states. The sampling and processing of the sensor data is done in the state *STATE_READ_SENSOR*.

4.1 Protocol Modifications

As shown in Sect. 2, the sensor data is encoded in a 64-bit data field. Since AES only operates on 128-bit blocks, the 64-bit sensor data has to be extended to 128 bit. We do this by padding the original sensor data with eight bytes containing 0xFF. The padded sensor data is then encrypted with AES-128. However, also the output of the encryption is 128 bits wide and cannot be accommodated in the 64-bit data field of the ID reply message. We therefore split the 128-bit cipher block into two 64-bit packets (packet0 and packet1). Figure 5 illustrates this procedure.

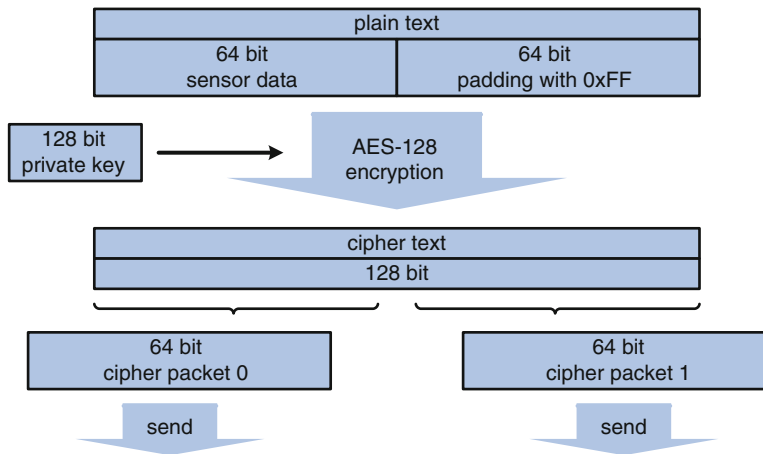


Fig. 5 Encrypting the sensor data and splitting it into two packets for transmission

To assure compatibility with original WISPs, we define new tag types to distinguish between unencrypted and encrypted ID replies. We leave to lowest four bits of the tag type unchanged and set the upper four bits of the tag type to 0×2 for the first encrypted packet and to 0×3 for the second encrypted packet. In that way, encrypted WISP messages can easily be detected, as all unencrypted tag types have the upper four bits set to 0×0 (see Table 1).


We used the AES in the electronic code book mode (ECB) [2]. It is the simplest mode of operation with the disadvantage of identical plain text blocks being encrypted into identical cipher-text blocks. However, this is not a problem in our setting. As described in Sect. 2, each measurement is stamped with a counter. So even if the measured data is identical, the input for the encryption function is always different.

4.2 Changes in the WISP Firmware

On the WISP tags we used a highly optimized assembler implementation of the AES provided by SIC.¹ To avoid fundamental changes of the firmware and the state flow, the encryption of the sensor data is directly performed within the state *STATE_READ_SENSOR*. Normally, this state gathers the sensor data by calling the subroutine of the activated sensor device. If encryption is enabled, the state *STATE_READ_SENSOR* has now to distinguish if new data has to be sampled from the sensor, or if we have to send the second part of our cipher text (packet1).

¹Stiftung Secure Information and Communication Technologies (SIC): AES Implementation for MSP430 Microcontroller. Available online at <http://jce.iaik.tugraz.at/sic/Products/Crypto-Software-for-Microcontrollers/Texas-Instruments-MSP430-Microcontrollers>

Fig. 6 Cipher packet0, cipher packet1, and the plain text packet of encrypted temperature data



If we need to sample the sensor, the measured sensor data is immediately encrypted. Afterwards, the first part of the cipher (packet0) is written to the ID reply buffer, and the second part is stored in a temporary buffer. If we enter the state *STATE_READ_SENSOR* and have a packet1 waiting in the buffer, we simply transfer the saved packet1 into the ID reply buffer and do not sample the sensor.

The firmware enters the state *STATE_READ_SENSOR* after a predefined number of ten state transitions, independently from the number of ID reply transmissions. However, our modified version samples the sensors only on every second entry. To achieve the same sensor sampling rate as the unmodified firmware, we decreased the number of state transitions before we enter the state *STATE_READ_SENSOR* to five.

4.3 Changes on the Application Side

If a WISP tag is within the reader's field and replies to the ID request of the reader, report packets with the replied information are sent to the demo application. More specifically, first the reader puts the tag reply into a report object, which is transferred through the low-level reader protocol (LLRP [4]) to the application. The application uses the *RFIDReader* object from the LLRP library to receive the report packets. It then parses the information into a *MyTag* object. The objects generated in the *RFIDReader* are passed to the *ReaderManager* and further on to the *TagHandler* where the tag information is handled depending on the contents of the tag-type field. For each tag type, the *TagHandler* provides a separate method which is called to parse and store the information of the *MyTag* object. The application then accesses this information through a getter function to display the sensor measurements.

To handle the encrypted sensor data additional functionality is required in the *TagHandler*. This is because one encrypted measurement of the WISP tag is transmitted using two ID replies. We added two new reply parsing methods for the encrypted tag types $0 \times 2^*$ (encrypted packet0) and $0 \times 3^*$ (encrypted packet1). When we receive a packet0, we simply buffer the data block. When we receive a packet1, we first check if we have previously received a packet0. If not, the corresponding packet0 got somehow lost and we discard the received packet. Otherwise, we combine the two packets and decrypt the data. Figure 6 shows an example of how two cipher packets are transformed into an unencrypted temperature data packet.

Table 2 The read rate for different distances between the antenna and the tag, for unencrypted and encrypted communication

	Update rate	Distance = 50 cm		Distance = 100 cm		Distance = 200 cm	
		Read rate	Percent	Read rate	Percent	Read rate	Percent
Without AES	10	37.45	100%	21.24	100%	9.80	100%
With AES	5	30.62	81.8%	16.95	79.8%	7.85	80.1%
Decrypted data	5	14.12	37.7%	7.27	34.2%	2.87	29.3%
With AES	10	31.87	85.1%	18.15	85.5%	8.24	84.1%
Decrypted data	10	13.21	35.3%	6.96	32.8%	2.66	27.1%

Decrypted data is the net data rate for the correctly reassembled and decrypted packets on the reader

Since it is possible that packets are lost during the transmission, it is not guaranteed that the correct packets are combined. Therefore, we check if the 0xFF padding is present after the encryption. In other words, the padding of the measured data acts as an integrity check. If the padding is intact, the first 64 bits of the plain text are valid sensor data. We handle this data equally to a non-encrypted ID reply: We put the decrypted data into a new *MyTag* object and pass it to the *TagHandler*, where it is processed like an unencrypted packet. If the 0xFF signature is not correct, we have combined two wrong packets and both have to be discarded.

4.4 Results

The security-enhanced firmware has a code size of 5,574 bytes, including the original WISP firmware, the AES-128 encryption routine, and the modifications to the firmware. The encryption code needs 83 additional bytes of RAM to save registers and store temporary values. The encryption of one sensor sample is performed in 5,432 clock cycles. We could not notice any differences in the maximal reading distance for unencrypted and encrypted transmissions. However, there are differences in the read rate. To show these differences between the transmission of unencrypted and encrypted sensor data, we measured the throughput for both cases at a distance of 50 cm, 100 cm, and 200 cm between the reader antenna and the tag. As sensing device we used the internal temperature sensor of the WISP. The read rate indicates the number of received tag replies per second at the reader. Table 2 presents the measured read rates for unencrypted and encrypted communications, as well as the net read rate for correctly recombined and decrypted packets. The difference in the reading rate between the transmission of plain data and AES-encrypted data is in total about 15%. All measurements were performed with the original sensor update-rate of ten and with our modified update-rate of five state transitions.

Due to the fact that two tag replies are necessary to transmit one encrypted sensor sample the reading rate is reduced by the 50% in theory. Because of lost packets,

and therefore wrongly combined and discarded packets, the effective read rate for decrypted sensor data is slightly lower than the theoretic value and depends on the distance. However, the main reason for the reduced performance with the AES encryption lies in the fact that we need two tag replies to transmit one encrypted sensor sample. This bottleneck could be resolved by transmitting the whole cipher text in a single ID reply. For that, the length of the ID reply must be increased. The EPC Class-1 Generation-2 UHF RFID protocol specification permits ID reply lengths of up to 496 bits, whereas the current firmware supports only 96 bits. Future WISP implementations should take the typical block size of cryptographic primitives into account and thus could avoid the transmission of multiple packets. With these changes an encrypted, secure communication can be achieved with an overhead of just about 15%.

5 Side-Channel Attacks on the WISP

Side-channel analysis attacks are an efficient method to successfully attack implementations of mathematically strong cryptographic algorithms. In the following, we present a short overview of SCA and explain how side-channel attacks work in practice. After that, we present a successful side-channel attack on the AES encryption on the WISP tag.

Traditional cryptoanalysis assumes that an attacker has only information about the input and the output of a cryptographic device. In contrast to that, in SCA, an attacker also exploits physical information leaked from the cryptographic device. Examples for such side channels are the power dissipation, timing information, or the electromagnetic emanation (EM).

The concept of power analysis attacks was first introduced by Kocher et al. [7]. In simple power analysis (SPA) the attacker tries to reveal the secret directly from a single power trace. This is typically possible in implementations that use key dependent branches.

Differential power analysis (DPA) attacks are more powerful than SPA attacks. Instead of visually inspecting a single trace, DPA uses several traces and applies statistic methods to reveal the secret key. The attack operates on intermediate results of the cryptographic algorithm. An attackable intermediate result must depend on a known data value and a small part of the key (sub-key). The known data is usually the plain text or the cipher text. The attacker then calculates the intermediate values for every possible value of the sub-key. These intermediate values are then mapped to hypothetical power-consumption values, by assuming a power model of the attacked device. Finally, these hypothetical power-consumption values are compared with the measured power or EM values from the real device. The most common way to compare these values is the use of the correlation coefficient. The correlation coefficient determines the linear relationship between the measured data and the calculated, hypothetical values. The highest correlation reveals the correct guess of the sub-key.

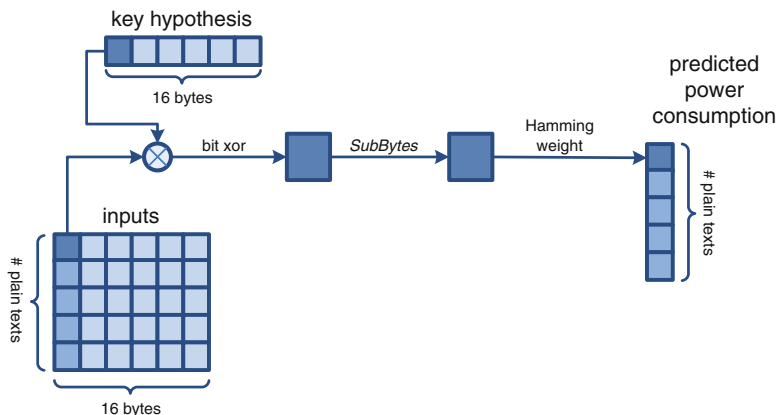


Fig. 7 Generation of the hypotheses matrix for the SubBytes operation in the first AES round

5.1 Differential Power and EM Attack on AES

To attack our AES implementation on the WISP, we use the output of the first SubBytes operation in the first round as the intermediate result. This intermediate result depends on the plain text (i.e., the sensor data) and the first byte of the secret key. As described in Sect. 5, we need to compute every possible intermediate value. In order to create this so-called hypotheses matrix, the known plain text has to be manipulated the same way as the AES algorithm on the attacked device manipulates the plain text. This means we have to process a hypothetical AES encryption with all 256 possible values of the sub-key and all provided plain texts up to the first SubBytes operation. In addition, we also need to record the power consumption or EM radiation while the WISP encrypts these provided plain texts.

Figure 7 depicts the generation of the hypotheses matrix. The matrix holds the output of the SubBytes operation of all possible sub-key values bitwise XORed with the known plain text. Before we can correlate the hypotheses matrix with the measured power traces we need to apply a power model, which is suitable for the MSP430, to the calculated values. For our attack, we have chosen the Hamming-weight model. This model assumes that the power consumption of the MSP430 is proportional to the number of bits in the processed value that are set to one. This model is often a good guess if no details about the internals of the attacked device are known.

The last step in the attack is the calculation of the correlation matrix. The size of the correlation matrix depends on the number of key-byte hypotheses in the hypotheses matrix and the number of recorded measurement points. The higher the correlation coefficient, the better the measured data and the hypothetic power consumption match. Therefore, the highest value in the correlation matrix reveals the correct sub-key. Further details can be found in [9].

5.2 Measurement Setup

This section describes the measurement setup that has been used to perform the DPA attack on the WISP. To facilitate the power and EM measurements, the attack was not performed within the reader field. Instead, the WISP has been actively powered. In addition, some changes to the WISP firmware alleviate the attack. The finite state machine in the WISP firmware was modified so that it operates in a single state. This state performs the AES encryption on predefined plain texts instead of real sensor data. In every encryption cycle, immediately before the first execution of *SubBytes* in the first round, the firmware asserts a pin (P3.1) on port three of the MSP430. This pin is used as a reliable trigger signal for the measurement on the digital storage oscilloscope. After the completion of *SubBytes*, the software resets the trigger pin. In addition, we disabled all sensors and the communication with the reader.

This setup does not provide a realistic attack scenario. However, the modifications allow us to quickly show that the platform is susceptible to this kind of attack. In addition, the setup is a first step towards more realistic attacks. The changes to the firmware are not a requirement for a successful attack as they only reduce the noise on the measurements and ease the finding of a usable trigger position. The susceptibility of passive RFID tags to EM analysis attacks, even if they are powered from the reader, has been shown in [11].

5.2.1 Current Consumption

The optimal place to measure the current consumption of the microcontroller is directly in the power-supply path of the MSP430. However, to measure at this position, we would have to desolder the chip from the board. In the noninvasive attack we accomplished, this procedure was not performed. The power consumption during the AES encryption on the WISP tag should be dominated by the microcontroller, as it is the only active device on the circuit board at the time of the measurement. Therefore, we put a $33\ \Omega$ resistor in the power-supply path of the WISP and measured the voltage drop over that resistor using a digital storage oscilloscope. To avoid smoothing of the power consumption by capacitors, we removed the capacitors between the supply branch and ground (C14 and C20).

5.2.2 Electromagnetic Emanation

For the EM attack no hardware changes were required. As for the attack on the current consumption, we used the active power supply to power the WISP. This setup allows to record the EM side-channel caused by the AES encryption without any influence of the reader's field. We placed a magnetic near-field probe at the point, where we could detect the highest amount electromagnetic emanations from




Fig. 8 Three recorded power traces

MSP430. This spot was at the corner of the MSP430 casing right between the pins 24 and 25. The best results were measured with the probe LF-B 3 [8], which detects electromagnetic emanations at a frequency from 100 kHz up to 50 MHz. Furthermore, we amplified the sensed signal using a 30 dB amplifier.

5.3 Results of the Differential Power Analysis Attack

We measured the voltage drop over the $33\ \Omega$ resistor in the supply branch during the first *SubBytes* operation of the first round. In total, we recorded the power traces of 600 encryptions. The traces of three different measurements are presented in Fig. 8. In the active state, the MSP430 has a current consumption below $100\ \mu\text{A}$. Because of this small power consumption, the signal-to-noise ratio of the measurements is critically low.

Nevertheless, there are a few interesting facts to notice in the power traces. The amplitude of the power consumption reaches the highest value at approximately $10\ \mu\text{s}$. At this position the trigger pin is set to high in the firmware. This change of the port state causes the maximal current drain during the whole encryption. In addition, there are several peaks at constant intervals in the power trace. A closer look at the position of the trigger signal reveals that the MSP430 shows the same power-consumption pattern in every trace. However, this pattern vanishes in noise after some clock cycles, as the signal is too weak for our equipment.

Because of the small signal-to-noise ratio, our efforts trying to align the different traces were unrewarded. We were not able to extract the correct key from the recorded power traces.




Fig. 9 EM emanation side-channel trace




Fig. 10 Misaligned electromagnetic emanation traces at a clock edge

5.4 Results of the EM Attack

With the EM near-field probe LF-B 3 [8], we measured the encryption of 200 different plain texts. We recorded the electromagnetic emanation during the encryption with a constant secret key.

In Fig. 9, a snippet of a recorded EM trace is shown. Every peak in this diagram between the nearly uniform parts is caused by a clock edge. As we can see in Fig. 10, these traces are also misaligned, similar to the current consumption traces. However, in this case, we have a sufficient high signal-to-noise ratio and can therefore align the traces in a preprocessing step.

In Fig. 11a, the correlation for the correct key-byte hypothesis has a peak at about 35 μ s. In contrast to that, Fig. 11b shows the correlation for a wrong key hypothesis. The correct key hypothesis is clearly distinguishable from the wrong hypothesis. To reveal the whole secret key, this process was repeated for every remaining sub-key. Thus, we successfully revealed the 128 bit secret key.




Fig. 11 Correlation between key hypotheses and aligned measurement traces. (a) Correct key hypothesis. (b) Wrong key hypothesis

6 Conclusions

In this chapter, we presented security and privacy enhancements for the WISP platform. Especially in untrusted environments these enhancements would allow for a new class of applications.

Our results show that standardized state-of-the-art encryption algorithms like AES can be implemented on the WISP without a noticeable reduction of its operation radius. However, the current protocol was not designed with cryptography in mind and therefore unnecessarily reduces the throughput of encrypted communication. Newer WISP protocol versions should take cryptography into account, as this could enable secure applications without considerable performance penalties.

We also presented first results on SCA attacks on actively powered WISP tags. Our results clearly show the vulnerability of the WISP against these attacks and that countermeasures against SCA need to be included when cryptography is added to the platform.

In future work, we will investigate on more realistic SCA attacks, where the WISPs are passively powered from field of the reader, and provide suggestions for countermeasures against SCA attacks.

References

1. S. Dominikus, E. Oswald, and M. Feldhofer. Symmetric authentication for RFID systems in practice. In *Workshop on RFID and Lightweight Crypto*, July 13–15, 2005, Graz, Austria, 2005.
2. M. Dworkin. Recommendation for Block Cipher Modes of Operation - Methods and Techniques. Technical report, National Institute of Standards and Technology (NIST), December 2001.
3. EPCglobal. EPC Radio-Frequency Identity Protocols Class-1 Generation-2 UHF RFID Protocol for Communications at 860 MHz–960 MHz Version 1.0.9, January 2005. Available online at http://www.epcglobalinc.org/standards/uhfclg2/uhfclg2_1.2_0-standard-20080511.pdf.

4. EPC Global Inc. Low level reader protocol (llrp). http://www.epcglobalinc.org/standards/llrp/llrp_1_0_1-standard-20070813.pdf, August 2007.
5. M. Feldhofer, J. Wolkerstorfer, and V. Rijmen. AES Implementation on a Grain of Sand. *IEE Proceedings on Information Security*, 152(1):13–20, October 2005.
6. A. Kerckhoffs. La cryptographie militaire. *Journal des sciences militaires*, IX:5–38, January 1883.
7. P.C. Kocher, J. Jaffe, and B. Jun. Differential Power Analysis. In Michael Wiener, (ed.) *Advances in Cryptology - CRYPTO '99, 19th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15–19, 1999, Proceedings*, vol. 1666 of *Lecture Notes in Computer Science*, pp. 388–397. Springer, 1999.
8. LANGER EMV-Technik GmbH. H-field probe lf-b 3. <http://www.langer-emv.de/en/products/disturbance-emission/near-field-probes/lf-1/>. Accessed 26 Dec 2012.
9. S. Mangard, E. Oswald, and T. Popp. *Power Analysis Attacks—Revealing the Secrets of Smart Cards*. Springer, 2007. ISBN 978-0-387-30857-9.
10. National Institute of Standards and Technology (NIST). FIPS-197: Advanced Encryption Standard, November 2001. Available online at <http://www.itl.nist.gov/fipspubs/>.
11. T. Plos. Susceptibility of UHF RFID Tags to Electromagnetic Analysis. In Tal Malkin, (ed.) *Topics in Cryptology - CT-RSA 2008, The Cryptographers' Track at the RSA Conference 2008, San Francisco, CA, USA, April 8–11, 2008, Proceedings*, vol. 4964 of *Lecture Notes in Computer Science*, pp. 288–300. Springer, April 2008.
12. Texas Instruments Inc. MSP430F21x2 MIXED SIGNAL MICROCONTROLLER. Available at <http://focus.ti.com/lit/ds/symlink/msp430f2132.pdf>, December 2009.

Part V
Wireless Power Beyond RFID

Power Optimized Waveforms that Enhance the Range of Energy-Harvesting Sensors

Matthew S. Trotter and Gregory D. Durgin

1 The Energy-Harvesting Environment of Passive Wireless Sensor Networks

Passive tags have no onboard power source and require energy-harvesting circuitry to absorb power from the reader transmissions. However, reader-transmitted power is limited, thereby limiting the range and reliability of the tags' response. The maximum equivalent isotropically radiated power (EIRP) from an unlicensed user is limited by governing agencies in different areas of the world:

- In the USA, the FCC limits reader EIRP to 4 W for UHF and microwave unlicensed intentional radiators [5]
- In the EU, the ETSI limits reader EIRP to 2 W for UHF and microwave unlicensed intentional radiators [4]

thus limiting the amount of power available to a passive tag at a given range. In addition, a passive tag's received power suffers from a number of other barriers [8, 13]:

1. Path loss: Power transmitted to the tag in free space falls by $1/\text{distance}^2$.
2. Path blockages: A non-line-of-sight channel forces a tag to harvest from scattered reflections.
3. Multipath fading: Destructive interference at the tag location results in received power fading.

M.S. Trotter (✉)

Department of ECE, Georgia Tech, 777 Atlantic Avenue NW, Atlanta, GA 30332, USA

e-mail: durgin@gatech.edu

G.D. Durgin

Disney Research, Pittsburgh, PA, USA

e-mail: mtrotter@disneyresearch.com

4. Tag impedance mismatch: A high-Q impedance match is difficult to achieve over large bandwidths.
5. Tag antenna pattern: Omnidirectional tag antennas are inherently unfocused resulting in poor gain.
6. Tag chip sensitivity: The amount of power necessary for tag chip operation helps determine the minimum required received power.
7. Tag detuning: Placing the tag on conductive, high-permittivity, or high-permeability materials can reduce tag antenna gain.
8. Tag diode threshold: Tag diodes limit the voltage amplitude available to the charge pump.

A tag must harvest the available energy efficiently for these substantial barriers severely limit the potential power available.

2 Previous Methods of Improving Range and Tag Sensitivity

Several researchers have proposed changing the basic continuous wave (CW) carrier to enhance performance of RF energy-harvesting circuits, mostly by experimenting with cycling between two high and low transmit power levels. For example, Greene et al. describe transmission of a square-wave-modulated CW to some of its passive energy-harvesting devices [7]. In similar research, Matsumoto and Takei improved the efficiency of charge pumps by using intermittent CW transmissions switched on and off periodically in a square-wave fashion [12]. This reduced the tag's power consumption by 25%. A 900 MHz CW transmission was turned on and off at rates of 1 kHz, 100 kHz, and 10 MHz. These intermittent CW waveforms are similar to square power optimized waveforms (POWs, discussed in Sect. 4.2).

In a technique that explores multi-carrier gains similar to POW transmissions, Park et al. found a doubling of the read range for passive RFID tags by using multiple CW transmitters [15]. A separate, auxiliary reader is triggered by a query command from the interrogating reader. The passive tag receives the CW transmitted from different directions. Another benefit is that multipath fading is reduced in their approach.

Other researchers have concentrated on new designs for RF tag energy-harvesting circuitry to enhance the range of passive devices. For example, Oliver et al. describe a charge pump architecture (implemented in the ImpinjTM Monza 3 RFID tag) that uses two separate charge pumps to overcome some voltage limitations [14]. A small charge pump provides a small bias voltage for the transistors (instead of diodes) of a large charge pump, which powers the chip. This design lowers the required threshold voltage of the large charge pump's transistors, which increases power efficiency.

POWs provide range increases and tag sensitivity improvements without changing tag architecture. In addition, POWs may be used as a carrier for data transmission and reception.

3 Charge Pump Circuit Models in Passive-Tag Energy Harvesting

The most common energy-harvesting circuit used in passive tags is the charge pump [2] shown in Fig. 1 for its simplicity, inexpensiveness, and high power efficiency. The charge pump simultaneously rectifies the received RF energy and boosts the resulting DC voltage to a usable level for the tag chip to operate. The charge pump’s power efficiency directly controls the tag’s receive power sensitivity and is dependent on the input RF signal amplitude. In fact, the charge pump operates most efficiently when input RF signal magnitude is much larger than diode threshold voltage [16].

3.1 Input/Output Relationship

The charge pumps shown in Fig. 1 are based on the Dickson charge pump model [1]. The input/output equation is adapted from that work to match the stage-counting convention of Fig. 1:

$$V_{DC} = \frac{N(\max(V_{RF}) - V_t)}{1 + \frac{N}{fR_L C}} u(\max(V_{RF}) - V_t), \tag{1}$$

where $u(\cdot)$ is the unit-step function. The output only exists when the input’s maximum voltage is greater than the diode threshold voltage. The charge pump rectifies the input RF signal; thus it is dependent on the maximum value of the input voltage. Efficient operation occurs when the input voltage greatly exceeds the diode threshold voltage, V_t .




Fig. 1 Energy-harvesting charge pumps for passive tags, comprised of diodes and capacitors, are low cost and efficient when the input amplitude is much larger than the diode threshold ($\max(V_{RF}) \gg V_t$) (Figure used from [17] with permission from IEEE)

The charge pump fails to produce an output voltage when the maximum of the input voltage is smaller than the diode threshold. The capacitors in Fig. 1 are not presented a voltage to charge to when the connecting diode is not forward biased by the input signal. Therefore, the charge pump's diodes set the limit on the necessary power a passive tag needs to turn on. This voltage limitation sets today's passive UHF RFID tags' power sensitivity at between -20 and -15 dBm [9–11].

4 Power Optimized Waveforms


The charge pump is a highly nonlinear device whose performance depends greatly on the shape of the exciting RF waveform. In fact, conventional CW excitation turns out to be one of the most inefficient signals for exciting a charge pump. A POW transmitted by the reader helps reduce the voltage limitation, in many cases pushing the range of operation for voltage-limited RF tags an order of magnitude beyond their CW range.

4.1 The POW Concept

POWs improve the range and reliability of voltage-limited passive tags by increasing the tag's charge pump efficiency over its normal efficiency with a continuous waveform (CW) input. A POW is any reader-transmitted waveform with large voltage peaks that maintain a specified signal power according to local transmitting rules. Figure 2 shows a reader transmitting CW and POW both at the same transmit power. The POW's voltage peaks are larger than the CW's. These large voltage peaks dominate the charge pump diodes' threshold voltage, thereby increasing efficiency according to Tanzawa et al. [16]. Thus, a POW basically refocuses the available average transmit power into short, impactful bursts of power to which the charge pump is more sensitive.

POW Effectiveness. Peak to average power ratio (PAPR) is a convenient metric that formulates the POW's power-focusing ability. It also allows comparisons between POWs. The PAPR of signal $y(t)$ is defined as

Fig. 2 A reader transmitting RF energy to power up a passive tag first in the form of CW and second in the form of POW. Both waveforms have the same signal energy, yet the POW's peaks are larger (Figure used from [18] with permission from IEEE)






Fig. 3 Passive tags with envelope detection to demodulate reader commands will have larger ripple with POW than CW, which can limit tag range despite receiving plenty of power (Figure used from [18] with permission from IEEE)

$$PAPR = \frac{\max(y(t)y^*(t))}{E[y(t)y^*(t)]}, \tag{2}$$

where $(\cdot)^*$ denotes the complex conjugate of the signal and $E[\cdot]$ is the expected value operator. POWs with large peaks have a higher PAPR than POWs with small peaks and greatly improve range and reliability.

POW Spectrum. The transmit spectrum of a POW is more spread out than a comparable CW of the same transmit power. A POW’s spectrum is the spectrum of its envelope upconverted to the passband. Since the POW envelope is nonconstant and contains voltage peaks and valleys, the POW’s spectrum is spread out from the center frequency. This is a limitation of POWs in systems that have tight spectral masks, but a POW can be customized to fit within the spectral mask by sacrificing either large voltage peaks or short POW time period.

Envelope Detection. Passive tags, including the Intel wireless identification and sensing platform (WISP), employ envelope detection to demodulate amplitude-modulated reader commands. The discharging time constant of the envelope detector is designed on the order of a CW time period rather than a POW time period. These envelope detectors may not reliably demodulate a POW with a large time period.

An envelope detector is a rectifier attached to the antenna, often in series with or parallel to the charge pump if they share a common antenna, as shown in Fig. 3. Its capacitor is sized to reliably detect symbols modulated on CW with little ripple but may not perform well when detecting a POW. The ripple magnitude is proportional to the POW’s period, voltage peak, and envelope detector’s capacitor size. Without sufficient envelope detection, a tag may not receive the reader’s commands even if there is enough received power. Certain POWs discussed in the next section,

Sect. 4.2, have better envelope detection characteristics than others. In fact, POWs with long, smooth peaks rather than quick, steep peaks are read most reliably.

4.2 Basic POW Shapes

Many types of POWs can be created as there is a broad definition for a POW. As long as the waveform has large voltage peaks and maintains a specified signal power, it can be considered a POW. Some convenient example POWs are rigorously defined in this section. For each POW, the time-domain waveform, peak power, average power, and PAPR are defined in the passband based on the POW's basic parameters.

M-POW. This straightforward POW consists of M equally spaced and equally powered baseband subcarriers that are then upconverted to the passband. Its passband time-domain equation is

$$\text{POW}(t) = \frac{1}{\sqrt{M}} \sum_{k=1}^M \cos\left(2\pi \frac{kt}{T_{\text{POW}}}\right) \cdot \cos(2\pi f_c t) \quad (V). \quad (3)$$

Here, M is the number of subcarriers in the baseband, which turns into $2M$ subcarriers centered around the center frequency in the passband. The POW's peak occurs when all cosines reach their maximum simultaneously at the beginning of each POW period, T_{POW} . Its passband peak signal power, average signal power, and PAPR are:

$$\text{Peak signal power} = M \quad (V^2), \quad (4)$$

$$\text{Average signal power} = \frac{1}{4} \quad (V^2), \quad (5)$$

$$\text{PAPR} = 4M. \quad (6)$$

Since the M-POW definition is normalized, the average signal power is constant (with units V^2). There is a trade-off between increasing PAPR and maintaining a small bandwidth.

Gaussian POW. Systems that seek to optimize a small bandwidth and large peak power should use the Gaussian POW. Its passband envelope is shaped like a Gaussian function (or normal distribution). The time-domain waveform in the passband is defined as

$$\text{POW}(t) = \frac{1}{\sqrt{2\pi\sigma^2}} \sum_{k=-\infty}^{\infty} e^{-\frac{(t-kT_{\text{POW}})^2}{2\sigma^2}} \cdot \cos(2\pi f_c t) \quad (V), \quad (7)$$

where σ is the standard deviation, T_{POW} is the POW period (in s), and f_c is the center frequency of the passband. The peak signal power, average signal power, and PAPR in the passband are:

$$\text{Peak signal power} = \frac{1}{2\pi\sigma^2} (V^2), \quad (8)$$

$$\text{Average signal power} \approx \frac{1}{4\sqrt{\pi}\sigma T_{\text{POW}}} (V^2), \quad (9)$$

$$\text{PAPR} \approx \frac{2T_{\text{POW}}}{\sqrt{\pi}\sigma}. \quad (10)$$

These equations assume the small-tail approximation where the Gaussian pulse tail adds a negligible amount to the previous period's or subsequent period's Gaussian pulse. Also, the POW time period is assumed to be longer than the standard deviation ($T_{\text{POW}} > 3\sigma$).

Square POW. Systems that operate on a large bandwidth may use a square POW since its spectrum is very wide compared to other POWs. The waveform is defined by

$$\text{POW}(t) = \left[B + \sum_{k=-\infty}^{\infty} (A - B) \text{rect} \left(\frac{t - kT_{\text{POW}}}{DT_{\text{POW}}} \right) \right] \cdot \cos(2\pi f_c t) (V), \quad (11)$$

where A is the high voltage and B is the low voltage (both in V). D is the duty cycle ($= \text{on time}/\text{off time}$, $0 \leq D \leq 1.0$, unitless). The peak signal power, average signal power, and PAPR in the passband are:

$$\text{Peak signal power} = A^2 (V^2), \quad (12)$$

$$\text{Average signal power} = \frac{1}{2} (A^2 D + B^2 (1 - D)) (V^2), \quad (13)$$

$$\text{PAPR} = \frac{2}{D + (B^2/A^2)(1 - D)}. \quad (14)$$

Average signal power in the passband is the geometric average of the rectangles $A^2 D$ and $B^2 (1 - D)$. Increasing duty cycle reduces PAPR, while increasing the high voltage, A increases PAPR. The FCC requires that a data transmission does not include any time period of "zero" transmission [5]. Thus, the square POW must be nonzero for all time. Figure 4 shows a 4-POW, Gaussian POW, and square POW compared side by side as a visual aid.

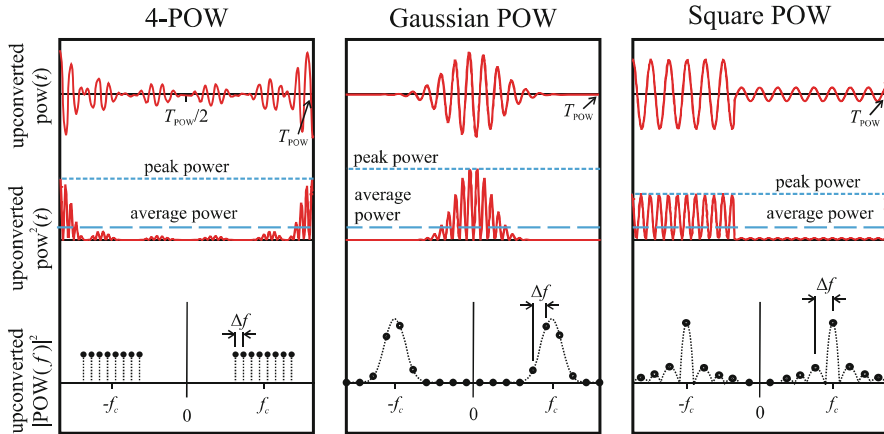


Fig. 4 Graphical comparison of signal, power, and spectrum of the 4-POW, Gaussian POW, and square POW (Figure used from [18] with permission from IEEE)

5 The Optimal POW

In general, POWs have trade-offs between PAPR, envelope detection reliability, and bandwidth. Increasing the peak voltage while holding average power constant increases PAPR, yet reduces envelope detection reliability since the envelope detector’s capacitor must charge and discharge to large voltage swings. Simultaneously, bandwidth increases since the voltage must swing faster to larger voltages than a low-PAPR POW, thus introducing higher-frequency components of the signal. Conversely, a POW with a smoother peak and smaller bandwidth that is easily detectable will necessarily have a smaller PAPR. This limits the POW gain experienced by the tag.

There is the well-known trade-off between bandwidth and RMS time width as well. A POW with large bandwidth has a shorter RMS time width than a POW with a small bandwidth. This trade-off is apparent from the well-known uncertainty principle [6], which states

$$T_{\text{RMS}}^2 \cdot B_{\text{RMS}}^2 \geq \frac{1}{16\pi^2}, \tag{15}$$

where T_{RMS} is the RMS width in the time domain, and B_{RMS} is the RMS bandwidth in the frequency domain. The equality in Eq. (15) only holds for Gaussian functions.

From the uncertainty principle, the optimal POW that maximizes POW gain for a given RMS width is the Gaussian POW. The RMS width is defined as the normalized second central moment of the POW time-domain signal. Likewise, the RMS bandwidth is defined as the normalized second central moment of the POW spectrum:

$$T_{\text{RMS}}^2 = \frac{\int_{-\infty}^{\infty} t^2 |\text{POW}(t)|^2 dt}{\int_{-\infty}^{\infty} |\text{POW}(t)|^2 dt}, \quad (16)$$

$$B_{\text{RMS}}^2 = \frac{\int_{-\infty}^{\infty} f^2 |\text{POW}(f)|^2 df}{\int_{-\infty}^{\infty} |\text{POW}(f)|^2 df}. \quad (17)$$

Both the RMS width and RMS bandwidth are normalized to signal energy. The denominators of both of these equations are equal according to Parseval's theorem. For the Gaussian function, the RMS width is commonly known as the variance.

6 Theoretical Power Gains

Charge pumps are nonlinear devices; thus they do not relate themselves well within linear circuit models. To get around this obstacle, an energy-delivery approach is used here to derive the power gain of using POW over using conventional CW as the power-transmission waveform. The *POW gain* is defined as the ratio of power delivered by a POW to power delivered by CW:

$$G_{\text{POW}} = \frac{P_{\text{DC,POW}}}{P_{\text{DC,CW}}}. \quad (18)$$

Compare two equal-power signals at the input to a charge pump:

$$V_{\text{CW}}(t) = \sqrt{2}V_{\text{RMS}} \cos(2\pi f_c t), \quad (19)$$

$$V_{\text{POW}}(t) = \sum_{i=-\infty}^{\infty} p(t - iT_{\text{POW}}) \cos(2\pi f_c t), \quad (20)$$

where $p(t)$ is the finite-energy baseband POW pulse with RMS width, T_{RMS} , and RMS bandwidth, B_{RMS} . The average power in Eq. (20) is effectively

$$P_{\text{avg}} = \frac{1}{R_{\text{in}}} \int_{-\infty}^{\infty} p^2(t) dt, \quad (21)$$

where R_{in} is the time-averaged input resistance to the charge pump.

Approximate $p(t)$ as a square pulse as shown in Fig. 5. The average square peak voltage, V_{SP} , is a good approximation to the average voltage level of the actual pulse during the energizing portion of the waveform; outside this interval, the POW's low energy levels are considered to be inconsequential. This square-interval approximation has the added benefit of being constant over the interval $2T_{\text{RMS}}$; thus we can use the charge pump output voltage Eq. (1).




Fig. 5 Approximating an arbitrary POW pulse as a square pulse with equal signal energy

For diodes with threshold voltage, V_t , the energy delivered to the charge pump load during one POW period, T_{POW} , is

$$\begin{aligned}
 E_{\text{CW}} &= \frac{V_{\text{DC,CW}}^2}{R_L} T_{\text{POW}} \\
 &= \frac{N^2 \left(V_{\text{RMS}} \sqrt{2} - V_t \right)^2 u \left(2V_{\text{RMS}}^2 - V_t^2 \right)}{R_L \left(1 + \frac{N}{f_c R_L C} \right)^2} T_{\text{POW}}, \quad (22)
 \end{aligned}$$

$$\begin{aligned}
 E_{\text{POW}} &= \frac{V_{\text{DC,POW}}^2}{R_L} (2T_{\text{RMS}}) \\
 &\approx \frac{N^2 \left(V_{\text{RMS}} \sqrt{\frac{T_{\text{POW}}}{T_{\text{RMS}}}} - V_t \right)^2 u \left(V_{\text{RMS}}^2 \frac{T_{\text{POW}}}{T_{\text{RMS}}} - V_t^2 \right)}{R_L \left(1 + \frac{N}{f_c R_L C} \right)^2} (2T_{\text{RMS}}). \quad (23)
 \end{aligned}$$

Energy is constantly delivered for the CW transmission because the input is always on. However, the POW delivers a large amount of energy for a shorter period of time, $2T_{\text{RMS}}$. The POW gain can equivalently be expressed as the ratio of energy delivered by a POW to energy delivered by CW:

$$\begin{aligned}
 G_{\text{POW}} &= \frac{E_{\text{POW}}}{E_{\text{CW}}} \\
 &\approx \frac{\left(\sqrt{P_{\text{in}}} - \sqrt{\frac{P_t}{\text{PAPR}}} \right)^2 u \left(P_{\text{in}} - \frac{P_t}{\text{PAPR}} \right)}{\left(\sqrt{P_{\text{in}}} - \sqrt{P_t} \right)^2 u \left(P_{\text{in}} - P_t \right)}, \quad (24)
 \end{aligned}$$




Fig. 6 POW gain across four regions of input power. Region I is the “below threshold Region.” Region II is the “max gain region.” Region III is the “common region.” Lastly, region IV is the “reverse breakdown region”

where the approximate relationship, $PAPR \approx T_{POW}/2T_{RMS}$, is used. Also, the input signal power is P_{in} , and P_t is the signal power required to forward bias the diodes.

Equation (24) illustrates four regions of POW gain vs. input power as shown in Fig. 6:

- Region I, “below threshold region”: Neither POW nor CW can forward bias the diodes since there is not enough input power. The device remains unpowered for both POW and CW.
- Region II, “max gain region”: The POW forward biases the diodes, but CW does not. POW gain is theoretically infinite in this region.
- Region III, “common region”: POW gain approaches 1.0 as the POW and CW both efficiently operate the charge pump when input power reaches very high power.
- Region IV, “reverse breakdown region”: Reverse-breakdown diode behavior begins to limit the POW’s effectiveness, but not CW’s. Note that a very large $PAPR \left(\geq \frac{V_{RB}^2}{V_t^2} \right)$ can completely collapse region III. Equation (24) does not model this behavior as the equation assumes the ideal diode model.

Equation (24) also shows that, for fixed RMS bandwidth, POW gain is maximized if and only if $p(t)$ is a Gaussian POW as a result of the uncertainty principle.

7 Tested Power Gains on WISP and Current Passive RFID Tags

The WISP implements key aspects of the EPCglobal, class 1, generation 2 standard [3] (also known as the “gen2” standard) for passive RFID tags [19], which operates in the UHF band from 860 to 960 MHz. A WISP tag’s energy harvester is a charge pump very much like a passive RFID tag’s and like the examples shown in Fig. 1. Therefore, it is a convenient platform on which to test POWs.

Table 1 Measured POW gains on WISP with various POWs

Power transmission signal	PAPR	Minimum TX power (dBm) ^a	POW gain (dB)
Gen2 CW	2.0	5.6	–
1-POW ^b	4.0	4.7	+0.9
2-POW	8.0	4.6	+1.0
3-POW	12.0	4.0	+1.6
4-POW	16.0	3.8	+1.8
Square POW	4.8	4.1	+1.5
Sharp Gaussian POW ^c	6.0	No response	–
Wide Gaussian POW ^d	3.0	9.0	–3.4

^aAntenna to WISP separation distance was 26.5 cm

^bPOW period $T_{\text{POW}} = 333.33$ ns for all POWs except Gaussian POWs

^c $T_{\text{POW}} = 846.67$ ns and $\sigma = 159.16$ ns

^d $T_{\text{POW}} = 846.67$ ns and $\sigma = 318.31$ ns

The objective of the test is to compare the range-increasing effectiveness of various POWs including N-POW, Gaussian POW, and square POW. This test is run in two different methods.

7.1 WISP Tag Sensitivity Measurement

The first method measures WISP tag sensitivity at a fixed range (26.5 cm in this case) from the transmitting antenna. Tag sensitivity is defined as the minimum power the WISP tag must receive in order to power up and backscatter a response. Equivalently, the minimum transmitted power required to power up the WISP tag can also be measured and is measured for this test method. First, the transmitter sends a gen2 *query* command modulated on a single-tone CW, and the minimum transmit power is recorded. Then, the same gen2 query command is modulated on each POW and the corresponding minimum transmit power recorded. The difference between the POW's minimum transmit power and the CW's minimum transmit power is recorded as the POW gain:

$$G_{\text{POW}} = \text{CW min. TX power} - \text{POW min. TX power.} \quad (25)$$

A positive value of G_{POW} means the tag is more sensitive to POW than CW thus it implies an improvement of range for equal transmit power. Conversely, a negative value of G_{POW} represents a reduction of range. The transmission remains on until the WISP responds reliably to every consecutive query command with either CW or POW as the carrier. The results are tabulated in Table 1.

POW gain increases when moving from 1-POW to 4-POW. This is explained by the increase in PAPR from the 1-POW's PAPR of 4 to the 4-POW's PAPR of 16. The sharp gaussian POW, which has a Gaussian envelope and a short variance, did not illicit a response from the WISP. Also, the wide gaussian POW exhibited a




Fig. 7 Graph of POW gain vs. transmitted power illustrating the max gain region and common region. All POWs except 1-POW and 2-POW start to decrease into the common region after a transmit power of -1 dBm

negative POW gain implying a reduction in range. However, in previous sensitivity measurements of passive RFID tags, the Gaussian POWs performed the best out of all POWs tested [18]. The square POW improved tag sensitivity by 1.5 dB, which is in between the performance of the 2-POW and 3-POW. The square POW is generally the most reliable POW since its peak lasts for a longer amount of time than the other POWs.

7.2 POW Gain Versus Transmit Power Measurement

According to the POW gain Eq. (24) in Sect. 6, POW gain is dependent on input power to the charge pump and operates in four regions. To measure this effect, transmit power is varied between -13 and $+10$ dBm to find the boundaries of the four regions.

The WISP tag is placed at a fixed range of 34 cm, and a voltmeter is connected to the DC output port of the charge pump to measure the DC output voltage. A settling time of 20 s allows the charge pump to reach steady state before a measurement is taken. A base test of transmitting CW to the WISP over the transmit power range gives the baseline reference. Then, each POW is measured over the transmit power range, and the POW gain is found by taking the ratio of DC output powers and converting to log scale:

$$G_{\text{POW}} = 20 \log_{10} \left(\frac{V_{O,\text{POW}}}{V_{O,\text{CW}}} \right). \tag{26}$$

The plot in Fig. 7 shows the measured results. Each POW is plotted on the same graph to illustrate that, in general, POW gain is higher for POWs with large PAPR

than for POWs with low PAPR. Notice that the 1-POW line is mostly lower than the rest of the POWs since it has the smallest PAPR.

The boundary between the max gain region and the common region is -1 dBm for most POWs shown. The 1-POW and 2-POW lines continue flat and modestly increase after transmit power of -1 dBm, which suggests that it is about to enter the common region.

In Fig. 7, the sharp Gaussian and wide Gaussian dominate the performance of the other POWs. This measurement is insensitive to envelope detection characteristics and therefore shows that the WISP tag did not accurately detect the reader command when Gaussian POW was used as the transmission waveform in the first measurement method (Table 1).

8 Backscatter Communications with POW as a Signal Carrier

The backscattered signal received by the reader consists of the tag’s backscatter modulated on the passband POW. Figure 8 shows a receiver chain that demodulates the tag’s backscatter. The receiver first downconverts the received signal to a baseband signal, which is sent through a matched filter with an impulse response of the time-reversed POW. The matched filter output is then sampled and sent to the symbol decision device.

An alternative implementation is subcarrier filtering and detection. Here, the POW carrying backscattered data is first downconverted to baseband. Then, each baseband subcarrier of the POW contains the backscattered data. All but one of these baseband subcarriers is filtered out, and the backscattered data signal is sent to the symbol decision device.

The matched filter method is preferred for POWs with long time periods or if the reader contains a high-speed analog-to-digital converter (ADC) which can adequately sample the baseband received signal. In this case, digital filtering is preferred over hardware-based filtering. Subcarrier filtering is preferred when the POW has a short time period and therefore large subcarrier spacing. Low-Q and inexpensive low-pass filters can eliminate all baseband POW subcarriers but the backscattered data.




Fig. 8 Matched filter demodulator for demodulating passive RFID tags’ backscatter with a POW carrier (Figure used from [18] with permission from IEEE)

9 Conclusions

A wireless sensor network based on current WISP nodes that transmits POWs can increase range on the order of 1 to 2 dB. POWs help absolve the voltage limitation of charge pump diodes, which limit the maximum range of passive tags. Example POWs such as the N-POW, Gaussian POW, and square POW were shown to provide extra DC power at the output of the charge pump on a WISP tag.

Challenges still remain for implementing POW successfully including bandwidth limitations of various wireless radio protocols such as the EPCglobal “gen2” standard [3] and ZigBee [20]. POW may be used as a location-finding technique for passive wireless sensor networks. The POW can be used similarly to a radar system. Further, a new medium-access control scheme can be invented if proven accurate at finding tag location.

The WISP tested in this work can experience *longer* range improvement by redesigning its energy harvester, although it is not necessary. In general, any passive tag can implement these design goals to maximize POW’s range improvement. First, the envelope detector’s time constant should be increased to handle the POW’s long time period rather than CW’s time period, thus ensuring reliable signal detection. Second, the voltage regulation circuitry should allow higher input voltages to accommodate large-PAPR POWs. Third, the tag antenna and matching section’s passband should match the POW’s passband spectrum, which preserves the POW’s shape and does not suppress POW subcarriers.

Implementing POW also affords a couple of design luxuries. The charge pump can be designed with fewer stages to achieve the same DC voltage, which saves space and cost. Higher-threshold diodes, which are typically less expensive than low-threshold diodes, may be used as well to achieve the same DC voltage.

References

1. Dickson, J. F.: On-Chip High-Voltage Generation in MNOS Integrated Circuits Using an Improved Voltage Multiplier Technique. *IEEE Journal of Solid-State Circuits*, 11, 374–378 (1976)
2. Dobkin, D. M.: *The RF in RFID: Passive UHF RFID in Practice*. Elsevier, Boston (2008)
3. EPCglobal: EPC (TM) Radio-Frequency Identity Protocols Class-1 Generation-2 UHF RFID Protocol for Communications at 860 MHz–960 MHz Version 1.2.0. <http://www.epcglobalinc.org/home/>. Cited July 2010
4. European Telecommunications Standards Institute: European Telecommunications Standards Institute Radio Frequency Identification Equipment operating in the band 865 MHz to 868 MHz with power levels up to 2 W. EN 302 308 <http://www.etsi.org/WebSite/technologies/RFID.aspx>. Cited July 2010
5. Federal Communications Commission: FCC Rules and Regulations Part 15 Section 247 (15.247), Operation within the bands 902–928 MHz, 2400–2483.5 MHz, and 5725–5850 MHz. <http://www.gpoaccess.gov/>. Cited July 2010
6. Gasquet C., Witomski P.: *Fourier Analysis and Applications: Filtering, Numerical Computation, Wavelets*. Springer, New York (1999)

7. Greene, C. E., Shearer, J. G., Harrist, D. W.: Pulse Transmission Method. US Patent Application Publication. Publication Number US 2007/0149162 A1. June 28, 2007
8. Griffin, J. D., Durgin, G. D.: Complete Link Budgets for Backscatter Radio and RFID Systems. *IEEE Antennas and Propagation Magazine*, 51, 11–25 (2009)
9. HiggsTM-3 EPC Class 1 Gen 2 RFID Tag IC. Alien Technology Corporation (2008)
10. Monza 3TM Product Brief. Impinj, Inc (2008)
11. Monza 4[®] Tag Chip Datasheet. Impinj, Inc (2010)
12. Matsumoto, H., Takei, K.: An Experimental Study of Passive UHF RFID System with Longer Communication Range. *Proceedings of Asia-Pacific Microwave Conference*, 1–4 (2007)
13. Nikitin, P., Rao, K. V. S.: Performance Limitations of Passive UHF RFID Systems. *Proceedings of the IEEE Antenna and Propagation Society International Symposium*, 1101–1014 (2006)
14. Oliver, R. A., Diorio, C. J.: RFID Tags with Power Rectifiers that have Bias. US Patent 7 561 866. September 26, 2005
15. Park, J. S., Jung, J. W., Ahn, S. Y., Roh, H. H., Oh, H. R., Seong, Y. R., Lee Y. D, Choi, K.: Extending the Interrogation Range of a Passive UHF RFID System With an External Continuous Wave Transmitter. *IEEE Transactions on Instrumentation and Measurement*, Accepted for future publication (2010)
16. Tanzawa, T., Tanaka, T.: A Dynamic Analysis of The Dickson Charge Pump Circuit. *IEEE Journal of Solid-State Circuits*, 32, 1231–1240 (1997)
17. Trotter, M. S.: Effect of DC to DC Converters on Organic Solar Cell Arrays for Powering DC Loads. Georgia Institute of Technology. (2009)
18. Trotter, M. S., Durgin, G. D.: Survey of Range Improvement of Commercial RFID Tags With Power Optimized Waveforms. *Proceedings of the 2010 IEEE International Conference on RFID*, 195–202 (2010)
19. The WISP Wiki. <http://wisp.wikispaces.com/>. Cited July 2010
20. Zigbee Alliance: ZigBee RF4CE Specification. <http://www.zigbee.org/Products/DownloadZigBeeTechnicalDocuments.aspx>. Cited July 2010

Wireless Ambient Radio Power

Alanson P. Sample, Aaron N. Parks, Scott Southwood,
and Joshua R. Smith

1 WARP TV Harvester

Early work on wireless power transfer using far-field, propagating electromagnetic waves transferred energy the same way radios transmit signals. In the 1960s, Brown pioneered the field of microwave power transmission using rectennas (or rectifying antennas) [2]. Demonstrations using high gain transmitters and rectennas for receiving and rectifying have shown wireless power transfer of 30 kW over a distance of one mile, with a total end-to-end efficiency of 84% [3]. Today, UHF RFID systems wirelessly power RFID tags, which have no batteries and an operating range of ~ 10 m [1, 5].

However, when considering the application of wirelessly powering and recharging distributed sensor nodes, these approaches have several drawbacks. First, both of the examples require a dedicated transmitter to energize a given region of space. Secondly, these systems suffer from the need for sophisticated tracking and alignment equipment to maintain a line-of-sight (point to point) connection in unstructured and dynamic environments.

An alternative approach is to harvest the ambient RF energy that permeates our environment from sources such as commercial broadcast TV and radio, as well as

A.P. Sample • J.R. Smith (✉)

Department of Computer Science and Engineering and Department of Electrical Engineering, University of Washington, Seattle, WA, USA
e-mail: alanson@u.washington.com; jrs@cs.washington.edu

A.N. Parks

Department of Electrical Engineering, University of Washington, Seattle, WA, USA
e-mail: anparks@uw.edu

S. Southwood

Department of Computer Science and Engineering, University of Washington, Seattle, WA, USA
e-mail: scott.southwood@gmail.com




Fig. 1 Position “A” indicates the location of the KING-TV tower which transmits 960 kW. The contour rings estimate the available receive power for the Seattle region (assuming line of sight and a 5 dBi receive antenna)

cellular and Wi-Fi transmitters. Ambient RF energy harvesting begins with a source, such as a TV transmitter, where power is continuously transmitted in the form of data signals. Generally speaking, the RF energy is radiated in a radially symmetric pattern, with power intensity decreasing as it travels from the source.

Consider the example depicted in Fig. 1, which shows a TV tower located at point “A.” The radiation pattern of the transmitter results in power contour lines that fall off as the reference point (or receiver) moves away from the transmitter. If, for example, an energy-harvesting device with an activation threshold of $100\mu W$ is brought towards the TV tower there will be a particular distance at which the device will begin to operate. Thus, the power requirements of the load application defines the maximum distance of operation. This distance represents a radius of operation around the TV tower which encircles an area, wherein the device will continually harvest energy for operation. Initial research on the wireless ambient radio power (WARP) project has demonstrated the feasibility of developing such a device, capable of having an operating area of several hundred square miles around the transmission tower [4].

Figure 2 shows an image of the WARP TV harvester, consisting of a commercially available digital TV antenna and an RF rectifier. To demonstrate the harvesting functionality, a RadioShack indoor/outdoor thermometer with hygrometer (Cat. No.



Fig. 2 Demonstration of a temperature and humidity meter (including LCD display) that is exclusively powered from ambient RF signals transmitted by the TV tower seen in the background, at a distance of 4.1 km

63-1032) is connected to the output of the rectifier as an example application. From the balcony at Intel Labs Seattle ($47^{\circ} 39' 41''$ N, $122^{\circ} 18' 60''$ W), I harvested RF power from the KING-TV tower ($47^{\circ} 37' 55''$ N, $122^{\circ} 20' 59''$ W), which broadcasts 960 kW EIRP on channel 48 at 674–680 MHz. The total distance is 4.1 km. The antenna used is a 5 dBi log periodic antenna designed for TV applications. The power harvester is a 4-stage, voltage-multiplying rectifier, the design of which is similar to that of the one used on the WISP. The RF front end is tuned to the desired channel and has a bandwidth of approximately 30 MHz. With the antenna manually oriented towards the transmit tower, the measured open circuit voltage at the rectifier output was 5.0 V (i.e., the only load on the power harvester was the voltmeter). Next, an $8\text{ k}\Omega$ load was attached to the rectifier output, resulting in an output voltage of 0.7 V, which corresponds to $60\text{ }\mu\text{W}$ of power harvested. This is equivalent to the net power budget of many of the WISP sensing applications.

Applying the Friis transmission formula with the parameters above, yields an expected power received of $220\text{ }\mu\text{W}$. Thus, the experimentally measured performance of the system is reasonably close to the theoretically expected performance. We then connected this ambient RF-harvesting system to the battery terminals of the RadioShack indoor/outdoor temperature and humidity meter (thermometer/hygrometer) with an LCD display. This device is normally powered by a 1.5 V AAA battery. The thermometer/hygrometer was measured to consume around $25\text{ }\mu\text{A}$, at 1.5 V, from a laboratory power supply. Approximately once per second, the current consumption briefly spiked up to around $50\text{ }\mu\text{A}$, presumably when the sensor measurements were made.

The thermometer/hygrometer functioned normally when connected to the power-harvesting circuit with the antenna oriented at the appropriate transmission tower; the display contrast appeared to be as good as when the system was powered by a battery. With the antenna oriented directly at the TV tower and the thermometer/hygrometer connected and operational, the loaded voltage was measured to be 1.7 V. As the antenna was oriented away from the tower to which it was tuned, the display contrast dropped and then, when the antenna was further mis-oriented, appeared to stop operating altogether.

2 WARP Sensor Node

The previous section showed the ability to harvest ambient RF energy and power a dedicated sensor and display. Although this example aptly demonstrates the possibility of RF energy harvesting, the application itself is somewhat limited in its scope. The next logical extension of ambient RF energy harvesting is to use this capability to enable new applications and usage models. One compelling area of application is wireless sensor nodes, which have computational, sensing, and bi-directional communication capabilities.

For as long as there have been wireless sensors, researchers have searched for ways to sustainably power them in remote locations. Today, there are numerous remote sensor applications deployed across a wide range of disciplines. These include environmental monitoring, such as volcano or hurricane sensors, to industrial monitoring of strain, vibration, consumable parts, and machine degradation. Solar power is generally viewed as the most usable source of ambient energy, and when present, it provides a viable source of power for remote sensors. However, solar-powered systems suffer from the limitation that they must be placed in well-lit areas and must have line of sight to the light source. Additionally, since there is no sunlight at night, operation is often quite restricted. Batteries can mitigate this restriction and allow devices to continue operating for a short time after ambient light is removed. However, dependency on batteries can be troublesome, and in many cases, it is too costly or time consuming to replace each node after the battery is empty.

Furthermore, we must consider the total system complexity of solar-powered sensor nodes. Typically, there is an integrated circuit for computation and sensing, an antenna for communication, a solar cell for power, and a battery for storage. In contrast, a passive RFID tag is simply made up of an antenna and an integrated circuit. If we extend the concept of a passive RFID tag to include energy harvesting from ambient RF sources, instead of a deliberate source such as an RFID reader, we can see that this system is inherently less complex. Although ambient RF power is not as abundant as solar power, the low manufacturing complexity and potentially uninterrupted power supply of a WARP-type solution will make it viable for a number of application scenarios.

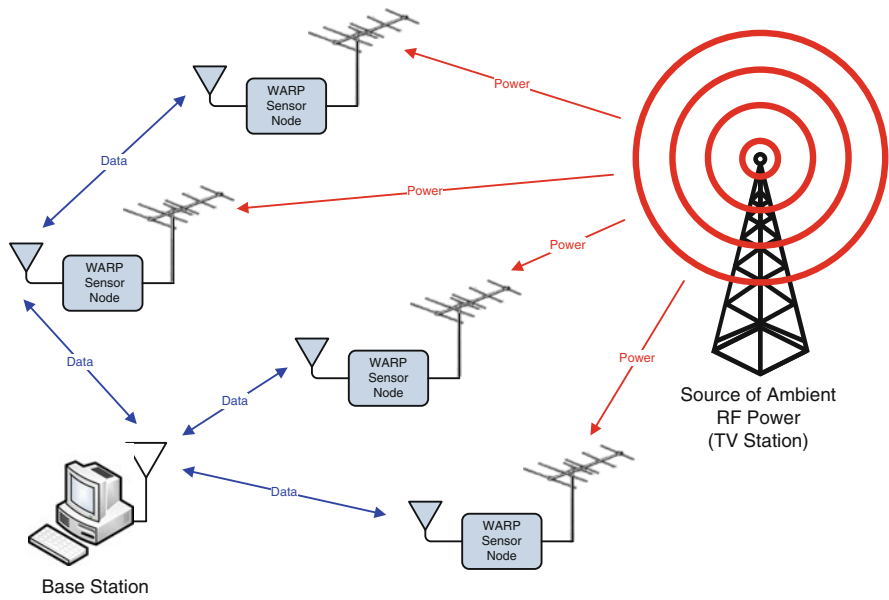


Fig. 3 Conceptual diagram of a deployment of WARP sensor nodes. The WARP nodes harvest ambient RF power from TV stations for operation and communicate back to a base station using an onboard radio

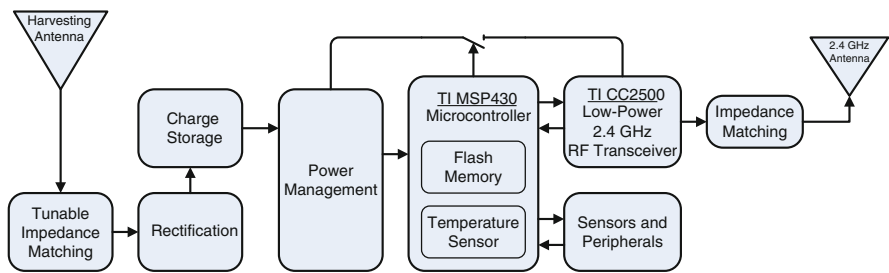


Fig. 4 Block diagram of the proposed WARP sensor node with 2.4 GHz transceiver

The WARP project aims to create wireless sensor nodes that are completely powered by harvested RF energy. In order to accomplish this goal, the existing WARP energy harvesters were extended by adding the capability to sense the environment, perform computation, and communicate wirelessly. The first target application is a weather-monitoring node. Figure 3 shows a block diagram of the system, which receives RF power from a TV transmitter and wirelessly reports that data back to a host computer.

Figure 4 shows a block diagram of the sensor node architecture. The harvesting antenna is connected with an SMA connector so that different frequency ranges and antenna gains can be easily tested. Received RF energy is fed to the rectifier via a tunable impedance matching network. Rectified energy is gathered over time

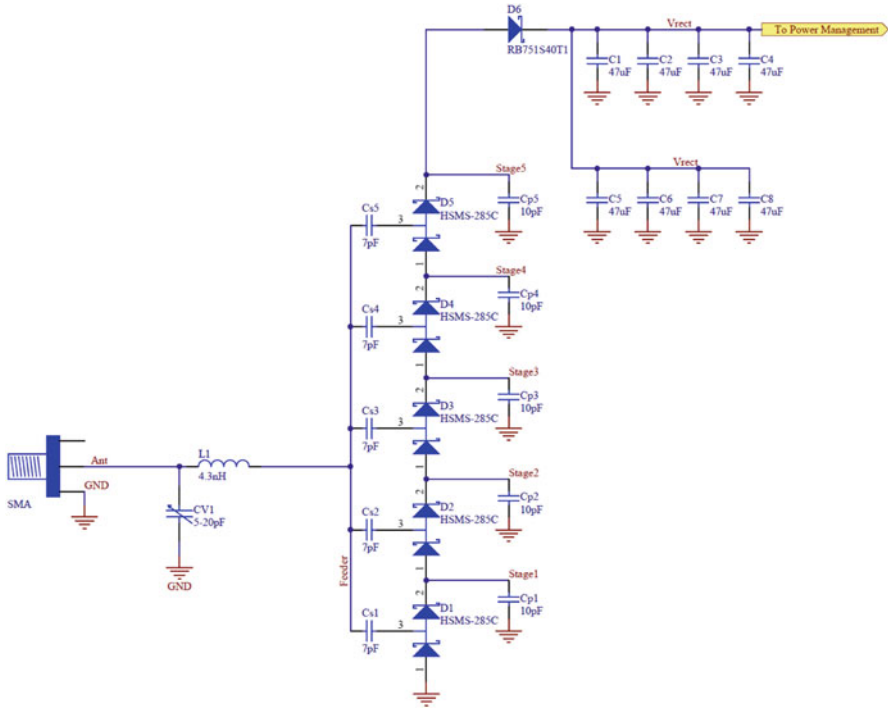


Fig. 5 RF power-harvesting circuit consisting of an LC impedance matching network, a 5-stage voltage multiplier, and a 376 μF capacitor storage bank

and stored on the capacitor bank until the voltage threshold for the application is met. The power management circuit controls basic duty cycling and regulation. The MSP430 microcontroller provides a higher level of power management by choosing when to take sensor measurements and controlling when to power the CC2500 low-power transceiver. When ready, the MSP430 takes sensor measurements, computes the data payload, and communicates back to the base station using the CC2500 radio IC. A back of the envelope calculation estimates the energy required for a burst of 16 bytes to be approximately 3 mJ.

2.1 RF Harvester and Power Management

Figure 5 shows a schematic of the WARP RF harvester consisting of an LC impedance matching network that conjugate matches the antenna’s impedance to the input of the multistage rectifier. For this application, it was necessary to maximize output voltage, so a five-stage voltage doubler was used instead of the four stage used in the previous application. The rectified voltage is stored on eight 47 μF low loss ceramic capacitors, for a total of 376 μF of charge storage.

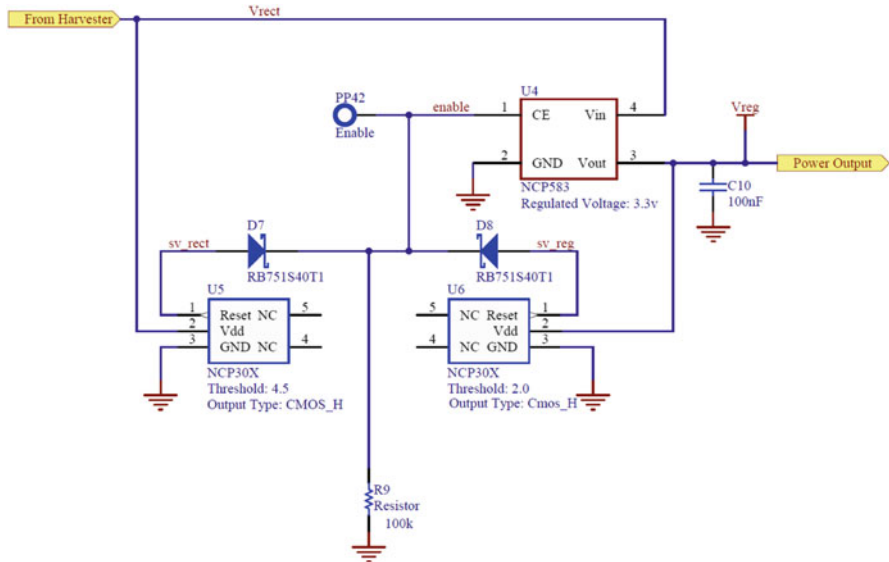


Fig. 6 The power management circuit allows charge to accumulate on the capacitor bank until 4.5 V is reached, at which point the circuit enables the regulator until the storage charge drops below 2.0 V

Figure 6 shows the power management circuit, which consists of a 3.3 V low-dropout (LDO) linear voltage regulator and two voltage supervisors with detection thresholds of 2.0 and 4.5 V, respectively. The reason for the two voltage supervisors is that the control signal that enables the output of the regulator needs a wide hysteresis window. Consider the following example: at startup the rectified voltage is small and both supervisors are low, causing the regulator to be disabled. Thus, the quiescent current draw of the system is small, which allows the circuit to harvest and store energy on the capacitor bank. When the rectified voltage reaches 4.5 V, the primary supervisor enables the regulator, which supplies voltage to the MCU. Typically, these supervisors have only 0.1 V of hysteresis, which means that when the voltage on the capacitor bank drops below 4.4 V, the primary supervisor will immediately disable the regulator, halting operation. To prevent this from happening, the secondary supervisor, which has a threshold of 2.0 V, is fed by the regulator output. This feedback circuit keeps the regulator enabled until the output drops below 2.0 V. Once the rectified voltage is below 2.0 V, the system enters its low quiescent current state and the cycle repeats.

In this system, there are two additional key voltage conditions that should be kept in mind. One, the LDO regulator, when enabled, will pass current if the input voltage is less than its specified regulated output voltage (3.3 V in this case). This means for $V_{in} \geq 3.3\text{ V}$, $V_{out} = 3.3\text{ V}$ and for $V_{in} \leq 3.3\text{ V}$, $V_{out} = V_{in}$. Two, the MSP430 will meet the required time constraints (for this application’s clock speed) down to $\sim 2.0\text{ V}$. Therefore, even though the voltage supplied to the MSP430 is unregulated from 3.3

to 2.0 V, reliable operation will occur. The result is that all the charge that is stored on the capacitor bank from 4.5 to 2.0 V is available for use. Future revisions of the WARP power management circuitry will focus on reducing the required MCU voltage down to 1.8 V by lowering the required clock speed for operation and by using a higher efficiency regulator.

Several sensors are also included on the WARP board. The analog to digital converter (ADC) and associated sensor circuits are constructed to minimize power consumption. The general technique used in creating the sensor circuits is that the microprocessor will control a line that enables the sensor by pulling a pin high. This prevents power leakage from the sensors when not in use. The sensors used in this design include an ambient light sensor, a thermometer, and a voltage divider connected to the capacitor bank. The voltage divider enables measurement of the energy stored in the capacitor bank.

2.1.1 Radio Transceiver and Embedded System

The WARP sensor node uses Texas Instrument's MSP430F2272 microcontroller for sensor measurement, computation, and to interface with the TI CC2500 low-power 2.4 GHz RF transceiver. The MSP430 is an excellent device for low-power sensing applications. It consumes 400 μ A in active mode at 8 MHz and 600 nA in standby mode and can operate from 3.3 V down to 1.8 V. The CC2500 low power radio also operates at 1.8 V and requires 15 mA while transmitting and 400 nA while in sleep mode. Additionally, the CC2500 supports the TI's SimpliciTI protocol for small, low-power RF networks.

The embedded software in this system performs three primary tasks. First, it manages the power usage and optimizes the circuit for low-power consumption. Second, it collects data and constructs packets to send back to the access point. Third, it handles data transmission and communication protocols. Since power is very limited, the code is constructed in a way that minimizes the impact of unexpected power failures by sampling the stored voltage and entering low-power mode until data transmission is possible.

When conducting the ADC measurements, there are three steps which help reduce power consumption. The sensor is normally disabled, so that it consumes no power. When a sensor measurement is to be made, the sensor-enabled pin is driven high to enable the sensor. Next, the microprocessor enters standby mode LPM1 for 250 cycles (31 μ s) to conserve power while the ADC reference voltage settles. Then, the microprocessor wakes up and reads the measured voltage.

The transceiver architecture is based on the TI eZ430-RF2500 wireless development tool kit. A second TI eZ430-RF2500 board acts as a base station hub when attached to a host computer. The WARP board wirelessly communicates with the host node at a range of up to 30 feet using the Texas Instruments SimpliciTI protocol. In order for the system to operate on ambient RF TV power, the protocol had to be modified in two ways to reduce power. In the SimpliciTI protocol, nodes randomly generate a 64 bit identification number and broadcast it to initiate linking. Then,



Fig. 7 WARP sensor node consisting of a rectifier, power management circuitry, sensors, and 2.45 GHz transceiver with ceramic chip antenna (power-harvesting antenna not shown)

the access point replies with its address. Linking requires the node to spend power on transmitting its ID number and receiving an access point address. The linking process requires both TX and RX, which both draw 15 mA. Due to the power requirements of this project, we decided to bypass this linking procedure and give each node a hard-coded, 8 bit identification number. These two changes allow the microprocessor to transmit data reliably while reducing the amount of harvested power needed to accomplish packet transmission.

2.2 Analysis of Power Usage

Figure 7 shows an image of the complete WARP sensor board. Initial development of the system firmware was done with bench top equipment to better control system variables. In order to emulate the harvested RF power, an 8 k Ω resistor was placed in series between a bench power supply and the capacitor bank (i.e., the V_{rect} node in Fig. 5). This resistor approximates the output resistance of the rectifier. By adjusting output voltage of the power supply, different amounts of power can be injected into the WARP board. Measurements show that approximately 50 μW of harvested power is required to charge the capacitor bank up to 4.5 V with a cycling period (charge/discharge) of 5 seconds.

Measurements of the microcontroller and transceiver showed that approximately 2.8 mJ of energy is consumed each time a packet is sent. The total transaction time




Fig. 8 Scope plot of a packet transmission cycle annotated with program events

is 8.3 ms. In order to understand the details of where power was consumed, the microcontroller was programmed to toggle a debug pin at each stage of operation. In this experiment, the system was powered with the same RC circuit used earlier to emulate harvested power. Once the 4.5 V supervisor triggered the transmit cycle, an oscilloscope was used to record the change in voltage on the capacitor bank. Figure 8 shows the annotated transmission cycle. The general microcontroller operation and analog to digital conversion consumed a very small percentage of the overall power used. Only 8.4% of the power consumed went towards running the microprocessor and bringing the device back into active mode. While 9.1% was consumed by the analog to digital sensor readings, 13.6% of the power was spent by the SimpliciTI protocol preparing to send data, 19.5% was consumed initializing the CC2500, and 49.4% was spent transmitting data.

2.3 Field Test of the WARP Sensor Node

Figure 9 shows the completed WARP sensor node, which uses the same 5 dBi log periodic antenna used in Sect. 1. A trial run of the WARP sensor node was conducted on the balcony of Intel Labs Seattle. The base station and host computer were located inside the Intel office, approximately 10 feet from the WARP node.




Fig. 9 Image of the fully assembled WARP sensor node consisting of a 5 dBi log periodic antenna and custom circuit board

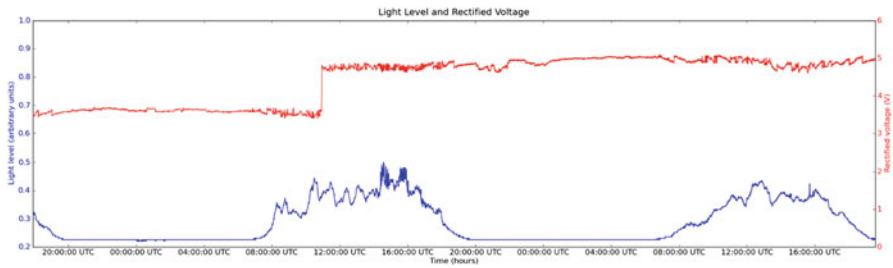


Fig. 10 Data recorded over 2 days while operating on harvested RF power

For this experiment the device was configured to take measurements and send data once every 5 s. This provides an extra margin of safety to ensure that harvested power remains higher than consumed power. During this trial run we successfully monitored ambient temperature, ambient light, and the voltage of the capacitor bank.

Figure 10 shows the collected data over two days. The ambient light level tracks the rising and setting of the sun which occurs at approximately 7:30am and 8:00pm each day. Variations in the recorded light level are due to passing cloud cover. In the figure, it can be seen that the energy harvested increases at 11:00am on the first day and remains high for the remainder of the experiment. It is hypothesized that this is due to the TV station boosting the signal strengths to account for changes in the atmosphere, which is a common occurrence in long-distance communication systems.

3 Summary

This project successfully demonstrated the ability to harvest wireless ambient radio waves and use them to power both a commercially available home weather station and a custom-built wireless sensor node. We tested the sensor node at a distance of 4.1 kilometers from the TV broadcast tower and wirelessly reported sensor data back to a host computer at five-second intervals. Future work will focus on refining the sensor node in several ways. First, optimizing the firmware to reduce the MCU clock speed will allow for operation at 1.8 V, thus saving power. Second, a switching regulator should be used in the power management circuit to more efficiently convert energy stored in the capacitor bank to the 1.8 V required for system operation. Finally, new antenna designs should be explored to optimize for size, bandwidth, and gain for a given application.

As hardware power requirements continue to drop, the viability of using ambient radio energy will continue to rise. The system described in this chapter is a first step towards harnessing this ambient RF energy. There are numerous wireless sensor networks in existence today, which are quite constrained by both battery power and the limitations of other harvesting techniques. The prospect of applying this system in new ways is exciting, given the abundance of ambient RF energy in urban areas. This platform enables wireless sensor systems to run perpetually, with no battery or data collection maintenance required.

References

1. S. Ahson and M. Ilyas. *RFID handbook: applications, technology, security, and privacy*. Boca Raton: CRC Press, 2008.
2. W.C. Brown. The history of power transmission by radio waves. *IEEE Transactions on Microwave Theory and Techniques*, 32(9):1230–1242, Sep 1984.
3. J.O. McSpadden and J.C. Mankins. Space solar power programs and microwave wireless power transmission technology. *IEEE Microwave Magazine*, 3(4):46–57, Dec 2002.
4. A. Sample and J.R. Smith. Experimental results with two wireless power transfer systems. In *2009 IEEE Radio and Wireless Conference. (RAWCON 2009)*, pages 16–18, jan. 2009.
5. A.P. Sample, D.J. Yeager, P.S. Powledge, A.V. Marnishev, and J.R. Smith. Design of an RFID-based battery-free programmable sensing platform. *IEEE Transactions on Instrumentation and Measurement*, 57(11):2608–2615, Nov. 2008.

A Portable Transmitter for Wirelessly Powering a Ventricular Assist Device Using the Free-Range Resonant Electrical Energy Delivery (FREE-D) System

Benjamin H. Waters, Jordan T. Reed, Kara R. Kagi, Alanson P. Sample, Pramod Bonde, and Joshua R. Smith

1 Wireless Power

Wireless power transfer using inductive coupling is becoming increasingly popular for consumer electronic devices. Emerging applications include wireless charging pads, electric toothbrushes, induction cookers, and electric car battery chargers. However, all of these applications are limited in their ability to transfer power in a noncontact fashion. Charging pads and electric toothbrushes require that the device be placed very close by or directly on top of the charging pad. This is because the efficiency for inductively coupled wireless power transfer systems drops off immediately as the distance between the transmitter (Tx) and receiver (Rx) increases.

Range and mobility can be increased with resonant coupling techniques. Resonant systems, like FREE-D, enable true wireless power transfer in that devices can be charged in free space, without direct physical contact between the Tx charger and Rx device. Wireless power transfer efficiency will eventually decrease as the distance between the Tx and Rx resonators increases, but unlike inductive coupling, this working range is significantly larger for resonantly coupled systems. The working range is dependent on the size of the Tx and Rx resonators. As a rule of thumb, power can be transferred at maximum efficiency if the distance between the Tx and Rx resonators is within one resonator diameter; beyond this working range, efficiency decreases [6].

B.H. Waters (✉) • J.T. Reed • K.R. Kagi • J.R. Smith • A.P. Sample
Department of Computer Science and Engineering and Department of Electrical Engineering,
University of Washington, Seattle, WA, USA
e-mail: bhw2114@uw.edu

P. Bonde
Yale School of Medicine, New Haven, CT, USA

The frequency of wireless power transfer systems for commercial applications must operate within the frequency bands that the FCC has reserved for industrial, scientific, and medical (ISM) applications. For resonant coupling, in order to operate at constant efficiency in the working range, the frequency must be actively tuned to track the maximum power transfer points. This is problematic because these optimal frequencies can sometimes lie outside the allowable bandwidth of the ISM bands.

The second difficulty with resonant coupling is that most electronic devices like laptops, cell phones, and electric cars that would be high-profile applications for wireless power integration contain lots of metal and other conductive materials. Although resonant coupling will allow for power transfer around metallic objects, when conductive materials are close to or surrounding either one or all of the resonators, wireless power transfer capabilities will be extremely limited. Techniques using ferrite materials, which have high magnetic permeability and low electrical conductivity, are being explored to isolate the resonators from nearby metal and minimize this undesirable effect.

Despite these regulatory and technological difficulties, the VAD remains an ideal application for integrating resonantly coupled wireless power transfer systems because the VAD requires long-range wireless power transfer, its power demands are relatively constant and there will not be highly conductive materials around the resonators.


1.1 The VAD Problem

Heart failure is a terminal disease with a very poor prognosis and constitutes Medicare's greatest area of spending with annual spending close to 35 billion [7]. The most desirable treatment for this disease remains a heart transplant; however, only a minority of patients (approximately 2,000 per year) can benefit from heart transplants due to continuing donor shortages.

An alternative treatment is the use of VADs [7]. The VAD continuously pumps blood throughout the body using a centrifugal, axial, or pulsatile rotor and is typically powered by a cable which penetrates the body wall to communicate with the device that is deeply implanted within the body cavity. The benefit of VAD technology was demonstrated by the REMATCH trial from 1998 to 2001, which showed a 50% survival benefit for patients assigned to VADs [8]. However, the same trial also showed the limitations of the first-generation VADs due to their large size and mechanical failure typically occurring within the first 2 years of implantation [8]. This prompted researchers and industry to concentrate their efforts on minimizing the size of the device, which has been achieved in the past decade with miniature axial and centrifugal pump technology [9,10]. Figure 1 shows a block diagram of the typical configuration for a VAD.

VADs differ from other metallic implants in a unique way: the presence of a transcutaneous driveline that penetrates the skin to transmit data to the external system controller and provide power to the VAD. However, as VAD technology

Fig. 1 Block diagram of the typical configuration for a VAD implantation. The double line represents the percutaneous driveline



continues to improve and moves towards even longer durations of patient support on VADs [11], the risk of exit site infection (ESI) from the driveline continues to increase temporally, hampering the patient's quality of life and leading to repeated hospitalizations for antibiotic treatment or surgical interventions [12] and in rare instances necessitating a pump exchange [13]. The net result of these effects due to ESI is the reduced survival and increased cost negating the intended benefit of VAD therapy.

A battery-powered solution is available for the existing VAD technologies. Rather than using a power base unit (PBU) as the permanent power supply, two batteries—a primary and secondary in case one fails—can be connected to the VAD system controller and provide power to the implanted pump for an extended period of time. Although this solution improves patient mobility and allows patients to leave the household, the percutaneous driveline is still susceptible to ESI. Integrating a portable, battery-powered solution with the FREE-D system will increase the reliability of a wirelessly powered VAD system and maximize patient mobility both in and outside the household, while eliminating the risk of ESI.

1.2 The Portable FREE-D System

The in-home FREE-D system has wirelessly powered both an axial and centrifugal VAD pump over meter distances [1–4]. The grand vision for implementing the in-home FREE-D system is shown in Fig. 2. The vision for the portable FREE-D system is shown in Fig. 3.

A Tx resonator is built into an exterior vest that will be worn by the patient. The vest will have a holster for both the batteries and the additional power management circuitry required to power the implanted VAD pump and system controller. A single battery will be capable of powering the system entirely on its own, and the secondary battery will only activate if the first expires or fails. To further improve the reliability of this system, a third battery will also be implanted along with the implanted receive resonator. The implanted battery will provide power directly to the VAD if the wireless power transferred to the implanted receive resonator fails. Also, there is an RF-DC rectifier and a DC-DC regulator to convert the transmitted AC signal into a DC voltage compatible with the VAD pump controller. Figure 4 shows a detailed block diagram of the portable FREE-D system.


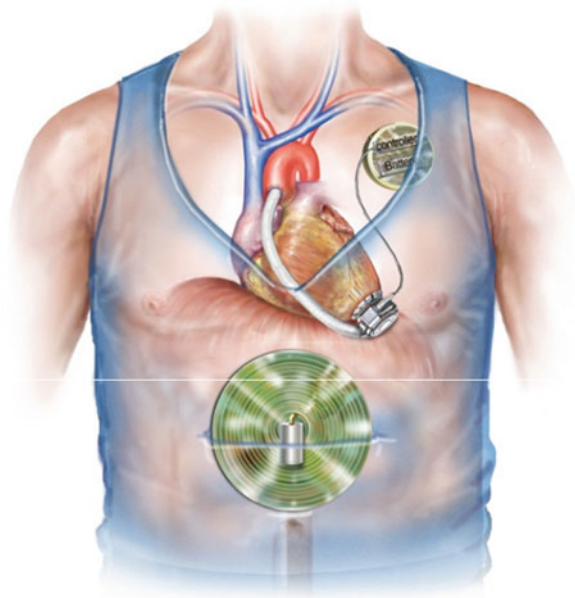


Fig. 2 Sketch of the grand vision for the in-home FREE-D system showing multiple transmit resonators delivering power wirelessly to the vest resonator, which relays power to the implanted receive resonator

Fig. 3 Sketch of the portable FREE-D system showing the battery-powered vest resonator and the implanted receive resonator connected to the VAD






Fig. 4 Block diagram of the portable FREE-D system




Fig. 5 Equivalent circuit model of the basic FREE-D wireless power system. Each of the four resonators is linked by the coupling coefficients k_{12} , k_{23} , and k_{34} (© 2011IEEE [6])

2 Overview of the FREE-D System

The FREE-D system uses magnetically coupled resonators to wirelessly transfer power from the Tx to Rx resonator. The wireless power circuit model and the additional circuitry for dynamic power management control are summarized in the following sections.

2.1 Wireless Power Circuit Model

The magnetically coupled resonator system can be represented in terms of lumped circuit elements (L , C , and R). Figure 5 shows a straightforward circuit diagram that can be used for hand analysis or for SPICE simulations.

The schematic consists of four resonant circuits, linked magnetically by coupling coefficients k_{12} , k_{23} , and k_{34} . Starting from the left, the drive loop is excited by

a source with finite output impedance R_s . A simple single-turn drive loop can be modeled as an inductor L_1 with parasitic resistance R_{p1} . A capacitor C_1 is added to make the drive loop resonant at the frequency of interest. The Tx resonator consists of a multi-turn air core spiral inductor L_2 , with parasitic resistance R_{p2} . The geometry of the coil determines its self-capacitance, which is represented as C_2 . Inductors L_1 and L_2 are connected with coupling coefficient k_{12} ; the Rx side is defined similarly. Finally, the Tx and Rx resonators are linked by coupling coefficient k_{23} . The implementation of the portable FREE-D system has the drive loop and Tx resonator built into a single device such that k_{12} is fixed. Similarly, k_{34} is also fixed. Thus, k_{23} is the remaining uncontrolled value, which varies as a function of the distances between the Tx and Rx resonators.

This circuit model provides a convenient reference for analysis of the transfer characteristics of a magnetically coupled resonator system. For the sake of simplicity, the cross-coupling terms (k_{13} , k_{24} , and k_{14}) are neglected in the following analysis. Kirchhoff's voltage law (KVL) can be applied to determine the currents in each resonant circuit. The four KVL equations are simultaneously solved for the voltage across the load resistor and yield Eq. (1), with the substitution in Eq. (2) and the coupling coefficient defined in Eq. (3).

$$\frac{V_L}{V_S} = \frac{j\omega^3 k_{12} k_{23} k_{34} L_2 L_3 \sqrt{L_1 L_4} R_L}{k_{12}^2 k_{34}^2 L_1 L_2 L_3 L_4 \omega^4 + Z_1 Z_2 Z_3 Z_4 + \omega^2 (k_{12}^2 L_1 L_2 Z_3 Z_4 + k_{23}^2 L_2 L_3 Z_1 Z_4 + k_{34}^2 L_3 L_4 Z_1 Z_2)} \quad (1)$$

$$\begin{cases} Z_1 = R_1 + R_s + j\omega L_1 + \frac{1}{j\omega C_1} \\ Z_2 = R_2 + j\omega L_2 + \frac{1}{j\omega C_2} \\ Z_3 = R_3 + j\omega L_3 + \frac{1}{j\omega C_3} \\ Z_4 = R_4 + R_L + j\omega L_4 + \frac{1}{j\omega C_4} \end{cases} \quad (2)$$

$$k_{ij} = \frac{M_{ij}}{\sqrt{L_i L_j}}, \quad 0 \leq k_{ij} \leq 1 \quad (3)$$

The system transfer function (1) is plotted in Fig. 6 for the circuit values shown in Table 1. This plot shows $|S_{21}|$ as a function of frequency and coupling coefficient k_{23} .

For consistency, power transfer will be represented in terms of linear magnitude scattering parameters $|S_{21}|$, which is important experimentally since it can be measured with a vector network analyzer (VNA) for comparison. The entire wireless power transfer apparatus can be viewed as a two-port network (one port being the input, fed by the source, and the other being the output, feeding the load). Using Eq. (1), one can calculate the equivalent S_{21} scattering parameter using [14, 15], which results in Eq. (4).

$$S_{21} = 2 \frac{V_{\text{Load}}}{V_{\text{Source}}} \left(\frac{R_{\text{Source}}}{R_{\text{Load}}} \right)^{1/2} \quad (4)$$

In Fig. 6, frequency splitting is clearly visible as the value of k_{23} is increased. A SPICE simulation reveals that, indeed, the lower frequency mode of the two




Fig. 6 Efficiency curve for the wireless power system resonators defined in Table 1 (© 2011 IEEE [6])

Table 1 Circuit values used to evaluate simplified model (© 2011 IEEE [6])

Parameter	Value
R_{Source}, R_{Load}	$50\ \Omega$
L_1, L_4	$1.0\ \mu\text{H}$
C_1, C_4	$235\ \text{pF}$
R_{p1}, R_{p4}	$0.25\ \Omega$
k_{12}, k_{34}	0.10
L_2, L_3	$20.0\ \mu\text{H}$
C_2, C_3	$12.6\ \text{pF}$
R_{p2}, R_{p3}	$1.0\ \Omega$
k_{23}	0.0001–0.30
f_0	10 MHz
Frequency	8–12 MHz

resonators is in phase, while the higher frequency mode is 180° out of phase. As the coupling between the resonators decreases, the frequency separation also decreases until the two modes converge at f_0 . This point is called the critical coupling point and represents the farthest distance at which maximum power efficiency is still achievable (since k_{23} is proportional to $1/\text{distance}^3$). When k_{23} is greater than $k_{critical}$, the system is said to be overcoupled, and operating at either resonance will result in maximum power transfer efficiency. Conversely, when k_{23} is less than $k_{critical}$, the system is undercoupled, and the amount of power delivered to the load begins to fall off precipitously with distance. The red dashed box outlined in Fig. 5

encloses the “magic regime” where near-constant efficiency versus distance can be achieved if the correct frequency is selected. This is dramatically different from typical far- or near-field systems where efficiency drops off rapidly with distance.

2.2 Battery-Powered Transmitter Circuit Model

A DC-RF inverter must be used at the Tx side in order to convert the DC voltage supplied by the portable FREE-D system’s battery to an RF signal oscillating at the resonant frequency of the FREE-D resonators. For magnetically coupled resonators, the mutual inductance of the resonators is constantly changing as the distance between the Tx and Rx resonators changes. Therefore, a Royer oscillator configuration has been selected so that the capacitance C_1 and the drive loop inductance L_1 determine the frequency of oscillation in Eq. (5). A Royer-type oscillator is a push–pull circuit that sends current from V_{DC} through inductors L_{M1} and L_{M2} to begin the oscillating cycle caused by power MOSFETS M_1 and M_2 turning on and off. Resistors R_1 and R_2 bias the gate of the M_1 and similarly for M_2 .

$$f_o = \frac{1}{2\pi\sqrt{L_1C_1}} \quad (5)$$

Rather than using a directional coupler and an MCU to perform frequency tracking, the Royer oscillator configuration automatically operates at the resonant frequency of each coil, which is also determined by Eq. (5). The oscillator is acting like a variable frequency oscillator (VFO) that uses a variable inductance—the mutual inductance of the FREE-D resonators M_{ij} —for frequency tuning. As the distance between the Tx and Rx resonators increases, the mutual inductance between the resonators decreases, which changes the oscillating frequency.

The VFO for the portable FREE-D system (Fig. 7) is slightly modified from an ideal Royer oscillator [16].

The FREE-D VFO uses inductors L_{M1} and L_{M2} to drive the MOSFETS M_1 and M_2 in order to generate the RF signal oscillating at the resonant frequency of the coils. Depending on the speed of the VAD pump, 5–15 W can be required to provide sufficient power to the pump. Therefore, all components for the FREE-D VFO must have sufficient power ratings to satisfy these requirements. The most critical components are the two power MOSFETS. The junction capacitances must be significantly smaller than the self-capacitance of the drive loop such that they do not affect the oscillation frequency. Also, the rise and fall times must be short enough such that the VFO can generate a high enough frequency—7.65 MHz for this VAD application.




Fig. 7 Circuit schematic of the FREE-D VFO. The components C_1 and L_1 set the frequency of oscillation

Fig. 8 A centrifugal pump (VentrAssist LVAD)



3 Model Implementation and Experimental Results

The portable FREE-D system has been tested at both high power levels using $2 \times 12\text{ V}$ nickel metal hydride (NiMH) batteries to power a centrifugal pump VAD (Fig. 8) and from a signal level using a vector network analyzer (VNA). Using the battery-powered system, the entire portable FREE-D system provides sufficient power to operate a centrifugal pump VAD.

Figure 9 shows the successful operation of the portable FREE-D system. The resonators used in the experimental implementation in Fig. 9 are different than the resonators modeled in Table 1. Table 2 characterizes the resonators shown in Fig. 9.

Next, the set of FREE-D resonators were connected to a VNA to analyze the frequency tracking capabilities of the VFO. The frequencies at which $|S_{21}|$ is a maximum are shown as a function of distance in Fig. 10. Only the overcoupled




Fig. 9 The portable FREE-D system wirelessly powering the VentrAssist centrifugal pump VAD. From *right to left*: 2×12 V NiMH batteries, FREE-D VFO, Tx drive loop, Tx coil, Rx coil, Rx receive loop, VentrAssist VAD, rectification circuitry, and VAD system controller

Table 2 Characteristics of the resonators used in the experimental implementation of the portable FREE-D system (Fig. 9)

Parameter	Tx loop	Tx coil	Rx coil	Rx loop
Outer diameter	19 cm	19 cm	19 cm	19 cm
Distance from drive loop	0 cm	4 cm	16 cm	26 cm
# Turns	1	9	9	1
Pitch	n/a	1 cm	1 cm	n/a

regime has been shown in Fig. 10, where strong coupling between the Tx and Rx resonators causes frequency splitting due to the two resonant modes of the system: the low-frequency mode (in phase) and the high-frequency mode (out of phase). Typically, the low-frequency mode corresponds to the maximum $|S_{21}|$ value. Therefore, if the transmitter can dynamically operate at the frequency corresponding to the low-frequency curve in Fig. 10, then maximum power will be transferred at constant efficiency independent of the distance between the Tx and Rx resonators.

To analyze how well the FREE-D VFO performs frequency tracking, an oscilloscope was used to determine the frequency of operation for the portable FREE-D system at high power. Figure 10 confirms that the FREE-D VFO tracks the in-phase mode, corresponding to the frequencies at which the maximum efficiency is observed. Also, both the low-frequency mode, and high-frequency mode and the oscillator frequency curves converge on the expected resonant frequency of the Tx and Rx coils, which are both tuned to resonate at 6.78 MHz.

The slight discrepancy between the low-frequency mode and the oscillator frequency curve is due to the contribution of the parasitic capacitance $C_{p, \text{VFO}}$ from the FREE-D VFO. As expected from Eq. (5), as capacitance increases, the oscillating frequency will decrease. In order for the FREE-D VFO to oscillate at the




Fig. 10 Plot of the frequency corresponding to the maximum values of $|S_{21}|$ for both the in-phase and out-of-phase modes as a function of distance between the Tx and Rx resonators. The plot compares the peak frequency data taken from the VNA (red and blue curves) with the oscillating frequency values that were measured during the full-power operation of the FREE-D VFO (green curve). The center frequency (purple dotted line) indicates the frequency that each coil was initially tuned to resonate (6.78 MHz)

resonant frequency of the coils, the tuning capacitance C_1 of the Tx drive loop must account for the additional parasitic capacitance contribution from the FREE-D VFO such that the oscillating frequency matches the resonant frequency as in Eq. (6).

$$f_o = \frac{1}{2\pi\sqrt{L_1(C_1 + C_{P,VFO})}} \tag{6}$$

This capacitance has been accounted for in the experimental implementation of the FREE-D VFO as in Fig. 9. The FREE-D VFO is capable of transmitting up to 15 W, which can satisfy the power requirements for all pump speeds of the centrifugal pump VAD, and has demonstrated the capability to perform frequency tracking that will ensure maximum power transfer efficiency from Tx to Rx.

4 Future Work and Conclusions

FREE-D technology can be expected to have many medical applications that include unobtrusive recharging of batteries in any implanted electronics, including pace-makers, implanted defibrillators, cochlear implants, retinal implants, and so forth. Outside the medical industry, some ideal consumer electronic applications include mobile phone charging, electric car charging, and kitchen appliance operation.

For use with VADs, the future work for the FREE-D system includes developing a dynamic impedance matching network [17] for single-frequency operation and improved implantable resonator designs to minimize localized temperature changes of the implanted components [18].

Lastly, exploring other oscillator configurations will be useful not only for this VAD application but also for RF transmitters. A multimode oscillator that can operate at several resonant frequencies could improve the frequency tracking capabilities of the FREE-D VFO by operating at either the high- or low-frequency mode and also be useful as a digital FM transmitter.

Acknowledgments The section entitled “Wireless Power Circuit Model” is based on Section III of [6], which is Copyright 2011 IEEE, and used by permission.

References

1. Smith J.R., Sample A.P., Waters B., Toyoda Y., Kormos R., and Bonde P., “Innovative Free-Range Resonant Electrical Energy Delivery System (FREE-D system) for a Ventricular Assist Device (VAD) Using Wireless Power.” ASAIO 31st Annual Conference, 2011, Washington, DC, June 10–12, 2011
2. Bonde P., Sample A.P., Waters B., Cooper E., Toyoda Y., Kormos R., and Smith J.R., “Wireless Power for Ventricular Assist Devices: Innovation with the Free-Range Resonant Electrical Energy Delivery System (FREE-D) for Mechanical Circulatory Assist.” AATS 91st Annual Scientific Meeting, 2011, Philadelphia, May 7–11, 2011
3. Waters B., Sample A.P., Bonde P., Smith J.R., “Tether-Free Existence with the Free-Range Resonant Electrical Energy Delivery (FREE-D) System for Ventricular Assist Device (VAD) Recipients.” Annual Meeting for the Bio-medical Engineering Society, 2011, Hartford, CT, Oct. 12–15, 2011
4. Waters B., Sample A.P., Kormos R., Smith J.R., Bonde P., “Powering a Ventricular Assist Device over Meter Distances Wirelessly: The Free-Range Resonant Electrical Energy Delivery (FREE-D) System.” Heart Failure Society of America 15th Annual Scientific Meeting, 2011, Boston, MA, Sep. 12–15, 2011
5. VentrAssist LVAS Clinical Instructions for Use. Ventracor 2010
6. Sample A.P., Meyer D., Smith J.R., “Analysis, experimental results, and range adaptation of magnetically coupled resonators wireless power transfer.” IEEE Trans Ind Electron. 2011 Feb;58:544–554
7. Deng M.C., Edwards L.B., Hertz M.I., Rowe A.W., Kormos R.L., “Mechanical circulatory support device database of the international society for heart and lung transplantation: first annual report – 2003.” J Heart Lung Transplant. 2003 Jun;22(6):653–62
8. Holman W.L., Pamboukian S.V., McGiffin D.C., Tallaj J.A., Cadeiras M., Kirklin J.K., “Device related infections: are we making progress?” J Card Surg. 2010 Jul;25(4):478–83
9. Miller L.W., Pagani F.D., Russell S.D., John R., Boyle A.J., Aaronson K.D., Conte J.V., Naka Y., Mancini D., Delgado R.M., MacGillivray T.E., Farrar D.J., Frazier O.H., “HeartMate II clinical investigators. Use of a continuous-flow device in patients awaiting heart transplantation.” N Engl J Med. 2007 Aug 30;357(9):885–96
10. Pagani F.D., Miller L.W., Russell S.D., Aaronson K.D., John R., Boyle A.J., Conte J.V., Bogaev R.C., MacGillivray T.E., Naka Y., Mancini D., Massey H.T., Chen L., Klodell C.T., Aranda J.M., Moazami N., Ewald G.A., Farrar D.J., Frazier O.H., “HeartMate II Investigators. Extended mechanical circulatory support with a continuous-flow rotary left ventricular assist device.” J Am Coll Cardiol. 2009 Jul 21;54(4):312–21

11. Slaughter M.S., Rogers J.G., Milano C.A., Russell S.D., Conte J.V., Feldman D., Sun B., Tatooles A.J., Delgado R.M. 3rd, Long J.W., Wozniak T.C., Ghumman W., Farrar D.J., Frazier O.H., "HeartMate II investigators. Advanced heart failure treated with continuous-flow left ventricular assist device." *N Engl J Med.* 2009 Dec 3;361(23):2241–51
12. Monkowski D.H., Axelrod P., Fekete T., Hollander T., Furukawa S., Samuel R. "Infections associated with ventricular assist devices: epidemiology and effect on prognosis after transplantation." *Transpl Infect Dis.* 2007 Jun;9(2):114–20
13. Wilson W., Taubert K.A., Gewitz M., Lockhart P.B., Baddour L.M., "Prevention of infective endocarditis." *J Am Dent Assoc.* 2008 Jan;139 Suppl:3S-24S
14. Mongia R. "Rf and Microwave Coupled-Line Circuits." City: Artech House Publishers, Boston, 2007
15. Chen J. "Feedback Networks: theory and circuit application." City: World Scientific Publishing Company, Singapore, 2007
16. Pressman A.I., Billings K., Morey T. "Switching Power Supply Design", pp. 266–270. The McGraw Hill Companies, New York, 2009
17. Waters B.H., Sample A.P., Smith J.R., "Adaptive Impedance Matching for Magnetically Coupled Resonators." *PIERS Proceedings*; pp. 694–701. Moscow, Russia, August 19–23, 2012
18. Christ A., Douglas M., Roman J., Cooper E., Sample A., Waters B., Smith J.R., Kuster N., "Evaluation of Wireless Resonant Power Transfer Systems With Human Electromagnetic Exposure Limits," *IEEE Transactions on Electromagnetic Compatibility*, October 2012

Part VI
Systems and Applications

PORFIDO: Using Neutrino Telescopes and RFID to Gather Oceanographic Data

Orlando Ciaffoni, Marco Cordelli, Roberto Habel, Agnese Martini, and Luciano Trasatti

1 Neutrino Telescopes

Neutrinos are the most elusive of elementary particles. They have no charge, no magnetic properties, and an extremely small mass. Moreover, they interact very little with the surrounding matter. Neutrinos may easily escape even from large stars, are hardly absorbed during their propagation, and are not deflected in magnetic fields. They are therefore the ideal candidate particles to investigate the most remote stars in the universe. And looking at greater distances in the universe also means looking backwards in time, closer and closer to the big bang.

However, the detection of low fluxes of these very elusive particles requires a detector with a huge sensitive volume, which should be effectively shielded from the overwhelming background from atmospheric cosmic rays. The only viable solution nowadays is to build a large array of light sensors in a transparent medium, such as seawater. Photomultipliers can then detect the Cerenkov light emitted by neutrino interactions inside the apparatus or in its immediate surroundings.

A neutrino telescope with a sensitive volume of the order of one cubic kilometer is needed for neutrino astrophysics. And it needs to be buried at a depth of about 3,000 m under the sea surface. A total of 10,000 optical modules (OM) will be installed, consisting of a photomultiplier tube (PMT) and related communication electronics enclosed in a 12 mm thick glass sphere with a diameter of 40 cm, a standard oceanographic instrument (see Fig. 1).

The NEMO [1] (neutrino Mediterranean observatory) collaboration was set up in 1998 with the aim to carry out the necessary R&D towards a km³ neutrino telescope in the Mediterranean Sea (see Fig. 2).

O. Ciaffoni • M. Cordelli • R. Habel • A. Martini • L. Trasatti (✉)
Laboratori Nazionali di Frascati, Istituto Nazionale Fisica Nucleare, Via E. Fermi 40, I-00044, Frascati, Italy
e-mail: orlando.ciaffoni@lnf.infn.it; marco.cordelli@lnf.infn.it; roberto.habel@lnf.infn.it; agnese.martini@lnf.infn.it; luciano.trasatti@lnf.infn.it



Fig. 1 Optical module

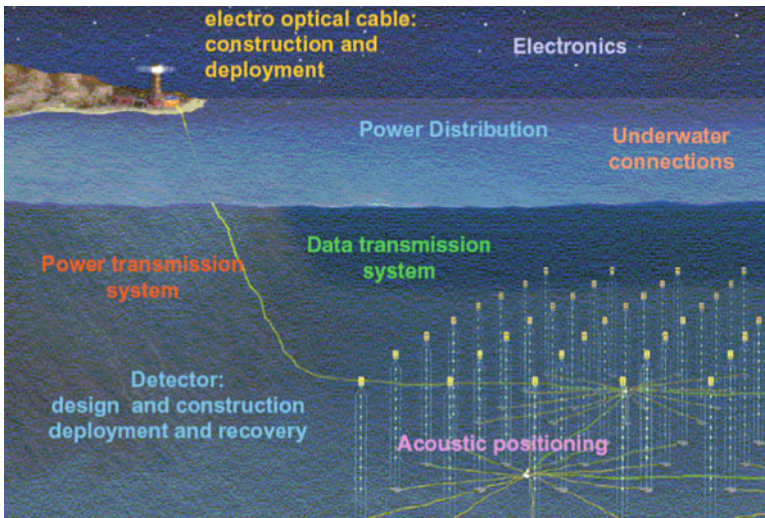



Fig. 2 NEMO km3 telescope conceptual design

Fig. 3 Optical module with PORFIDO probe schematic



Such an installation offers as a by-product new and revolutionary possibilities to oceanographers. Typical oceanographic data are collected by deploying self-contained instruments anchored to the bottom of the sea and recovering them after a period of weeks or months, when the batteries have died, to collect the acquired data. In contrast, a neutrino telescope installation offers to the oceanographic community the possibility to gather data continuously and in real time, since an essential part of the installation is an underwater electro-optical cable, which carries power to the telescope and data from it to shore. A very small fraction of both the power and of the bandwidth available even to a small telescope is sufficient to collect relevant oceanographic data. Clearly, before building the final gigantic apparatus, a number of tests with a relatively small number of optical modules are necessary.

The problem is to interfere as little as possible with the neutrino detector, which means using small probes and avoiding the need of connectors or penetrators, that are very expensive and that offer low reliability. In an apparatus placed in such an inaccessible environment, reliability is an essential requirement.

We have built such a system, PORFIDO [2], using the well-established technique of RFID to gather data through the glass spheres of the optical modules and to supply power to the sensors with the RF itself.

PORFIDO is made up of two elements: the sensor, which is glued to the outside of the optical module and gathers data directly from the seawater. The reader, which sits inside the OM, reads the measured data through the glass using RFID and communicates with the OM electronics to send the data to shore (see Fig. 3). We will test the system on the NEMO Phase II tower that is being built in Catania (Sicily) and will be deployed in 2012. It will consist of 32 optical modules, and four PORFIDO probes will be installed in four different optical Modules.




Fig. 4 The WISP

2 RFID (Radio-Frequency Identification)

Radio-frequency identification (RFID) is a technology that has been developed for access control and is spreading widely in this and other fields.

In the standard setup, a reader emits an RF beam, and the responding unit (tag) answers with its own identity code, deriving power from the RF itself and thus eliminating the need for batteries. Recent developments have focused on adding to the tag the possibility to take measurements in the environment and transmit them to the reader together with its ID code. The EPC Class 1 GEN 2 (C1G2) protocol [3], developed by EPCglobal, includes the possibility of sensors in RFID tags.

3 The RFID Tag: WISP

We have used as an RFID tag the WISP (see Fig. 4), developed in Seattle at Intel Labs and University of Washington (<http://wisp.wikispaces.com/>) [4, 5]. It is passive (no batteries), has a thermometer and an accelerometer on board, and is designed with an open architecture to include new sensors. Software for the integration is available from the designers. It can stand very well the extreme conditions that we require, namely 30 MPa (300 atm) pressure, a temperature of 15°C, and the exposure to seawater, if protected adequately. For this purpose we potted the WISP in a two-component epoxy.

4 The RFID Reader

The RFID reader to be installed inside the optical modules has to be small, does not need a lot of RF power, and must be cheap. Several firms offer such instruments, and we have chosen the ThingMagic [6] M5e-compact reader for its small footprint and good performance with the WISP.

Fig. 5 The PORFIDO pressure test



5 Functionality and Compatibility Tests

We ran a long series of tests to verify the feasibility of the system and the absence of interaction with the electronics of the NEMO Phase II tower.


The glass sphere on which PORFIDO had been mounted was immersed in seawater then put under a pressure of 40 MPa, equivalent to 4,000 m depth in water (see Fig. 5).

We also wanted to prove that the RF field generated by the reader does not interfere with the PM tube and with the electronic boards inside the optical module.

We built a black box and installed in it a working OM and the PORFIDO system, together with a LED light pulser. Turning on and off the reader no effects were detected on the measurement of the light pulse frequency.

6 An Improved Transceiver Antenna

The dipole antennas that are used for RFID proved adequate but difficult to handle, being very sensitive to the presence of conductive materials in the vicinity of the reader. At this point we realized that our setup, with an emitter and a receiver placed at a distance of only 12 mm and with glass in between, was more similar to a capacitor than to a transmitter–receiver pair.

Fig. 6 The sensor**Fig. 7** Two views of the sensor and the reader

Therefore, we discarded the long dipoles, cut the wings of the WISP, and built two capacitors using square copper pads, $25 \times 25 \text{ mm}^2$, facing each other on the two sides of the glass (see Figs. 6 and 7).


The results were excellent. The system was much more stable and immune to the presence of metal in the surrounding space. We were able to decrease the reader RF power by a factor of 10, and the interference decreased accordingly.

One more advantage of this new design is that the dimensions of PORFIDO decreased to the size of the reader alone, $70 \times 40 \text{ mm}$. This will make it much easier to incorporate the system in any new designs for the optical modules, since the area covered by PORFIDO is now very small compared to the total surface of the glass sphere (see Fig. 8).

7 Power Requirements

We have strict restrictions on the power available for PORFIDO. Since the reader draws about 0.5 A at 5 V for 200 ms , we have installed a 1 F capacitor to store

Fig. 8 PORFIDO mounted on the optical module



the necessary energy. The capacitor recharges slowly while the reader is inactive. Commands to the reader are sent transparently on a simulated serial connection in the optical fiber data link of NEMO.

8 Future Development

The WISP comes with a completely open architecture, so that different kinds of sensors can easily be implemented to work with PORFIDO. On the first implementation we only used a temperature sensor to prove that the system was functional, and we are starting to work on a salinity sensor and on a dissolved oxygen monitor.

9 Conclusions

The NEMO group has approved the installation of 4 PORFIDO probes measuring seawater temperature on the Phase II tower that will be deployed in 2012 at a depth of 3,500 m in the Capo Passero site (see Fig. 9), and we are now working to integrate it in the system.

We believe that the oceanographic community could greatly benefit by the use of this kind of instrument, an unobtrusive parasite on the Cerenkov neutrino telescopes which will be built in the future.

Acknowledgments We would like to thank Joshua R. Smith and the WISP team for the continuing support. We are grateful to the LNS and INFN-Roma1 components of the NEMO group who lent us a lot of essential support. Finally, we would like to thank the Catania University NEMO group for the use of their very efficient pressure chamber.

Fig. 9 The NEMO-KM4 underwater site



References

1. E. Migneco. Progress and latest results from Baikal, Nestor, NEMO and KM3NeT, Neutrino 2008 XXIII International Conference on Neutrino Physics and Astrophysics, www.slac.stanford.edu/econf/C0805263/Slides/Migneco.pdf
2. M. Cordelli, R. Habel, A. Martini, L. Trasatti. PORFIDO: oceanographic data for neutrino telescopes, LNF-10/02 (P), February 4, 2010 and Nuclear Instrumental and Methods in Physics Research A, 626–627, pp. S109–S110, 2011
3. EPCglobal Class 1 Generation 2, <http://www.gs1.org/epcglobal>
4. J.R. Smith, A. Sample, P. Powledge, A. Mamishev, S. Roy. A wirelessly powered platform for sensing and computation, Proceedings of Ubicomp 2006: 8th International Conference on Ubiquitous Computing. Orange Country, CA, USA, September 17–21 2006, pp. 495–506
5. A.P. Sample, D.J. Yeager, P.S. Powledge, A.V. Mamishev, J.R. Smith. Design of an RFID-based battery-free programmable sensing platform. IEEE Transactions on Instrumentation and Measurement, 57(11), 2008, pp. 2608–2615
6. ThingMagic, A Division of Trimble. Four Cambridge Center, 12th floor, Cambridge, MA 02142 United States

RFID-Vox: A Tribute to Leon Theremin

Pavel V. Nikitin, Aaron Parks, and Joshua R. Smith

1 Introduction

Radio-frequency identification is an automatic wireless data collection technology with a long history [1] which is usually traced back to World War II British aircraft identification transponders [2] and the seminal paper by Harry Stockman on principles of modulated backscattered communication [3]. A person who deserves a special mention in RFID history is Leon Theremin [4], the inventor of the first passive RFID tag-like device known as the Great Seal bug. Leon Theremin is more widely known as the inventor of the thereminvox, a contactless musical instrument (also known as simply the theremin). Interestingly enough, combining basic principles of the thereminvox with current passive UHF RFID technology allows one to create a long-range contactless musical instrument (“RFID-vox”) which can be played by moving RFID tags in the far field of the instrument antennas.

Section II presents a biographical overview of Leon Theremin and his work. Section III describes the RFID-vox concept and example implementations. Conclusions are drawn in Section IV.

P.V. Nikitin (✉)

Intermec Technologies Corporation, 6001 36th Ave W, Everett, WA 98203, USA
e-mail: nikitin@ieee.org

A. Parks

Department of Electrical Engineering, University of Washington,
Box 352350, Seattle, WA 98195, USA
e-mail: anparks@uw.edu

J.R. Smith

Department of Computer Science and Engineering, University of Washington,
Box 352350, Seattle, WA 98195, USA
e-mail: jrs@cs.washington.edu

2 Leon Theremin

On April 25, 1930, Carnegie Hall was very busy. Everyone wanted to see the concert performed by ten musicians, each simultaneously playing a thereminvox, a new electronic musical instrument invented by a Russian who organized the concert and was playing in it himself [5]. Thirty years later, in 1960, US ambassador Henry Cabot Lodge, Jr. was showing at the United Nations meeting a passive eavesdropping device that was discovered in the US embassy in Moscow, Russia.

What links those two events? The creator of both the thereminvox and the Great Seal bug was the same person—a prominent inventor and musician, Leon Theremin. We invite the reader to take a brief look at his extraordinary life. For more detailed information, we refer a reader to the article [6], the book [7], and the excellent documentary movie [8].

Lev Sergeyevich Termen (he became known as Leon Theremin after he came to America in 1927) was born on August 15, 1896, in St. Petersburg, Russia. He started learning music and physics at an early age and then went on to study physics and astronomy at the University of St. Petersburg. He also studied cello at the St. Petersburg Music Conservatory. During World War I, he was drafted, graduated from the Officers' Electro-Technical School, and served as an officer. After the Russian Revolution of 1917, he worked on equipment for the first radio stations of Soviet Russia. His life, like the lives of many others, was affected by the revolution: he spent time in prison in 1919–1920 after being accused of counterrevolutionary activity.

In 1920, he joined Ioffe Physical Technical Institute where he became the head of the new experimental laboratory and started working on high-frequency measurement methods. He found that the movement of one's hands affects the capacitance of electronic circuits and thus can be used to control oscillator pitch and volume. Using this effect, Leon Theremin created the first contactless musical instrument, originally called the etherphone (and later the thereminvox or simply the theremin). The basic physics of the thereminvox has been analyzed in detail in [9]. Based on the same capacitance-sensing effect, he also invented an alarm system. In 1922 he showed both inventions to Vladimir Lenin who liked them very much. In this period of time, Leon Theremin also worked on mirror-drum-based mechanical television and successfully demonstrated prototypes. He traveled across the country to demonstrate his inventions and was often referred to as the "Soviet Edison."

To promote his inventions, Leon Theremin went on an international tour, visiting Germany, England, France, and arriving in the USA in 1927. The thereminvox created a sensation there [10]. Leon Theremin established a laboratory in Manhattan where he worked on the thereminvox and other electronic musical instruments (electronic cello, terpsitone, etc.). In 1928, he received a US patent for the thereminvox [11] and sold it to RCA which began producing his instrument [12]. The instruments were expensive: in 1930 the cost of the complete system with speakers was approximately \$230 which corresponds to \$3,000 in today's dollars. Only about 500 instruments were produced and sold.

Fig. 1 Leon theremin plays theremin (1924) (courtesy Wikimedia foundation)



In 1930 Leon Theremin demonstrated ten thereminvoxes on the concert stage at Carnegie Hall [5], and in 1932 he conducted the first electronic orchestra performance there [13]. In that time period, Theremin closely interacted with many famous scientists and musicians, including Albert Einstein (who played violin), composer Joseph Schillinger, and thereminvox virtuoso Clara Rockmore, who helped him to promote his instrument (Fig. 1).

In 1931 he became a vice president of Teletouch Corporation which sold his patented “radio watchman” [14], a capacitance-based alarm system. One of its customers was Alcatraz prison. In 1936, he received his third US patent, for an electrical clock run by DC current [15]. In 1938 Leon Theremin married African American ballet dancer Lavinia Williams who bore him twin daughters (He had divorced his first wife, Katia Konstantinova, soon after he arrived to the USA).

In September 1938, he abruptly returned to the Soviet Union. Whether he returned voluntarily or was forced to is a subject of debate. In March 1939 he was arrested and sentenced to 8 years in prison. He was sent to the camp in Magadan, Kolyma, one of many in Gulag prison and labor camp system [16, 17]. Leon Theremin would have probably died there, as the survival rate in Stalin’s camps was very low. Fortunately, in 1940, he was transferred to a Moscow secret research and development laboratory, an elite part of the prison system, where he remained until 1947 working on various military projects. He worked with Sergei Korolev, who later became a key figure in the Soviet space program. Korolev went on to develop the rocket for the Sputnik launch and started the Soviet lunar program [18].

One of the projects which Leon Theremin worked on at prison became known as the Great Seal bug [19]. In 1945, Soviet Young Pioneers (analogous to boy scouts and girl scouts) presented to the US ambassador in Moscow a carved wooden replica of the Great Seal of the USA. This gift contained a passive listening device which was finally discovered by accident only in 1952, 7 years later. The device consisted of a monopole antenna connected to a resonator with a flexible sound-sensitive membrane. The sound in the room caused geometric deformations of the microphone, which in turn changed the antenna's RF reflection coefficient. The net effect was that the apparatus backscattered an incident carrier wave radio-frequency signal (originating from a transmitter outside the embassy) while modulating it with the voice of those present in the room. This device was essentially the first long-range passive UHF RFID tag. Modern RFID tags use basically the same operating principles of modulated backscatter. US ambassador Henry Cabot Lodge, Jr. demonstrated the Great Seal bug during the 1960 UN General Assembly session as an example of Soviet espionage. A replica of the Great Seal bug is currently on display at NSA National Cryptologic Museum in Annapolis Junction, MD [20] (Fig. 2).

In 1947, Leon Theremin was freed and awarded the Stalin Prize of the first degree for his work on eavesdropping devices (his other invention was Buran, an eavesdropping system which used an microwave beam to detect glass vibrations caused by sounds inside the room). After that, Leon Theremin continued to work on different military projects. He married for the third time and had twin daughters. In 1964, he joined the Moscow Conservatory where he worked on various electronic musical instruments. In 1967, an American journalist found him in Russia, interviewed, and published an article about him in New York Times [21]. It was the Cold War era. After the article came into print, Leon Theremin was immediately fired, his laboratory was closed, and most of his instruments were destroyed. For some time, he could not find any job. Finally, with help of his friends, he started working as a technician and lab assistant at the Physics Department of Moscow State University where he remained for the rest of his life.

In 1991, Steven M. Martin filmed a famous documentary about Leon Theremin [8] and brought him to visit the USA where Leon met again Clara Rockmore, after more than 50 years. Lavinia Williams, his wife whom he never saw after 1938, died in 1989, just two years before his visit. In 1991, Stanford University awarded Leon Theremin a Centennial Medal for contributions to electronic music. Leon Theremin died in Moscow on November 3, 1993, at the age of 97.

3 RFID-Vox

The thereminvox still holds an important place in electronic musical instruments. It was used for composing music by the Beach Boys (*Good Vibrations*, 1966), in Hollywood movies (*The Day the Earth Stood Still*, 1951; *It Came from Outer Space*, 1953), etc. Theremin amateur societies are abundant today [22–24], thereminvoxes

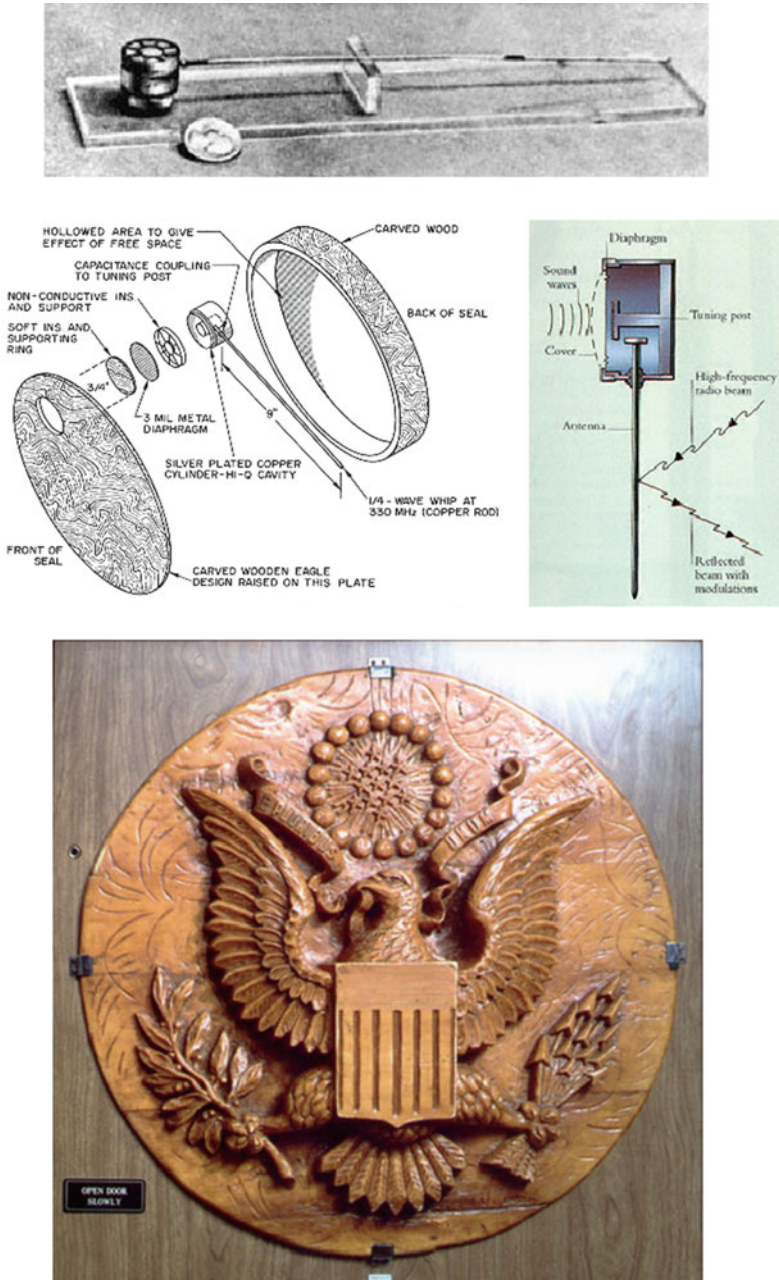


Fig. 2 The great seal bug: the antenna with the resonator (top), principle of operation (middle), the US Great Seal with the embedded device (bottom) (courtesy Wikimedia foundation)

continue to be built [25,26], patents based on original Theremin idea continue to be filed and issued [27], and even Thereminist robots are being developed [28]. There, of course, exists a variety of other electronic musical instruments, both thereminvox-based [29–31] and using other principles [32–34].

Leon Theremin once said in an interview about thereminvox: “I conceived of an instrument that would create sound without using any mechanical energy, like the conductor of an orchestra. The orchestra plays mechanically, using mechanical energy; the conductor just moves his hands, and his movements have an effect on the music artistry” [35]. The near-field nature of the original thereminvox construction required a player to be in the direct vicinity of volume- and pitch-controlling antennas (it is a form of electric field sensing [36]). Interestingly enough, current UHF RFID technology presents another way to realize Leon Theremin’s vision of creating a contactless musical instrument which can be played remotely, like a conductor guiding an orchestra.

The main concept of this instrument (we call it “RFID-vox”) is illustrated in Fig. 3. A person plays it by moving one or more tags (passive or semi-passive) in the far field of the reader antenna system. Tag signals from the RFID reader control the electronic sound characteristics, such as volume and pitch, like in a classical thereminvox. The reader can provide to the musical controller either analog or digital signals, or both. Examples of analog tag signals can be the received signal strength (RSSI) and the phase of backscattered tag signal, readily available [37,38] in commercial Gen2 (ISO-18000 6C [39]) UHF RFID readers. Examples of digital tag signals can be tag ID or data from sensors integrated into the tags. The Gen2 protocol allows for the tag to be read hundreds of times per second. Mapping and linking the received input from each tag to the sounds produced by the instrument is by itself a rich area of computer music research [40,41]. Multiple tags could be used to control many musical parameters or play an orchestra of such instruments. Using the latest passive Gen2 ICs with a sensitivity on the order of -20 dBm [42], such an instrument can be played at a distance of more than 20 m. With the semi-passive (battery-assisted) ICs [43,44], the operating range can be extended even further.

A simple analog version of RFID-vox can easily be realized using a Gen2 reader (which provides tag RSSI and phase readings) and a computer with a sound card. The RFID-vox can be played, for example, using two sticks, similar to conductor’s batons, with embedded dipole-like RFID tags, as shown in Fig. 3. The tag in each baton will have its own ID, which allows one to associate the received RSSI and phase with the particular tag. The received RSSI and phase of the tag signals change with the tag position (in free space, both change monotonically with the distance to the reader antenna) and can be directly linked to control the sounds (volume and pitch). As an alternative, tag location can be calculated using various methods [37,38] and mapped to a virtual piano keyboard space in front of the player. Note that unlike some well-known wireless controls for computer, such as Wii Music [45], these batons can be purely passive devices without any batteries. Of course, playing such an analog instrument will require a certain skill, just as a certain skill is required to play a thereminvox. In the thereminvox, one’s body and hands strongly interact with the near field of thereminvox antennas and affect

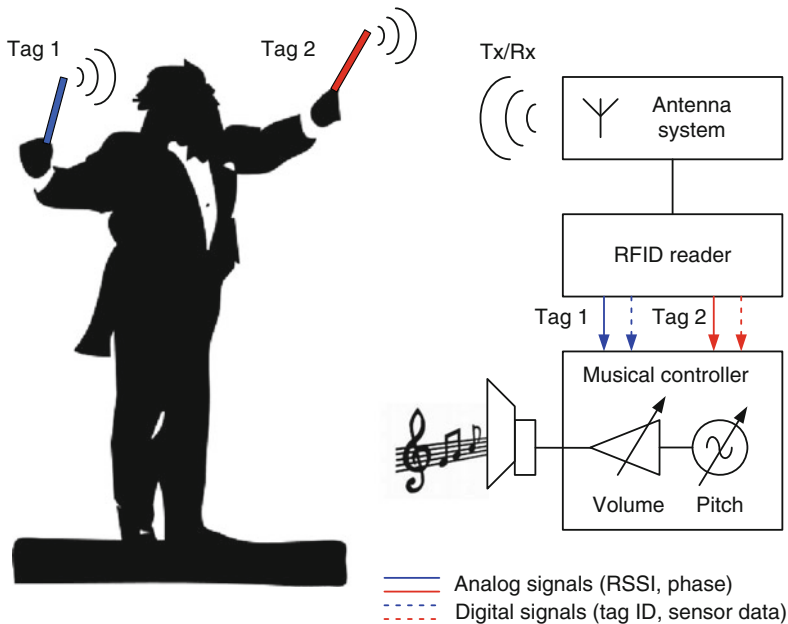


Fig. 3 RFID-vox concept: analog and digital (courtesy Wikimedia foundation for person silhouette)

the way the instrument responds to the player. In the RFID-vox, the interaction happens in the far field, and if the tags are detached from the hands (like in baton sticks), they are not affected by one’s body as much. However, the complexity of the propagation channel (including reflections, polarization mismatch, and other effects) will probably make the task of learning how to play RFID-vox equally challenging.

Digital RFID-vox instruments could also be implemented. The digital identification capability of RFID suggests using different physical objects as input devices, each of which could be associated with a unique sound or instrument. Rather than implementing the sensing in an analog fashion, one could build sensors into the RFID tags and communicate the sensor data as well as the ID digitally.

A digital version of an RFID-vox was realized using our single-bit tilt sensor called “the α -WISP” [46]. An α -WISP consists of two UHF RFID tag ICs multiplexed by tilt-sensing switches to a single antenna. Due to this arrangement, in one tilt state, the α -WISP returns its first ID to the reader; in the other tilt state, it returns a second ID to the reader. If the reader software knows that the two IDs are associated with the same object, then the reader learns one bit of sensor information, encoding the tilt of the object, from each read event.

An α -WISP was affixed to a cup. When the cup transitions from the first tilt state to the second, a musical sound associated with the new tilt state is triggered by the reader. Thus as the object is tipped back and forth, it plays its two sounds.

A second α -WISP, affixed to a second cup, is able to trigger completely different sounds: the unique IDs provided by the RFID tag allow each distinct physical object to be associated with its own characteristic sound set. A video of the α -WISP-based musical instrument is available online [47].

A more sophisticated instrument was created using a more sophisticated RFID sensor tag. Our WISP [48] sensor is a fully passive (RF powered), programmable UHF RFID tag that includes a three-axis accelerometer, as well as other sensors. The sensor values can be mapped to pitch and volume, by analogy with the thereminvox, but these controllers could also be used in countless other ways as well. As a first exploration of such an instrument, we mapped the WISP's tilt to a continuously generated pitch. As the WISP tilts, the pitch being played changes. A video of this experiment is available online [49]. A value thresholding method is used to map one dimension of the acceleration vector produced by the WISP to a musical scale degree (note index in a musical scale). Any particular musical scale (e.g., major or minor) may be selected, or alternately all twelve tones may be used. A small amount of threshold hysteresis makes for cleaner note transitions in the presence of measurement noise. A second dimension of the acceleration vector adjusts the volume of the tones produced.

The relatively low cost of RFID tags could enable many collectible input devices (which might even look like toys rather than musical instruments, as has been described in [50]), each of which could trigger unique musical behaviors. RFID-vox devices could enable new types of musical input devices that could find many applications. For example, novel musical controllers, such as the guitar controller used in the music game *Guitar Hero* [51], combined with proper music mapping algorithms, can also allow non-skilled musicians to enjoy the experience of playing music or used for children's musical education [52].

4 Conclusions

This book chapter was written as a tribute to Leon Theremin, a great inventor and an extraordinary person who lived an amazing life and stood at the roots of electronic musical instruments as well as modulated backscatter technology. Quoting Nick Holonyak, inventor of the light-emitting diode (LED): "Theremin (Lev Termen) should be known for more than just a musical instrument" [53].

It is exciting to see that new technologies (such as Gen2 UHF RFID) combined with old principles (such as the ones used in thereminvox) can open up a window for novel applications, such as wireless contactless musical instruments. As an example of such an application, we described RFID-vox instruments (both analog and digital versions) that can be readily built with existing technology.

References

1. J. Landt, The history of RFID, *IEEE Potentials*, 24(4), 2005, pp. 8–11
2. L. Brown, *A Radar History of World War II: Technical and Military Imperatives*. Institute of Physics Publishing, Bristol, 2000, p. 129
3. H. Stockman, Communication by means of reflected power, *Proceedings of the IRE*, 36, 1948, pp. 1196–1204
4. Wikipedia, Leon Theremin, [Online]. Available: http://en.wikipedia.org/wiki/Leon_Theremin
5. Theremin presents, Ether-wave recital; Russian scientist gives elaborate concert on new instruments of Carnegie Hall, *New York Times*, Apr. 26, 1930. p.18
6. B. Galeyev, Special Section: Leon Theremin, Pioneer of Electronic Art, *Leonardo Music Journal*, MIT, USA, 1996
7. A. Glinsky, *Theremin: Ether Music and Espionage*. University of Illinois Press, Champaign, 2000, pp. 253–254
8. *Theremin - an electronic odyssey*, documentary movie, 1995, director Steven M. Martin
9. K. Skeldon, L. Reid, V. McNally, B. Dougan, C. Fulton, Physics of the Theremin, *American Journal of Physics*, 66(11), 1998, pp. 945–955
10. Music from electricity; German inventor shows device that imitates many instruments, *New York Times*, Sep 14, 1927. p. 24
11. L. S. Theremin, Method of and apparatus for the generation of sounds. US Patent 1661058, issued Feb 28, 1928
12. RCA Theremin [Online]. Available: <http://rcatheremin.com>
13. Theremin's electric symphony, *New York Times*, Mar 27, 1932
14. L. S. Theremin, Signaling apparatus, US Patent 1658953, issued Feb 14, 1928
15. L. S. Theremin, Timing system, US Patent 2047912, issued July 14, 1936
16. V. Shalamov, *Kolyma Tales*, Penguin Books, London, 1994
17. A. Solzhenitsyn, *The Gulag Archipelago (1918–1956)*, Harper Perennial Modern Classic, 2002
18. *Sputnik Biographies – Sergei P. Korolev* [Online]. Available: <http://history.nasa.gov/sputnik/korolev.html>
19. The Great Seal bug [Online]. Available: [http://en.wikipedia.org/wiki/Thing_\(listening_device\)](http://en.wikipedia.org/wiki/Thing_(listening_device))
20. National Cryptologic Museum [Online], Available: http://www.nsa.gov/about/cryptologic_heritage/museum/index.shtml
21. Music: Leon Theremin; inventor of instrument bearing his name is interviewed in the Soviet Union, *New York Times*, Apr 26, 1967. p. 40
22. Theremin World [Online]. Available: www.thereminworld.com
23. Theremin Center for Electroacoustic Music [Online]. Available: <http://theremin.ru>, <http://asmir.theremin.ru>
24. Prometheus Institute [Online], Available: <http://prometheus.kai.ru>
25. B. Colwell, Me and my theremin, *Computer*, 36(2), 2003, pp. 8–9
26. Moog Music Etherwave Theremins [Online]. Available: <http://www.moogmusic.com/theremin/>
27. Visual display for music generated via electric apparatus, US patent 6137042, issued Oct. 24, 2000
28. W. Yan, P. Kuvnichkul, P. Cheung, Y. Demiris, Towards anthropomorphic robot Thereminist, *IEEE International Conference on Robotics and Biomimetics*, 2010, pp.235–240
29. G. Velasquez, The Aria, *WESCON Conference Proceedings*, 1997, pp. 527–530
30. C. Geiger, H. Reckter, D. Paschke, F. Schulz, Evolution of a Theremin-based 3D-interface for music synthesis, *IEEE Symposium on 3D User Interfaces*, 2008, pp. 163–164
31. Tsung-Ching Liu, Shu-Hui Chang, Che-Yi Hsiao, A modified Quad-Theremin for interactive computer music control, *International Conference on Multimedia Technology*, 2011, pp. 6179–6182
32. J. A. Paradiso, Electronic music: new ways to play, *IEEE Spectrum*, 34(12), 1997, pp. 18–30

33. A. Benbasat, J. Paradiso, A compact modular wireless sensor platform, Fourth International Symposium on Information Processing in Sensor Networks, 2005, pp. 410–415
34. M. M. Wanderley, P. Depalle, Gestural control of sound synthesis, Proceedings of the IEEE, 92(4), 2004, pp. 632–644
35. O. Mattis, An interview with Leon Theremin [Online], available: <http://www.thereminvoc.com/story/495/>
36. J. R. Smith, T. White, C. Dodge, J. Paradiso, N. Gershenfeld, D. Allport, Electric field sensing for graphical interfaces, IEEE Computer Graphics and Applications, 18(3), 1998, pp. 54–60
37. R. Miesen, R. Ebel, F. Kirsch, et al., Where is the tag?, IEEE Microwave Magazine, 12(7), 2011, pp. S49–S63
38. P. V. Nikitin, R. Martinez, S. Ramamurthy, H. Leland, G. Spiess, K. V. S. Rao, Phase based spatial identification of UHF RFID tags, IEEE RFID Conference, Apr 2010
39. Electronic Product Code (EPC) Class 1 Gen 2 standard. Available: <http://www.epcglobalinc.org/>
40. B. Miessner, Electronic music and instruments, Proceedings of the IRE, 24(11), 1936, pp. 1427–1463
41. H. Le Caine, Electronic music, Proceedings of the IRE, 44(4), 1956, pp. 457–478
42. Impinj Monza 5 Tag IC. [Online]. Available: http://www.impinj.com/Monza_5_RFID_Tag-Chips.aspx
43. S. Muller, Getting around the technical issues with battery-assisted UHF RFID tags, Wireless Design Magazine, 16(2), 2008, pp. 28–30
44. EM Microelectronic EM4324 RFID IC [Online]. Available: <http://www.emmicroelectronic.com/products.asp?IdProduct=284>
45. Wii Music [Online], available: http://en.wikipedia.org/wiki/Wii_Music
46. M. Philipose, J. Smith, B. Jiang, A. Mamishev, S. Roy, K. Sundara-Rajan, Battery-free wireless identification and sensing, IEEE Pervasive Computing, 4(1), 2005, pp. 37–45
47. J. R. Smith, Alphawisp – UHF RFID musical controller, <http://www.youtube.com/watch?v=mRc3peXsR8o>
48. A. Sample, D. Yeager, P. Powlledge, A. Mamishev, J. R. Smith, Design of an RFID-based battery-free programmable sensing platform, IEEE Transactions on Instrumentation and Measurement, 57(11), 2008, pp. 2608–2615
49. A. N. Parks, J. R. Smith, RFID-Vox: a UHF RFID digital Theremin, <http://www.youtube.com/watch?v=laTsFCsMM7kA>
50. K. Hsiao, J. Paradiso, A new continuous multimodal musical controller using wireless magnetic tags, Proceedings of the International Computer Music Conference, Oct 1999, pp. 24–27
51. Guitar Hero [Online]. Available: [http://en.wikipedia.org/wiki/Guitar_Hero_\(series\)](http://en.wikipedia.org/wiki/Guitar_Hero_(series))
52. J. Rudi, Computer music composition for children, IEEE Signal Processing Magazine, 24(2), 2007, pp. 140–143
53. N. Holonyak, Theremin oscillators and oscillations, American Journal of Physics, 67(5), 1999, p. 369

Index

A

- ACK, 169, 194
- Activity inference, 140
- ADC. *See* Analog to digital converter (ADC)
- Advanced encryption standard (AES), 22, 138, 139, 179, 190–195, 197–200, 203
- AES. *See* Advanced encryption standard (AES)
- AFE. *See* Analog front end (AFE)
- Amplitude shift keying (ASK), 34, 37, 40, 80, 84, 87–89, 91–93, 103–105, 135, 176
- Analog front end (AFE), 16, 20, 35–37, 45, 53, 54, 72, 80, 82, 87–93, 96, 97, 100, 101, 103–105, 107, 176, 177
- Analog to digital converter (ADC), 33, 38, 47, 58–62, 64, 67, 72, 73, 147, 150, 162, 164, 166, 178, 191, 220, 230, 232
- ASIC, 97, 101, 179
- ASK. *See* Amplitude shift keying (ASK)

B

- Backscatter-Anything-to-Tag (BAT), 25–26, 131–141
- BAT. *See* Backscatter-Anything-to-Tag (BAT)
- Battery/batteries, 5, 11, 14, 15, 18, 23, 39, 48, 50, 52, 55, 57, 59, 74, 76, 79–107, 113, 117–120, 122, 131, 143, 147, 176, 179, 223, 225, 226, 234, 235, 237, 238, 242–245, 253, 254, 264
- 16-Bit random number (RN16), 146–147, 149–151, 154, 169, 194
- BlockWrite, 136–137

C

- Capacitive sensor/capacitive sensing/capacitance sensor, 21, 53, 54, 260
- Chopper-stabilized amplifier, 61, 68–70
- Computational RFID (CRFID), 11, 20–21, 26, 131, 132, 134, 136–141, 175
- CRFID. *See* Computational RFID (CRFID)
- Cyber-physical systems, 143

D

- Data MULE, 118
- Demodulator, 16, 34, 37–38, 40, 55, 65–66, 70, 72, 80, 87–89, 91–92, 103, 121, 124, 138, 148–149, 158, 160, 163, 164, 176–178, 211, 220
- DH. *See* Diffie–Hellman (DH)
- Diffie–Hellman (DH), 139, 140
- Duty cycle, 34, 42–44, 46, 113, 181, 182, 184, 213

E

- ECC. *See* Elliptic curve cryptography (ECC)
- EEG. *See* Electroencephalogram (EEG)
- EEPROM, 35, 52, 68, 80, 87, 93, 97, 179
- Electroencephalogram (EEG), 11, 61
- Electromagnetic interference (EMI), 61, 72
- Electromyograms (EMG), 11, 61, 76
- Elliptic curve cryptography (ECC), 21, 22, 140
- EMG. *See* Electromyograms (EMG)
- EMI. *See* Electromagnetic interference (EMI)
- Energy consumption, 44–45, 114, 117, 119, 121, 122, 128, 149, 154, 158, 169, 182

Energy efficiency, 3–7, 10, 11, 113, 114,
116–119, 121, 128, 208
Energy efficiency scaling, 3, 4,
8, 11
Energy scaling, 6, 10

F

FCC. *See* Federal Communications
Commission (FCC)
Federal Communications Commission (FCC),
7, 146, 157, 207, 213, 236
Field-programmable gate array (FPGA), 20,
72, 96–98, 100–101, 103, 105
Flash, 121, 183–185
FM0, 71
FPGA. *See* Field-programmable gate array
(FPGA)
Frame, 19, 37–38, 133, 134, 136–138, 146,
147, 149–151, 153, 155, 162
Friis/Friis transmission formula, 7, 10, 24, 39,
42, 157, 183, 225

G

Gen 1/EPC C1G1/EPC Class 1 Generation
1/EPC Class 1 Gen 1, 14–17, 45,
179
Gen 2/EPC C1G2/EPC Class 1 Generation
2/EPC Class 1 Gen 2, 14, 17, 20, 22,
37, 39, 40, 44, 48, 52, 54, 55, 57–76,
80, 84, 85, 95–99, 106, 107, 119,
131–133, 136, 137, 143–155, 157,
159, 160, 169, 170, 176, 177, 179,
191, 192, 194, 198, 217, 218, 221,
254, 264, 266

I

ID modulation, 15–17
Input impedance, 36, 69–70, 80, 81, 83, 89, 90,
92, 93, 99–100
Instructions per micro
Joule, 4
Intelligent transportation system (ITS),
118
ITS. *See* Intelligent transportation system
(ITS)

L

LO. *See* Local oscillator (LO)
Local oscillator (LO), 60

M

MAC. *See* Medium access control and Message
authentication code (MAC)
MCU. *See* Microcontroller (MCU)
Medium access control and Message
authentication code (MAC), 17, 33,
113, 128, 135, 138, 143–155, 169,
221
Microcontroller (MCU), 5–6, 10, 11, 15,
16, 20–21, 33–35, 38, 39, 42, 44,
49, 55, 59, 113, 117, 119, 121,
131, 138–139, 147, 158–164, 167,
175–178, 180, 181, 183–185, 189,
191, 200, 228–232, 234, 242
Miller, 71, 135, 166
Modulator, 15–18, 34, 37–38, 66, 69, 70, 73,
76, 80, 81, 84–88, 92, 93, 103, 105,
116, 119–121, 135, 145, 148, 149,
158, 163, 164, 176–178, 191, 208,
211, 218, 220, 259, 262, 266
Moo, 20–21, 135, 137–139
Mote, 117, 119–124, 127, 147, 148
Moth, 20, 74, 75
MSP430, 6, 15–17, 20–21, 33, 34, 37–42,
44, 47, 121, 135, 137, 138, 140,
147, 163, 164, 176–178, 184, 191,
199–201, 228–230
Muscles, 20, 72–74

N

Neutrino, 251–258
Non volatile memory (NVM), 23, 52, 59, 62,
68
NVM. *See* Non volatile memory (NVM)

P

Passive data logger (PDL), 17–18, 51, 52
Payload, 134–138, 144, 228
PDL. *See* Passive data logger (PDL)
Public key, 138, 140

Q

Query, 21, 34, 41, 42, 44–46, 62, 71, 80, 134,
146, 147, 149–151, 154, 164, 165,
169, 182, 194, 208, 218
QueryRepeat, 147, 149, 151, 154

R

RC5, 20, 22, 51, 176, 178, 180–182, 184
Reciprocity, 54

Rectifier/rectification, 7, 9, 15, 24, 34–36,
41–45, 47, 49, 54, 59, 61, 62, 65, 66,
73, 80–82, 87, 88, 91, 92, 176, 177,
209, 211, 223–225, 227–229, 231,
237, 244

Reflection coefficient, 66, 81, 119–120, 262

RF harvesting, 15, 18–19, 24–26, 34, 143, 145,
176, 208, 223–231, 233

RN16. *See* 16-bit random number (RN16)

S

S11, 193

S21, 240, 243–245

Scattering parameters (S-parameters),
240

Schottky diode/Schottky, 15, 16, 36, 73, 89,
97, 115, 116, 177

Shared key, 138–140

Sleep, 16, 34, 36, 38, 39, 41–44, 73, 113, 114,
119, 121–123, 144, 148, 149, 154,
168–169, 191, 230

Smart dust, 143

Solar harvesting, 18–19

Supercapacitor, 35, 52

T

Tag ID, 16, 191, 192, 264, 265

Tari, 137–138

Thermocouple, 61, 68, 73, 74, 76

Tmote/Sky mote, 119–122

V

Voltage regulator/voltage regulation/LDO/low
drop-out regulator, 7, 38, 54, 64, 87,
88, 96, 97, 99–101, 181, 221, 229

W

Wake/wake-up/wake-up radio, 16, 23, 38,
113–128, 144, 148, 149, 168–169,
182, 191, 230

Wireless identification and sensing platform
(WISP)

accelerometer WISP, 6, 14, 15, 17, 21,
26, 34–35, 47–51, 147, 169, 254,
266

α WISP, 14

data logger WISP/WISP-DL/WISP-PDL,
17–18, 35, 51, 52, 137

LED WISP, 19, 35

π -WISP, 15

SOCWISP, 20, 179

Solar WISP, 18–19

WISP. *See* Wireless identification and sensing
platform (WISP)

WISP-Mote, 117, 119–124, 127, 147,
148

Workload, 3–9, 11, 26, 151, 180, 184

Z

ZebraNet, 118