Power, Communication and Sensing Solutions for Energy Constrained Platforms

Vamsi Talla

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2016

Reading Committee:

Joshua R. Smith, Chair

Shyamnath Gollakota

Matthew S. Reynolds

Program Authorized to Offer Degree:
Electrical Engineering

University of Washington

Abstract

Power, Communication and Sensing Solutions for Energy Constrained Platforms

Vamsi Talla

Chair of the Supervisory Committee:
Associate Professor Joshua R. Smith
Computer Science & Engineering and Electrical Engineering

We live in a world where mobile devices such as smartphones, smart watches and tablets are commonplace. With rapid strides in technology, we have come a long way from the first main frame computer ENIAC, which occupied 167m^2 and consumed 150 kW of electricity. But, we cannot stand still. We are still ways off from the vision of ubiquitous computing where, devices permeate our surroundings and function perpetually, without the need for maintenance or any user interference. Our current devices are either tethered to power cords or use batteries which require constant supervision and maintenance.

In this work we introduce power, communication and sensing technologies to help us achieve the vision of ubiquitous computing. We note that batteries are too restricting for a large number of applications. First, we present RF energy harvesting solutions to replace batteries with power harvested from ambient RF signals. We show that we can use ambient TV and RFID signals to power computing and sensing devices. Next we show that we can transform a Wi-Fi router, a ubiquitous part of the wireless infrastructure into a source of far field wireless power, but without significantly compromising the performance of the Wi-Fi communication.

For communication, we observe that traditional radio based communication is extremely power hungry which limits the lifespan of the device and makes energy harvesting impractical. We show that, using backscatter communication techniques, we can leverage ambient RF signals such as TV and RFID for power and communication between battery-free devices. This enables ubiquitous

communication where devices can communicate among themselves at unprecedented scales and in locations that were previously inaccessible. Next, to bring the benefits of backscatter communication to mainstream applications, we demonstrate that using backscatter, we can synthesize Wi-Fi packets at 3–4 orders of magnitude lower power than Wi-Fi radios. These Wi-Fi transmissions at 1-11 Mbps data rate can be received on standard Wi-Fi radios, bringing low power connectivity to the Wi-Fi space.

Finally, we demonstrate that backscatter techniques can also be applied to sensing to reduce power consumption. We use analog backscatter to directly transmit sensor information to a reader at *zero* power and combine this technique with digital backscatter to develop a digital addressable battery-free microphone.

We believe the technologies developed in this work will bring up a step closer to a world where we are surrounded by perpetually operating devices.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I wish to express sincere appreciation to the Department of Electrical Engineering and Department of Computer Science and Engineering at the University of Washington. I especially want to thank my advisor, Prof. Josh Smith for the mentorship and giving me the opportunity and the freedom to work on problems that excited me throughout the graduate school. Additionally, I have enjoyed working closely with Shyam and the collaboration with Shyam's group has been an integral part of the research outlined in this dissertation.

The past and present members of the Sensor System Lab and the Mobile Systems Lab have been a great source of collaboration, help, support and mentorship over the past 5 years. Finally, I would be remiss not to thank Julie, our group' administrator for the help and support over the course of past five years. I think it's safe to say without her, we would not have met a majority of our deadlines.

Finally, it goes without saying, but I would like to thank my family and friends for the support during my graduate school experience.

DEDICATION

to my mom, dad and *little* sister

Chapter 1

INTRODUCTION

Tiny sensing and computing devices are permeating our surroundings. This phenomenon known as the internet of things (IoT) or internet of everything (IoE) envisions a world where numerous such devices will enrich our surroundings and become an integral part of our daily lives. We are at the forefront of the IoT era; wearable devices such as Google Glass, Jawbone UP, FitBit, smart watches like Toq, Galaxy Gear, portable medical devices like insulin and blood pressure monitors are becoming commonplace [23, 25, 31, 51, 58]. Increasingly, buildings, furniture and appliances (such as smart TV and smart refrigerator) are being equipped with embedded sensing and computation. Smart home monitoring solutions such as Nest and SmartThings are a few examples of commercially successful devices [38]. And finally implanted medical devices such as pacemakers, cochlear implants, retinal implants and neural implants are becoming commonplace and enabling longer and higher quality life [39, 127].

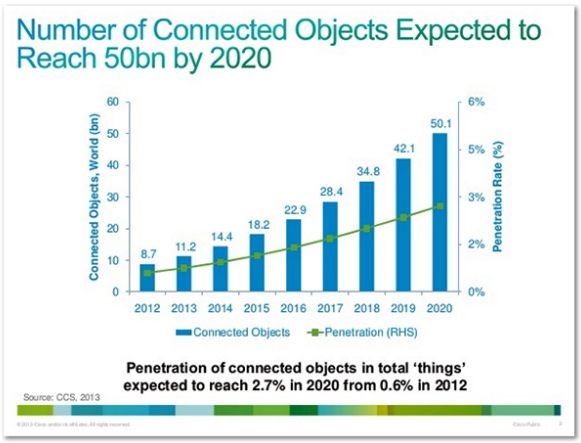

Figure 1.1: Expected number of connected devices by the year 2050. [41]

Cisco® estimates that by 2020, there will be more than 50 billion connected devices (see Fig. 1.1) [41]. This exponential rise in number of devices is a culmination of tremendous progress over the last decades in a wide range of technologies. Digital CMOS scaling, advances in MEMS, increasingly efficient and reliable wireless protocols, advanced chip packaging, printed circuit board and assembly solutions have all contributed to reduction in cost and size which has made IoT feasible and practical (See Fig. 1.3). CMOS technology scaling has resulted in smaller, cheaper and more power efficient silicon chips. Improvement in MEMS fabrication has enabled smaller and power efficient sensors. Advanced packaging solutions have led to integration of multiple sensors, silicon chips, antennas and passive components into a single small form factor solution. As an example, the FitBit Flex contains a vibration motor, accelerometer, battery, two antennas, a Cortex M3 micro-controller and a Bluetooth radio all in a half-dollar coin form factor [29].

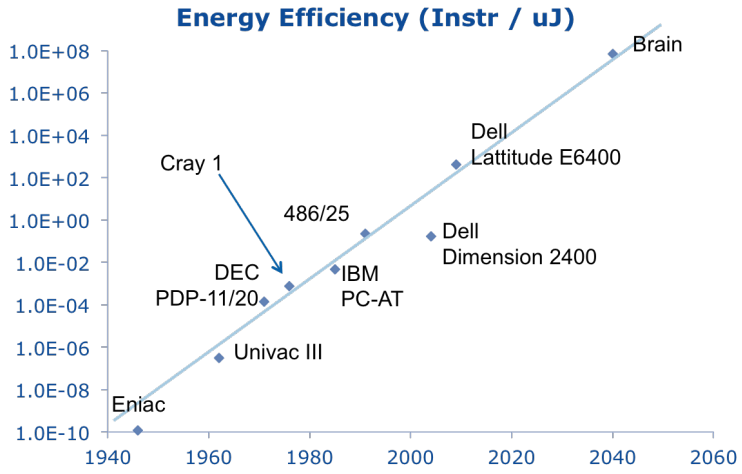

Figure 1.2: **Energy Efficiency of Computation** Historical data of number of instructions per μJ of energy for computational platforms [123].

As we look forward towards the future with a billion connected devices, we are must ask the following question: how are we going to deploy, power and maintain these devices at such large scales? Wires are often not feasible and too restricting. Batteries add weight, cost, and require periodic recharging or replacement that adds maintenance cost and are difficult at large scales. To replace batteries with long term sources of power, we must first understand the characteristics of

typical devices. Most devices have three primary functionalities: computation, communication and sensing[1]. The computation core has undergone tremendous improvements over the past 50 years due to CMOS technology scaling (also known as Moore's law). In addition to reducing cost and size, CMOS scaling has reduced the energy consumption of computation. As a result, in typical IoT devices, computation has a tiny power footprint. However, unlike computation, communication front end is primarily analog in nature and analog circuits do not scale with Moore's law. The power associated with active radio based communication is prohibitively expensive and drains the energy storage unit such as a battery or capacitor, limiting its lifespan and making radios impractical for majority of applications. Even the most energy efficient radio protocols consume milliwatts of power [56, 93] making communication, the major power draw in most devices [65, 143]. To enable ubiquitous connectivity, there is a need for ultra-low power communication protocols.

Lastly, with advancements in MEMS fabrication and packaging technology, sensors have become smaller, more power efficient and easier to integrate. However, despite this progress, the power associated with ADC operation i.e. digitization of data and, transmission of digital sensor data is prohibitively high. In energy constrained platforms, this approach requires heavy duty cycling and makes it impractical for high data rate and interactive applications such as audio and vision.

The goal of the thesis is to solve these challenges associated with power, communication and high data rate sensing. Specifically, we discuss how we can:

- Power large number of devices at long ranges
- Enable ultra-low power wireless connectivity
- Design ultra-low power high data rate sensing platforms

A positive answer to these questions would enable the deployment of IoT devices at unprecedented scale. In the rest of the chapter, we give an overview of the techniques which help us achieve these goals.

[1] In some cases, actuation is also included

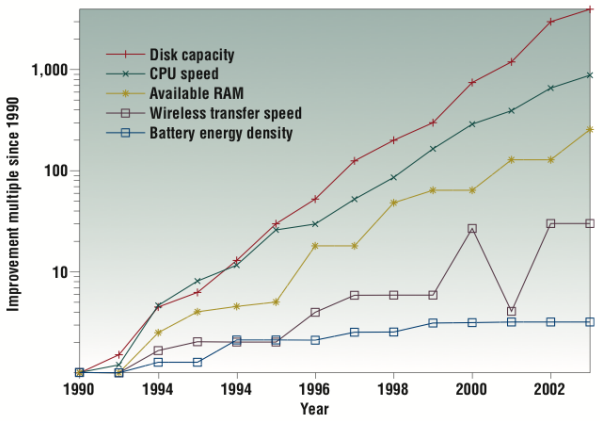

Figure 1.3: Historical trend in improvement in data storage, computation, wireless communication and battery technology [163].

1.1 Powering large number of devices at long ranges

As the number of devices increase and become more widespread, providing uninterrupted power to devices becomes a huge challenge. Wires are often not feasible and too restricting. Batteries are the most popular choice, but they add weight, cost, and require periodic recharging or replacement that increases maintenance cost and are difficult at large scales. But unfortunately, unlike electronics which has become smaller and more efficient over time, batteries haven't followed the same trend. Fig. 1.3 shows the historical trend of different technologies and as the plot indicates, battery technology is the slowest trend in mobile computing. Battery technology hasn't seen the exponential improvements we associate with semiconductors, in fact, the battery energy density has reached a saturation point [163].

Clearly, there is a need for uninterrupted sources of power which can either replace or at least supplement batteries. Due to reduction in power consumption of computation, we are at a stage where devices can be powered using energy from harvesting sources. Fig. 1.2 plots the historical trends in power consumption (number of instructions per μJ of energy) of computation. The energy efficiency of digital computation i.e. number of instructions per μJ of energy have doubled every 1.57 years. This is a consequence of CMOS technology scaling also known as Moore's law.

Moore's law states that the number of transistors in a given silicon area double every one and half years. This is achieved by shrinking the feature size of the transistor, which reduces the gate capacitance of MOS transistors and consequently reduces switching energy (CV^2). In summary, in addition to reducing cost and size, CMOS scaling also reduces the energy consumption of computation. As a result, the power requirements for digital computational have drastically reduced and now the systems can be designed to consume 10-100 μWs of power during active operation [36].[2]

Given this power budget, it's now practical to use solar, vibration, thermal and far field RF harvesting solutions to power tiny sensing and computational platforms and replace or supplement batteries [67,125,163,166,178,179,186]. In this thesis, the focus will be on harvesting energy from RF sources due to it's advantages over other harvesting source. All devices require an antenna for communication, and in principle, the same antenna can be reused to harvest incoming RF signals and have negligible impact on the form factor of the device. On the other hand, solar, thermal and motion based harvesting need additional transducers, which increase size, cost and complexity of the system [163, 179]. Unlike RF sources, available power from solar, thermal and motion based harvesting is inconsistent and unpredictable due to the dependence on external factors such as light, temperature and motion respectively. All these factors make RF signals an attractive source for harvesting power.

We are constantly surrounded by RF signals such as cellular, TV and Wi-Fi, which are all potential means to power devices and make them truly pervasive [131, 166, 196]. Due to their broadcast nature, RF signals can enable power delivery at long distances and can simultaneously power multiple devices. As a first step towards achieving the goal of ubiquitous power delivery, Chapters 2, 3 and 6 illustrate the use of ambient TV signals and UHF RFID signals to power devices outdoors and indoors respectively. The chapters demonstrate how RF harvesting solutions can be leveraged to design and develop battery-free computation, sensing and communication platforms.

Although a great start, it turns out that in typical indoor environments, there is a paucity of harvestable TV, cellular and RFID signals. TV signals do not penetrate buildings, cellular signals

[2] As mentioned earlier, radio based communication is the major power draw in most systems. In §1.2 we show how we can replace radios with backscatter communication which consumes μWatts of power.

are too weak and RFID reader infrastructure is too expensive to setup. As a result, long range wireless power delivery is still a pipe dream for a majority of applications. However, indoor environments such as homes, offices and public spaces have extensive Wi-Fi connectivity. Wi-Fi is hugely popular and due to economies of scale, there has been a steady decline in the cost of Wi-Fi chipsets. Repurposing existing Wi-Fi networks for power delivery can ease the deployment of RF-powered devices without additional power infrastructure. In fact, it is not unreasonable to expect that if better Wi-Fi coverage is required, additional Wi-Fi access points can be deployed with minimal effort and cost overhead. Hence, Wi-Fi presents a unique opportunity to wirelessly power devices at large scales in indoor environments. However, it turns out there is a fundamental tradeoff between wireless power delivery and wireless communication. In Chapter 4, we introduce Power Over Wi-Fi which introduces techniques to enable a standard 802.11 Wi-Fi router to provide far field wireless power without significantly compromising the network's communication performance. Using this technique, we prototype and power battery free temperature sensor, cameras and battery recharging devices from standard compliant Wi-Fi radio chipsets.

1.2 Enabling ultra-low power wireless connectivity

Conventional active radio based communication is extremely power intensive. To understand the reason behind this, consider a typical radio deployment shown in Fig. 1.5. On a high level, the design choice and architecture of radio based communication has remained more or less stagnant over the past 30 years. Radio links consist of two nodes communicating with each other with similar architectures, complexity and power profile. Each node has an RF transmitter chain of Digital to Analog converter, an up conversion mixer and a power amplifier. Similarly, the RF receiver has a low noise amplifier, followed by a down conversion mixer and an Analog to Digital converter. The RF carrier is generated using a frequency synthesizer. All these components operate at RF frequencies and, are primarily analog in nature and unlike digital circuits, do not scale with Moore's law. In summary, traditional radios are based on a symmetric architecture where both nodes operate a power hungry RF front end. This architecture is unsuitable for energy constrained systems.

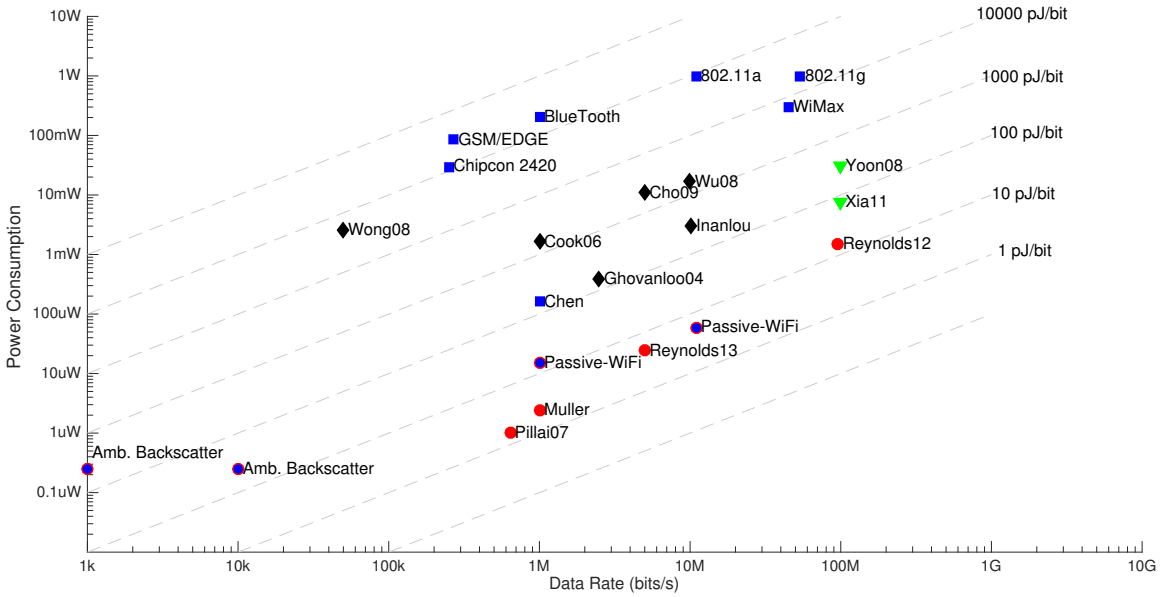

Figure 1.4: **Energy Efficiency of Communication:** An overview of the energy efficiency of wireless communication systems from literature. The bottom right corner of the graph represents the highest energy efficient performance (a combination of low power and high data rate). The circles represent backscatter communication systems.

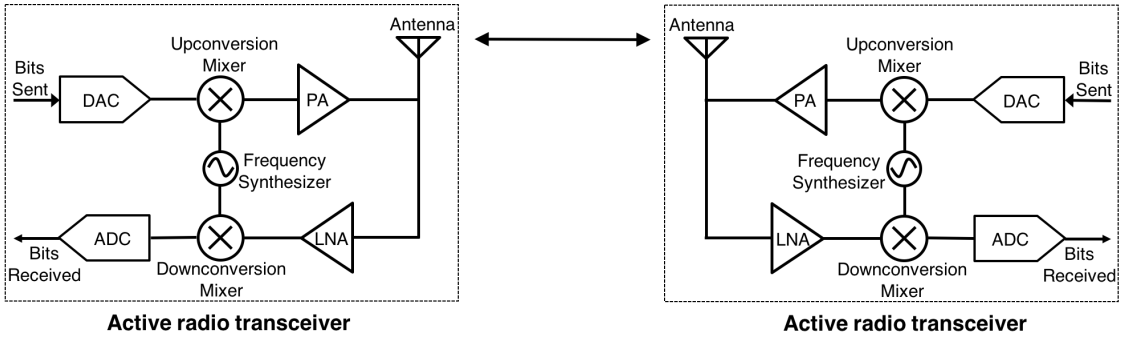

Figure 1.5: **Active radios based communication.** Two active radios communicate with each other using conventional means.

The key to reducing the power consumption of communication is to break this symmetry by eliminating the power hungry RF front end from one node i.e. the energy constrained device. Instead of every device in a network being identical, we delegate the complex and power hungry transceiver to a single device called the reader. Other devices in the network called the tags have

no RF transceiver and use *backscatter* for communication with the reader. The reader transmits an RF carrier wave and the backscatter tags use a switch to toggle the impedance of the antenna between different states which modulates the signal reflected by the tag's antenna. The reader uses a full duplex radio to suppress self-interference and demodulate the reflected RF signals to decode bits transmitted by the tag. Since backscatter systems only require a switch (typically a digital FET), they are orders of magnitude more energy efficiency than active radio communication. The topology of a system with a powered reader and a battery-free/energy constrained backscatter tag is shown in Fig. 1.6. Fig. 1.4 shows a comparison of the energy efficiency of various communication protocols. It can be seen that the red circles representing backscatter communication are orders of magnitude more energy efficient than active radios such as Wi-Fi, ZigBee and Bluetooth.

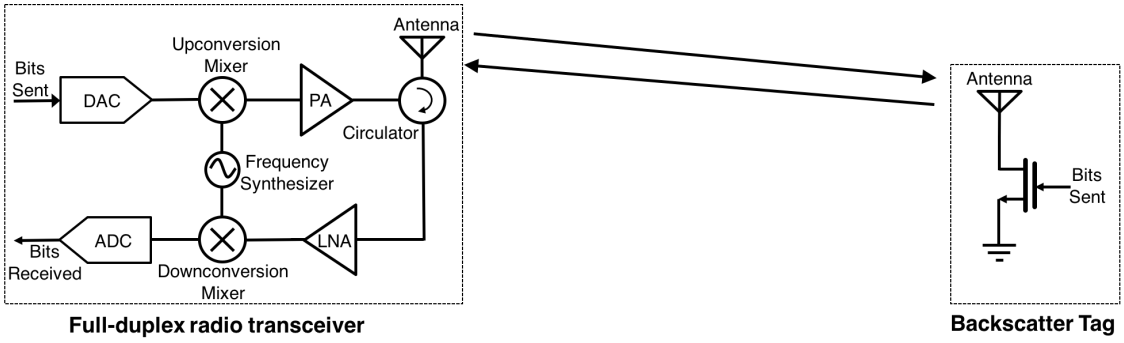

Figure 1.6: **Communication between a reader and backscatter tag.** The reader transmits a continuous wave and the backscatter tag selectively controls the signal reflected by it's antenna to send data back to the reader.

However, the need for special purpose hardware such as (RFID) readers limits the usability of backscatter systems. RFID readers use full duplex radio architecture and consequently are expensive and hard to deploy. Instead, if can use existing infrastructure and protocols such as Wi-Fi, we can bring the benefits of backscatter communication to a wide range of applications. In Chapter 5, we introduce Passive Wi-Fi where backscatter techniques are used to synthesize Wi-Fi packets which can be decoded on standard Wi-Fi routers and smartphones. Fig. 1.7 shows the network topology where a single transmitter generates a continuous wave carrier signal and backscatter tags backscatter this continuous wave signal to synthesize Wi-Fi packets which can be decoded on standard Wi-Fi receivers. We eliminate the need for a full duplex radio by physically

separating the continuous wave generator and the receiver. This technique requires a deployment of a single continuous wave generator and can enable devices to transmit 1-11 Mbps Wi-Fi packets using only 15-60 μWs. A power comparison of Passive Wi-Fi to traditional Wi-Fi radios is shown in Fig. 1.4.

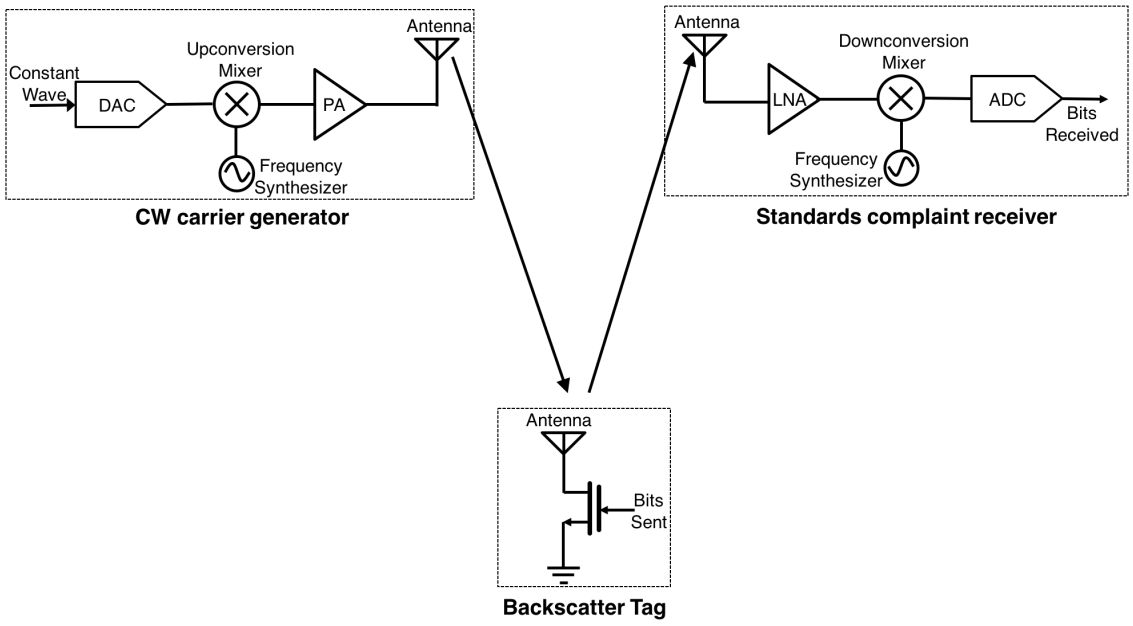

Figure 1.7: **Backscatter communication using a dedicated RF generator and standard receiver.** The dedicated RF signal generator transmits a continuous wave and the tag transmits information by selectively reflecting the signal to synthesize standard complaint digital packets which can be decoded on unmodified receivers.

Finally, there are scenarios where the need for any dedicated transmitter or receiver is too limiting. In Chapter 2 and 3 we show that ambient TV signals or ambient RFID transmissions can be backscattered to enable communication between battery-free devices. Fig. 1.8 illustrates the scenario where two battery-free devices communicate with each other using a backscatter transmitter and a passive ultra-low power receiver. Ambient backscatter and full duplex backscatter techniques can enable 100 bps to 10 kbps data rates between battery-free devices with extremely high energy efficiency (See Fig. 1.4).

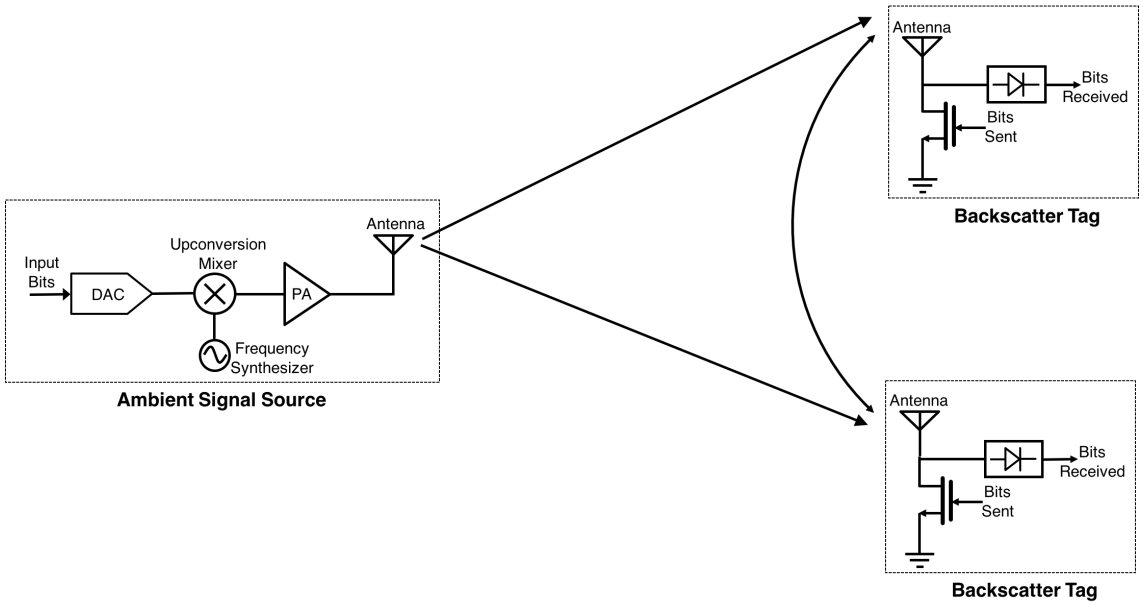

Figure 1.8: **Tag to Tag Communication.** In the presence of ambient RF signals, two battery-free devices communicate by backscattering ambient RF signals and decoding them using a passive receiver.

1.3 Designing ultra-low power high data rate sensing platforms

Finally, sensor technology has rapidly progressed in the last decade. With advancements in MEMS fabrication and packaging technology, sensors are becoming smaller, more power efficient and easier to integrate. As an example, Invensense Nine-axis MEMS motion tracking device integrates a Gyro, accelerometer and compass in a single 3x3x1mm package with ultra low power consumption (8 μA in sleep mode) [30]. The common approach to sensing is to use an ADC to sample the analog data and then transmitted it as digital packets. ADC based sampling in combination with duty cycling has been very successful with slowly varying quantities such as temperature, humidity, pressure, light, acceleration, etc. However, for applications such as audio and vision which require interactivity and high duty cycles, digital approaches are impractical on battery-free and energy constrained platforms.

In Chapter 6, we build on the RF energy harvesting and backscatter communication techniques and develop, a battery-free high data rate acoustic sensing platform. We introduce analog backscatter, a technique to continuously transmit high data rate audio data from a battery-free platform to

a reader. We combine analog backscatter with a digital platform to design, a digitally addressable hybrid analog-digital battery-free sensing platform.

1.4 Contributions and Organization of the Thesis

In this thesis, we present power, communication and sensing solutions for energy constrained platforms. We introduce novel hardware techniques and build on this hardware to design the network stack and application layer and develop system level solutions. The outline of the thesis is as follows.

- Chapter 2 introduces Ambient Backscatter, a new communication primitive where devices communicate by backscattering ambient RF signals. The chapter describes the concept of ambient backscatter, the hardware design, the network stack to enable battery-free devices to share the wireless channel and sample applications to illustrate the use cases for ambient backscatter communication.

- Chapter 3 builds on ambient backscatter and introduces full-duplex backscatter, which enables instantaneous feedback between two battery free devices with no additional power penalty. We describe the hardware design of the full-duplex prototype and use the feedback channel to design a network stack which minimizes the energy wastes that occur due to collisions and packet errors making deployment of battery-free systems robust and practical.

- Chapter 4 presents the first *power over Wi-Fi* system that delivers power to low-power sensors and devices using existing Wi-Fi chipsets. First we describe techniques which enables the Wi-Fi router, a ubiquitous part of wireless communication infrastructure, to provide far field wireless power without significantly compromising the network's communication performance. Building on this design, we prototype battery-free temperature sensors, camera and battery charging prototypes which are all powered using 802.11 Wi-Fi chipsets.

- Chapter 5 discusses Passive Wi-Fi, the first system to demonstrate that 802.11b Wi-Fi packets can be synthesized using backscatter communication while consuming 3–4 orders of magni-

tude lower power than traditional Wi-Fi chipsets. Passive Wi-Fi transmissions can be decoded on any standard complaint Wi-Fi device. The chapter describes the principle behind Passive Wi-Fi, as well we the hardware design and network stack design that enables Passive Wi-Fi transmitters to coexist with other devices in the ISM band, without incurring the power consumption of carrier sense and medium access control operations.

- Chapter 6 introduces hybrid analog-digital backscatter sensing mechanism which enables battery-free high data rate acoustic sensing. We describe the principle behind analog backscatter, analyze, optimize and integrate analog backscatter with digital platform.

- Chapter 7 concludes the thesis and outlines some areas of future research.

Chapter 2

AMBIENT BACKSCATTER COMMUNICATION

2.1 Introduction

Small computing devices are increasingly embedded in objects and environments such as thermostats, books, furniture, and even implantable medical devices [127, 142, 156]. A key issue is how to power these devices as they become smaller and numerous; wires are often not feasible, and batteries add weight, bulk, cost, and require recharging or replacement that adds maintenance cost and is difficult at large scales [203].

In this chapter we describe how we can enable devices to communicate using *ambient RF signals* as the only source of power? Ambient RF from TV and cellular communications is widely available in urban areas (day and night, indoors and outdoors). Further, recent work has shown that one can harvest tens to hundreds of microwatts from ambient RF signals [166, 186]. Thus, this system would enable ubiquitous communication at unprecedented scales and in locations that were previously inaccessible.

Designing such systems, however, is challenging as the simple act of generating a conventional radio wave typically requires much more power than can be harvested from ambient RF signals [166]. We introduce *ambient backscatter*, a novel communication mechanism that enables devices to communicate by *backscattering* ambient RF. In traditional backscatter communication (e.g., RFID), a device communicates by modulating its reflections of an incident RF signal (and not by generating radio waves). Hence, it is orders of magnitude more energy-efficient than conventional radio communication [77].

Ambient backscatter differs from RFID-style backscatter in three key respects. Firstly, it takes advantage of existing RF signals so it does not require the deployment of a special-purpose power infrastructure—like an RFID reader—to transmit a high-power (1W) signal to nearby devices. This

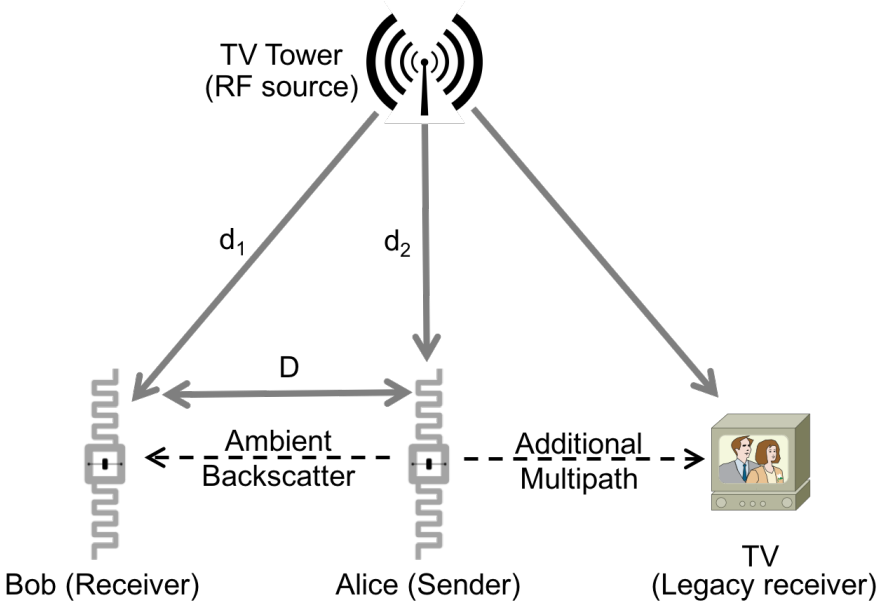

Figure 2.1: **Ambient Backscatter:** Communication between two battery-free devices. One such device, Alice, can backscatter ambient signals that can be decoded by other ambient backscatter devices. To legacy receivers, this signal is simply an additional source of multi-path, and they can still decode the original transmission.

avoids installation and maintenance costs that may make such a system impractical, especially if the environment is outdoors or spans a large area. Second, and related, it has a very small environmental footprint because no additional energy is consumed beyond that which is already in the air. Finally, ambient backscatter provides device-to-device communication. This is unlike traditional RFID systems in which tags must talk exclusively to an RFID reader and are unable to even sense the transmissions of other nearby tags.

To understand how ambient backscatter works, consider two nearby battery-free devices, Alice and Bob, and a TV tower in a metropolitan area as the ambient source, as shown in Fig. 2.1. Suppose Alice wants to send a packet to Bob. To do so, Alice backscatters the ambient signals to convey the bits in the packet—she can indicate either a '0' or a '1' bit by switching her antenna between reflecting and non-reflecting states. The signals that are reflected by Alice effectively create an additional path from the TV tower to Bob and other nearby receivers. Tv and cellular use radio based receivers and are designed to compensate for multi-path wireless channels, and can potentially account for the additional path. Bob, on the other hand, can sense the signal changes

caused by the backscattering, and decode Alice's packet.

Designing an ambient backscatter system is challenging for at least three reasons.

- Backscattered signals are weak and traditional backscatter uses a constant signal [153] to facilitate the detection of small level changes. Ambient backscatter uses uncontrollable RF signals that already have information encoded in them. Hence it requires a different mechanism to extract the backscattered information.

- Traditional backscatter receivers such as RFID readers use coherent radio based architecture which use power-hungry components such as oscillators and ADCs and decode the signal with relatively complex digital signal processing techniques. These techniques are not practical for use on a battery-free device.

- Ambient backscatter lacks a centralized controller such as an RFID reader to coordinate all communications. Thus, it must operate a distributed multiple access protocol and develop functionalities like carrier sense that are not available in traditional backscattering devices.

Our approach is to co-design the hardware elements for ambient backscatter along with the layers in the network stack that make use of it. The key insight we use to decode transmissions is that there is a large difference in the information transfer rates of the ambient RF signal and backscattered signal. This difference allows for the separation of these signals using only low-power analog filtering operations which can be implemented using resistors, capacitors and comparators. Similarily, we are able to realize carrier sense and framing operations with low-power components based on the physical properties of ambient backscatter signals. This in turn lets us synthesize network protocols for coordinating multiple such devices.

To show the feasibility of our ideas, we have built a hardware prototype, shown in Fig. 2.2, that is approximately the size of a credit card.[1] Our prototype includes a power harvester for TV signals and backscatter which are tuned to communicate using UHF TV signals centered at 539 MHz. The

[1]We design and build a COTS (commercial off-the-shelf) prototype. An integrated circuit would achieve better results and be of an much smaller form factor (down to 1 mm^2 plus the antenna).

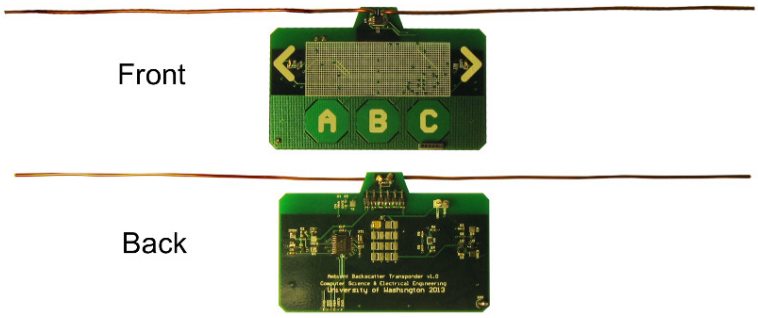

Figure 2.2: **Ambient Backscatter Prototype:** A photo of a PCB prototype that can harvest, transmit and receive without needing a battery or powered reader. It also includes touch sensors (the A, B and C buttons), LEDs (placed near the two arrows) and programmable microcontroller that operate using harvested energy.

harvested energy is used to provide the small amounts of power required for ambient backscatter and to run the microcontroller and the on-board sensors. Our prototype also includes a low-power flashing LED and capacitive touch sensor for use by applications.

We experiment with two proof-of-concept applications that show the potential of ambient backscatter in achieving ubiquitous communication. The first application is a bus pass that can also transfer money to other cards anywhere, at any time. When a user swipes the touch sensor in the presence of another card, it transmits the current balance stored in the microcontroller and confirms the transaction by flashing the LED. The second is a grocery store application where an item tag can tell when an item is placed in a wrong shelf. We ask 10 tags to verify that they do not contain a misplaced tag and flash the LED when they do.

The system is evaluated in both indoor and outdoor scenarios and at varying distances between the transmitter and receiver. To account for multi-path effects, we repeat our measurements with slight perturbations of the receiver position for a total of 1020 measurements. Results show that our prototypes can achieve an information rate of 1 kbps between two ambient backscattering devices, at distances of up to 2.5 feet in outdoor locations and 1.5 feet in indoor locations. Furthermore, we test a variety of locations and show that our end-to-end system (which includes communication, an LED, touch sensors and a general-purpose microcontroller) is able to operate battery-free at distances of up to 6.5 miles from the TV tower. Finally, we test the interference of ambient

backscattering and find that, even in less favorable conditions, it does not create any noticeable glitches on an off-the-shelf TV, as long as the device is more than 7.2 inches away from the TV antenna.[2]

Our Contributions: We make the following contributions:

- We introduce ambient backscatter, the first wireless primitive to let devices communicate without either requiring them to generate RF signals (as in conventional communications) or reflect signals from a dedicated powered reader (as in RFID).

- We develop a network stack that enables multiple ambient backscattering devices to co-exist. Specifically, we show how to perform energy detection without the ability to directly measure the energy on the medium and hence enable carrier sense.

- We present designs and a prototype which show how all of the above, from ambient backscatter through to the multi-access protocols of our network, can be implemented on ultra-low-power devices using simple analog components.

While the performance of our prototype is a modest start, we hope that the techniques we present will help realize ubiquitous communication, and allow computing devices embedded into the physical world to communicate amongst themselves at an unprecedented scale.

2.2 Background on TV Transmissions

In principle, ambient backscatter is a general technique that can leverage RF signals including TV, radio and cellular transmissions. In this paper we have chosen to focus on demonstrating the feasibility of ambient backscatter of signals from TV broadcast sources.

TV towers transmit up to 1 MW effective radiated power (ERP) and can serve locations more than 100 mi away from the tower in very flat terrain and up to 45 mi in denser terrain [2]. The coverage of these signals is excellent, particularly in urban areas with the top four broadcast TV

[2] At such close distances, it is in the near-field of the TV antenna.

channels in America reaching 97% of households and the average American household receiving 17 broadcast TV stations [11]. It is this pervasive nature of TV signals that make them attractive for use in our first ambient backscatter prototype.

There are currently three main TV standards that are used around the world: ATSC (N. America and S. Korea), DVB-T (Europe, Australia, New Zealand, etc.) and ISDB-T (Japan, most of S. America) [19]. While our prototype targets ATSC transmissions, our method for communicating using ambient signals leverages the following properties of TV signals that hold across all standards:

Firstly, TV towers broadcast uninterrupted, continuous signals at all hours of the day and night. Thus, they provide a reliable source of both power and signal for use in ambient backscatter. Secondly, TV transmissions are amplitude-varying signals that change at a fast rate. For example, in ATSC, which uses an 8-level vestigial sideband (8VSB) modulation to transmit one of eight amplitude values per symbol, symbols are sent over a 6 MHz wideband channel, resulting in a very fast fluctuation in the signal.

Lastly, TV transmissions periodically encode special synchronization symbols that are used by the receiver to compute the multipath channel characteristics [63]. In ATSC, the 8VSB symbols are organized first into data segments of 832 symbols and then fields of 313 segments. Before every data segment, the transmitter sends a data segment sync that consists of four symbols and is intended to help the receiver calibrate the 8VSB amplitude levels. Before every field, the transmitter sends a field sync data segment that is also used by the receiver to compute the channel information. Since ambient backscatter effectively creates additional paths from the transmitter to the TV receiver, the existing ability of TV receivers to account for multipath distortion make them resistant to interference from backscattering devices that operate at a lower rate than these sync segments. We note that the other common TV standard in the world—DVB-T, which uses OFDM modulation—includes cyclic prefixes and guard intervals, and hence has an even higher resistance to multipath distortion compared to the ATSC standard [3].

Legality: In general, it is illegal to broadcast random signals on spectrum reserved for TV (or cellular) channels. However, battery-free backscattering devices (e.g. RFID tags) are unregulated

and not tested by FCC because the emission levels from such devices is very low [40] and because they are only modulating their reflection of a pre-existing signal rather than actively emitting a signal in reserved spectrum. Ambient backscatter also falls into this category, and would therefore be legal under current policies.

In the rest of the chapter, we show how ultra-low-power devices can communicate by backscattering these ambient signals.

2.3 Ambient Backscatter Design

Ambient backscatter is a new form of communication in which devices communicate without the need for any infrastructure such as a dedicated reader. An ambient backscattering device reflects existing RF signals such as broadcast TV or cellular transmissions to communicate. Since the ambient signals are pre-existing, the added cost of such communication is negligible.

Designing such devices, however, is challenging for three main reasons: First, the ambient signals are random and uncontrollable. Thus, we need a mechanism to extract the backscattered information from random ambient signals. Second, the receiver has to decode these signals on a battery-free device which significantly limits the design space by placing a severe constraint on the power requirements of the device. Third, since there is no centralized controller to coordinate communications, these devices need to operate a distributed multiple access protocol and develop functionalities like carrier sense. In the rest of this section, we describe how our design addresses the above challenges.

2.3.1 Overview

Fig. 2.3(a) shows the block diagram of the ambient backscattering device design. The design consists of an RF harvester and power management unit, receiver, backscatter transmitter, digital logic unit and peripherals such as LEDs and user capacitive touch button. The digital control unit is implemented on a programmable microcontroller and is powered by harvested RF energy. The digital control unit implements the backscatter communication protocol and controls the LEDs and

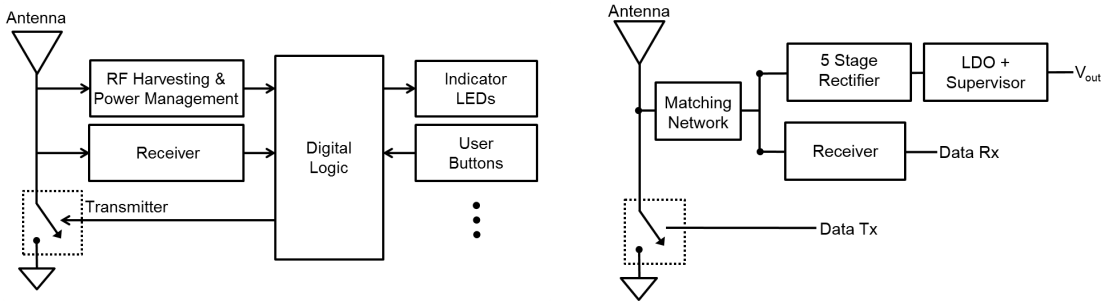

(a) Block diagram of ambient backscattering device. (b) RF front end of ambient backscattering device.

Figure 2.3: **Architecture of the ambient backscattering device.** An ambient backscattering device consists of a power harvester, receiver and transmitter, all connected to a single antenna. These form the RF front end. The digital logic unit is implemented on a microcontroller and runs on harvested power. The digital logic runs the communication protocols and controls peripherals such as LEDs and user touch buttons.

capacitive touch input buttons.

A more detailed configuration of the RF front of the ambient backscatter device shown in Fig. 2.3(b). The backscatter transmitter is directly interfaced with the antenna to maximize the changes in the backscattered signal. The RF energy harvester and the receiver are impedance matched to the antenna using an LC matching network. The harvester extracts energy from ambient TV signals to provide small amounts of power required for ambient backscatter communication, sensors and the digital logic units (e.g., microcontroller). The harvester is based on a 5 stage Dickson charge pump topology from [186] and, has a sensitivity of -9 dBm. Since the receiver and harvester are directly connected to the antenna, the transmitter and receiver/harvester operate independently. When the transmitter is active and backscattering signals, the receiver and harvester cannot capture much signal/power and vice-versa. Next, we describe our design of the ambient backscattering transmitter and receiver in more detail.

2.3.2 Signal Propagation

We use Friss path loss equation to model RF signal propagation to understand the ambient backscatter scenario [153]. The coverage of ambient backscatter is a function of the distance between the TV tower, Alice and Bob. As shown in Fig. 2.1, consider that the TV transmitter is at a distance of

d_1 from Alice and d_2 from Bob. Alice and Bob are separated by a distance D. In typical scenarios, the distance between Alice and Bob $D \ll d_1, d_2$ and we can safely assume that $d_1 \approx d_1 = d$. So, for a TV transmitter operating at an output power P_t with an antenna gain G_t, the power density of ambient TV signals at a distance d from the tower can be written as

$$S = \frac{P_t G_t}{4\pi d^2} \tag{2.1}$$

Battery-free devices Alice and Bob with an antenna gain G_r will receive a direct path TV signal whose signal strength is given by

$$P_{direct,Alice} = P_{direct,Bob} = S\frac{\lambda^2 G_r}{4\pi} = \frac{P_t G_t}{4\pi d^2}\frac{\lambda^2 G_r}{4\pi} \tag{2.2}$$

(2.2) shows that the signal strength at Alice and Bob's location is inversely proportional with the distance from the tower i.e. a location closer to the TV tower will have a higher signal strength.

2.3.3 Ambient Backscattering Transmitter

In order to backscatter ambient TV signals, ambient backscattering transmitter builds on conventional backscatter communication techniques. Communication is achieved by changing the impedance connected to the antenna in the presence of an incident signal. A switch connected to the antenna modulates the impedance between two states and causes a change in the amount of energy reflected by the antenna. More formally, given an incident power density S (2.1) and assuming a 50 % duty cycle for transmitted bits, the power of the differential backscattered signal is given by

$$P_{diff,backscatter} = S\frac{\lambda^2 G_r^2}{4\pi}\left|\frac{\Delta\Gamma}{2}\right|^2 = \frac{P_t G_t}{4\pi d^2}\frac{\lambda^2 G_r^2}{4\pi}\left|\frac{\Delta\Gamma}{2}\right|^2 \tag{2.3}$$

where

$$\Delta\Gamma = \Gamma_1^* - \Gamma_2^* = \left(\frac{Z_a^* - Z_{L1}}{Z_a^* + Z_{L1}}\right) - \left(\frac{Z_a^* - Z_{L2}}{Z_a^* + Z_{L2}}\right) \tag{2.4}$$

Z_a denotes the impedance of the antenna and Z_{L1} and Z_{L2} are the two impedance states that are connected to the antenna corresponding to bit '0' and bit '1'. In our implementation, when the

input is zero, the transistor is off and the impedances Z_{L1} is close to Z_a i.e. a matched state to maximize the energy captured by the RF energy harvester. When the switch input signal is one, the transistor is on and the impedance Z_{L2} is close to zero, shorting the antenna which maximizes the signal scattered by the antenna. Thus, the switch toggles between a matched and short states to convey bits to the receiver.

The transmitter is implemented using the ADG902 RF switch [4] connected directly to the antenna. In this design, we use a 258 millimeter dipole antenna. The antenna is tuned for the 539 MHz UHF band and has a 50 MHz bandwidth. Other antenna topologies such as meandered antennas [181] and folded dipoles [174] can result in smaller dimensions, and further design choices can be made to increase the bandwidth of the antenna in order to make it capable of utilizing a larger frequency band. However, exploring alternate antenna designs is not within the scope of this work.

2.3.4 Ambient Backscattering Receiver

Consider the scenario shown in Fig. 2.1 where Alice is backscattering ambient TV signals to transmit bits to Bob. The RF signal at Bob's receiver can be written as combination of the direct signal from the TV and backscattered signal from Alice. Mathematically, we can represent this by combining (2.2) and (2.3)

$$P_{Bob} = P_{direct,Bob} + P_{diff,backscatter}\left(\frac{1}{4\pi D^2}\frac{\lambda^2 G_r}{4\pi}\right) \tag{2.5}$$

$$= \frac{P_t G_t}{4\pi d^2}\frac{\lambda^2 G_r}{4\pi} + \frac{P_t G_t}{4\pi d^2}\frac{\lambda^2 G_r^2}{4\pi}\frac{|\Delta\Gamma|^2}{4}\left(\frac{1}{4\pi D^2}\frac{\lambda^2 G_r}{4\pi}\right) \tag{2.6}$$

The challenge here for Bob's receiver is to decode bits backscattered by Alice on top of ambient signals which already encode information. Secondly, as we can see from (2.6), the backscattered signal is very weak compared (attenuated as a function of $\frac{1}{D^2}$) to direct path TV signal. Finally, Bob's receiver is severely power constraint and cannot use power-hungry hardware components such as ADCs and oscillators. Since, there is no oscillator, the receiver cannot recover phase information and has access to only amplitude information. Bob uses an envelope detector based

receiver design to decode the backscattered bits by tracking changes in the amplitude of the incident signal. To understand this in detail, we can rewrite the received RF signals shown in (2.6) in terms of electric fields

$$\begin{aligned} E^{i}_{Bob} &= E_{direct,Bob} + E_{backscatter} \\ &= E_{direct,Bob} + E_{direct,Alice} k \frac{e^{-j\beta D}}{D} \left(1 + \Gamma^{*}_{1/2}\right) \\ &= E_{direct,Bob} \left(1 + k \frac{e^{-j\beta D}}{D} \left(1 + \Gamma^{*}_{1/2}\right)\right) \end{aligned}$$

Here k models the scattering phenomena at Alice's antenna. Now assuming that Bob's receiver is matched to the antenna, the received voltage for the two cases when Alice transmits '0' bit and a '1' bit can be written in terms of voltages.

$$\begin{aligned} V^{0}_{Bob} &= \frac{l_e}{2} E_{direct,Bob} \left(1 + k \frac{e^{-j\beta D}}{D} \left(1 + \Gamma^{*}_{1}\right)\right) \\ V^{1}_{Bob} &= \frac{l_e}{2} E_{direct,Bob} \left(1 + k \frac{e^{-j\beta D}}{D} \left(1 + \Gamma^{*}_{2}\right)\right) \end{aligned}$$

where l_e is the electrical length of the antenna. In the rest of the section, we will show how we can decode the bits from the received RF signal in the presence of strong ambient TV signals ($E_{direct,Bob}$). First we will illustrate technique using conventional digital receiver and then, describe how to implement it in hardware using an ultra-low-power receiver.

Extracting Backscatter Information from Ambient Signals

Ambient signals like TV and cellular transmissions encode information and hence are not controllable. To illustrate this, Fig. 2.4(a) shows an example of the time-domain ambient TV signal captured on a USRP operating at 539 MHz. It can be seen that amplitude of the ambient TV signal varies significantly with time. This is expected because the captured ATSC TV signals encode information using 8VSB modulation, which changes the instantaneous voltage of the transmitted signal. Thus, the receiver should be capable of decoding the backscattered signals in the presence of these fast changing signals.

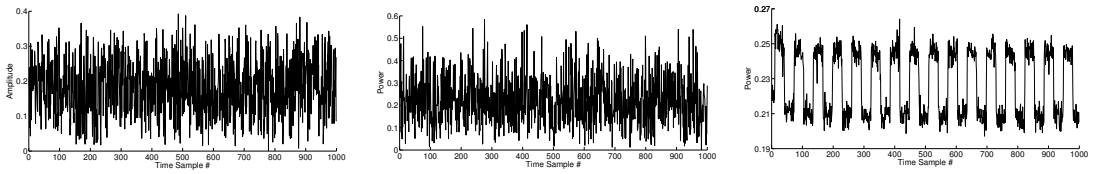

(a) Captured TV Signal (b) TV plus Backscatter signal (c) Signal After Averaging

Figure 2.4: **Comparison of TV signals.** We show the ambient TV signals incident on the receiver without any backscatter data, with backscatter data and with averaging of the backscattered data.

Our key insight is that if the transmitter backscatters information at a lower rate than the ambient signals, then one can design a receiver that can separate the two signals by leveraging the difference in communication rates. Specifically, ambient TV signals encode information at a bandwidth of 6 MHz, so if we ensure that the transmitter backscatters information at a larger time-scale than 6 MHz, then the receiver can extract the backscattered information using averaging mechanisms. Intuitively, this works because the wideband ambient TV signals change at a fast rate and hence adjacent samples in TV signals tend to be more uncorrelated than the adjacent samples in the backscattered signals. Thus, averaging the received signal across multiple samples effectively removes the variations in the wideband ambient TV signals, allowing the backscattered signals to be decoded.

Formally, we can describe this using a digital receiver that samples the received signal at Nyquist-information rate of the TV signal. The received samples, $y[n]$, can then be expressed as a combination of the wideband TV signals and the backscattered signals, i.e.,

$$y[n] = x[n]\left(1 + \alpha B[n]\right) + w[n]$$

where $x[n]$s are the samples corresponding to the TV signal as received by the receiver (corresponding to $E_{direct,Bob}$), α is the complex attenuation of the backscattered signals relative to the TV signals (corresponding to $\frac{e^{-j\beta D}}{D}$), $B[n]$ are the bits transmitted by the backscattering transmitter (corresponding to $\Gamma^{*}_{1/2}$) and $w[n]$ is the noise. Since the receiver samples at the TV Nyquist rate, the adjacent samples in $x[n]$ are uncorrelated. Now, if the backscatterer conveys information at a fraction of the rate, say $\frac{1}{N}$, then $B[Ni+j]$s are all equal for $j = 1$ *to* N.

For the design of the receiver we have two primary choices. We can use an energy detector

(square law model) and recover bits by comparing the energy of the incoming signal [96]. Alternatively, we can use an envelope detector where absolute value of the voltage of the incoming signal [145] is used to decode bits. We use the later since it corresponds to the hardware used in this work.

If the receiver averages the instantaneous absolute voltage in the N receiver samples corresponding to a single backscattered bit, then we get:

$$\frac{1}{N}\sum_{i=1}^{N}|y[n]| \;=\; \frac{1}{N}\sum_{i=1}^{N}|x[n]\,(1+\alpha B[n])+w[n]|$$

where B is either '0' or '1'. Since the TV signal, $x[n]$, is uncorrelated with noise, $w[n]$, we can rewrite the above equation as:

$$\frac{1}{N}\sum_{i=1}^{N}|y[n]| \;=\; \frac{|1+\alpha B|}{N}\sum_{i=1}^{N}|x[n]| + \frac{1}{N}\sum_{i=1}^{N}w[n]$$

Say V_{avg} is the average of the absolute value of the voltage in the received TV signal, i.e., $V_{avg} = \frac{1}{N}\sum_{i=1}^{N}|x[n]|$. Ignoring noise, the average voltage at the receiver is $|1+\alpha|V_{avg}$ and V_{avg} when the transmitter is in the reflecting and non-reflecting states, respectively. The receiver can distinguish between the two voltage levels, $|1+\alpha|V_{avg}$ and V_{avg}, to decode the information from the backscattering transmitter. Thus, even in the presence of changes in the TV signal, the receiver can decode information from the backscattering transmitter. We apply the above mechanism to the ambient ATSC TV signals [3]. Specifically, we set our ambient backscattering transmitter to transmit an alternating sequence of ones and zeroes at a rate of 1kbps. Fig. 2.4(b) plots the received signal on an USRP that is placed one foot from the transmitter. Fig. 2.4(c) plots the effect of averaging every 100 received samples. As the figure shows, averaging reduces the effect of the fast-varying ambient TV signals. Further, the receiver can now see two average power levels which it can use to decode the backscattered information.

We note that ambient backscatter can either increase or decrease the average voltage of the received signal. Specifically, the channel, α, corresponding to $\frac{e^{-j\beta D}}{D}$ is a complex number and hence

$|1+\alpha|$ can be either less than or greater than one. This means that a zero bit can be either a lower voltage or can have a higher voltage than the average. Intuitively, this is because the additional multi-path created by the backscattering transmitter can either constructively or destructively interfere up with the existing signal. We use differential coding to eliminate the need to know the extra mapping between the power levels and the bits (see §2.4.1).

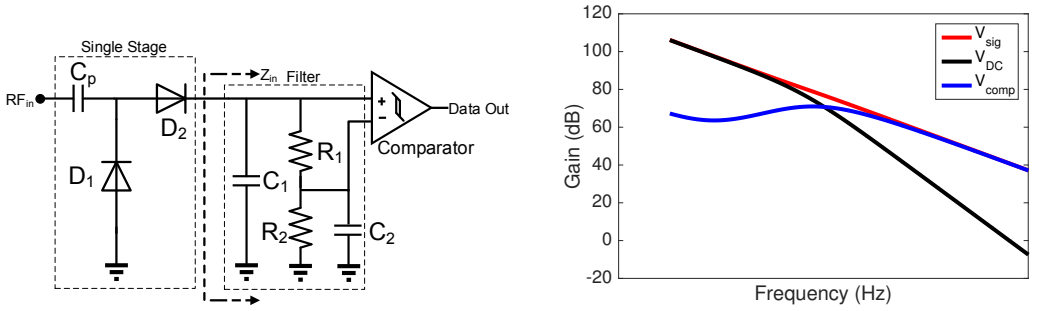

(a) Architecture of the receiver

(b) Bode plot of the filter on the receiver

Figure 2.5: **Circuit diagram and filter response for the receiver:** The receiver has two stages: an envelope detector with filter which computes the average envelope of the signal, and a threshold stage that compares averaged signal with a threshold value to output bits. The frequency response of the passive filter is also shown which illustrates the low pass nature of the receiver.

Decoding on an Ultra-Low-Power Device

The above design assumes that the receiver can get digital samples on which it can perform operations like averaging and comparison of power levels. However, acquiring digital samples requires an analog-to-digital converter (ADC) which can consume a significant amount of power and is typically avoided in ultra-low-power designs [206]. In this section, we implement the above operations in analog hardware by selecting an appropriate analog circuit topology.

The architecture of the receiver is shown in Fig. 2.5. It has two stages. First an envelope detector and a filter smoothens out the natural variations in the TV signal to output two voltages: an average envelope signal and a threshold voltage. Then a comparator circuit takes the two voltages as inputs to compute the difference and output data bits.

Envelope detector and filter: This circuit is implemented using a diode charge pump based

topology. The filters are implemented using a passive resistor and capacitor network as shown in Fig. 2.5(a). We use a single stage Dickson charge pump based voltage doubler as our envelope detector. We model the envelope detector using first order analysis. Given that the envelop detector is terminated by a low pass filter, we can assume that the output of the envelope detector varies very slowly or is near constant as compared to the input RF signal. For an input RF signal $\left(1+\alpha B_{0/1}\right) A_{TV} cos(2\pi ft)$, the average current flowing through the diode can be written as [145]:

$$I_{diode} = k * \left|\left(1+\alpha B_{0/1}\right) A_{TV}\right|$$

Here B_0 and B_1 correspond to '0' and '1' bit respectively. The diode current follows into the filter with the input impedance given by

$$Z_{in} = \frac{1}{sC_1} || \left(R_1 + R_2 || \frac{1}{sC_1}\right) \tag{2.7}$$

Next, the two voltages V_{sig} and V_{DC} can be written in the frequency domain as

$$V_{sig} = I_{diode} Z_{in} = k * \left|\left(1+\alpha B_{0/1}\right) A_{TV}\right| \frac{1}{sC_1} || \left(R_1 + R_2 || \frac{1}{sC_1}\right)$$

$$V_{DC} = V_{sig} \frac{\left(R_1 + R_2 || \frac{1}{sC_1}\right)}{\frac{1}{sC_1} || \left(R_1 + R_2 || \frac{1}{sC_1}\right)}$$

The bode plots for the frequency characteristics of the two voltages is shown in Fig. 2.5(b). V_{sig} is a first order low pass filter which attenuates the natural amplitude variations in the TV signal but passes the backscatter data. V_{DC} is a second order low power filter which attenuates both the natural amplitude variations in the TV signal and the backscatter data to determine the threshold voltage. By choosing the right values of the passive components the frequency response of the filter can be optimized for different data rates.

Comparator stage: The output of the envelope detector and filter outputs two voltages V_{sig} and V_{DC}. In principle, an ADC can be used to distinguish between the two signal levels by processing the digital samples. However, ADC's are power hungry and instead we use a 1-bit ADC i.e. a

comparator which takes $V_{comp} = |V_{sig}| - |V_{DC}|$ as the input and outputs '1' if V_{comp} is positive and '0' if negative. Fig. 2.5(b) shows the frequency response of V_{comp} which exhibits a band pass filter response. The band pass filter eliminates the low frequency components which correspond to channel variations and high frequency components of the amplitude variations in the ambient TV signals. It only passes the frequency components corresponding to the backscattered bits.

We note that the bit rate of the prototype dictates the choice of values for the RC circuit elements (e.g., a receiver operating at 10 kbps requires different RC values than one at 1 kbps). This is because the V_{comp} bandpass filter needs to be optimized for the frequency response corresponding to the bit rate. We use TS881 [59], as the ultra-low-power comparator in our design. We implement different bit rates by setting the capacitor and resistor values, R_1, R_2, C_1, and C_2 in Fig. 2.5, to $(150\ k\Omega, 10\ M\Omega, 27\ nF, 200\ nF)$ for 100 bps, $(150\ k\Omega, 10\ M\Omega, 4.7\ nF, 10\ nF)$ for 1 kbps, and $(150\ k\Omega, 10\ M\Omega, 680\ pF, 1\ \mu\ F)$ for 10 kbps. Finally, while in theory we can distinguish between any two power levels by sufficient averaging, each comparator has a minimum threshold below which it cannot distinguish between the two voltage levels. The threshold determines the maximum distance at which two devices can communicate with each other.

2.4 Network Stack Design

The network stack design for ambient backscatter communication is closely integrated with the properties of the circuits and the hardware described so far. In this section, we explore the physical layer and the link layer design for ambient backscatter.

2.4.1 Physical Layer

The physical layer for ambient backscatter communication addresses questions such as what modulation and coding to use, how to perform packet detection, and how to find bit boundaries.

Modulation and Bit Encoding: Since a backscattering transmitter works by switching between reflecting and non-reflecting states, it effectively creates an ON-OFF keying modulation. However, as described earlier, the backscattered signal could either constructively or destructively interfere

with the ambient TV signal. Thus, depending on the receiver's location, a '1' bit could appear as either an increase or a decrease in the received power. To address this issue, the physical layer uses FM0 coding [132]. FM0 coding turns every bit into two symbols and encodes information using symbol transitions [132]. FM0 has a symbol transition at the beginning of every bit period along with an additional mid-bit transition to represent a '1', and no such transition in the '0' bit. Thus, bits are encoded using transitions in the voltage level, rather than the actual voltage levels; further, it guarantees an equal number of '0' and '1' symbols.

Detecting the Beginning of a Packet Transmission: At the beginning of each packet transmission, an ambient backscattering transmitter sends a known preamble that the receiver detects using bit-level correlation on the digital hardware (in our case, the microcontroller). However, unlike RFID communication, where the tags correlate only when they are powered by a nearby reader, an ambient backscatter device does not know when nearby devices will transmit and hence might have to continuously correlate, which is power-consuming and impractical for a low-power device.

We avoid continuous correlation by only activating the relatively expensive correlation process when the comparator detects bit transitions. The comparator hardware takes very little power and has a built-in threshold before it detects bit transitions (in our implementation, this threshold is 2.4 mV). It is only when the power difference crosses this threshold that an interrupt is sent to the digital hardware to wake it up from its idle state (to perform correlation). Since the averaging circuit eliminates the large variations in the ambient TV signal, it is unlikely that ambient signals alone create changes in the power level in the absence of a packet transmission.

To provide the hardware with sufficient leeway to wake up the digital hardware, as shown in Fig. 2.6, the transmitter sends a longer preamble that starts with an alternating 0-1 bit sequence before sending the actual preamble. The alternating bit sequence is long enough (8 bits in our implementation) to wake up the digital hardware, which then uses traditional mechanisms to detect bit boundaries and perform framing.

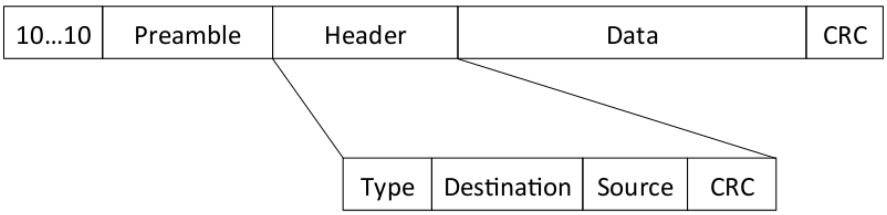

Figure 2.6: **Packet Format:** Each packet starts with an alternating sequence of '1's and '0's followed by a preamble that is used by the receiver to detect packets. The preamble is followed by a header and then the data, which both include CRCs used to detect bit errors.

2.4.2 Link Layer

Next we describe the following aspects of an ambient backscatter link layer design: error detection, acknowledgments, and carrier sense for mediating access to the channel.

Fig. 2.6 depicts the high-level packet format for ambient backscatter systems. The packet starts with a few bits of the preamble that are used to wake up the receiver's hardware; the rest of the preamble is then used by the receiver to detect the beginning of a packet. The preamble is followed by a header containing the type of packet (data/ACK), destination and source addresses, and the length of the packet. This is followed immediately by the packet's data. Both the header and the data include CRCs, which the receiver can use to detect bit errors in either field. Data may also be protected using simple error correction codes that do not consume significant power, e.g., hamming codes, repetition codes, etc. [187]. The receiver successfully receives a packet when both the CRC checks pass. It then sends back an acknowledgment within a pre-set time that is determined by the time it takes to successfully decode the packet at the receiver and switch to a transmitting state. In the rest of this section, we design carrier sense to arbitrate the wireless medium between these backscattering transmitters.

Carrier Sense

The discussion so far focuses on the communication aspects of a single ambient backscattering transmitter-receiver pair. However, when many of these devices are in range of each other, we need mechanisms to arbitrate the channel between them. In traditional RFID, a centralized, pow-

ered reader acts as the arbitrator for the wireless medium. Ambient backscatter communication, however, cannot rely on such a powered reader and thus requires a different set of mechanisms to provide media access control.

The advantage we have over traditional backscatter is that ambient backscattering devices can decode each other's transmissions. Thus, they can potentially perform carrier sense: detect the beginning of other packet transmissions (preamble correlation), and detect energy in the middle of a packet transmission (energy detection). Preamble correlation for carrier sense is operationally similar to that performed by the receiver for decoding packets. Energy detection, however, is challenging because the digital hardware does not have access to the power levels.

To see this, let us look at communication systems like Wi-Fi where energy detection is performed by computing the average power in the signal and detecting a packet when the average power is greater than a threshold. Such operations require a full ADC to get the digital samples on which to operate. Since an ambient backscattering device does not have access to a full ADC it does not have access to these power levels.

We show that one can perform energy detection by leveraging the property of the analog comparator. Specifically, unlike a traditional receiver where, even in the absence of nearby transmitters, it sees random changes in the received signal due to environmental noise; the bits output by our analog comparator is constant in the absence of a backscattering transmitter. This is because, as described in §2.4.1, the analog comparator has a minimum threshold below which it does not register any changes. Since the averaging circuit smoothens out the variations in the ambient signals, they typically do not create signal changes that are above this threshold. This means that in the absence of a nearby backscattering transmitter, the comparator typically outputs either a constant sequence of ones or a constant sequence of zeros. A nearby transmission, on the other hand, results in changes that are greater than the comparator's threshold and therefore bit transitions at the comparator's output. Since the transmitted bits have an equal number of ones and zeros (due to FM0 encoding), the comparator outputs the same number of ones and zeros. Thus comparing the number of ones and zeros allows the receiver to distinguish between the presence and absence of a backscatter transmission. More formally, the receiver performs energy detection by using the

following equation:

$$D = 1 - \frac{|\#ones - \#zeros|}{\#ones + \#zeros}$$

where $\#ones$ and $\#zeros$ denote the number of zeros and ones seen at the receiver over some time interval. In the presence of a backscattering transmitter, the average number of ones and zeros is about the same, and hence D is close to one. But in the absence of any close-by backscattering transmitters, the bits output by the comparator are either mostly ones or mostly zeros; thus, D is close to zero. Our results in §2.6 show that the above ideas hold even with mobility and in dynamic environments.

We note that the transmitter performs carrier sense only when it has data to transmit and before it starts transmitting. Upon detection of a competing transmission, microcontrollers (including the one used in our prototype) are able to sleep for the duration of the packet by masking interrupts caused by bit transitions.[3] Thus, the power drain of the above operations is minimal.

2.4.3 Further Discussion

So far we described the key functionalities (carrier sense, start-of-frame detection, etc.) required to build a network out of ambient backscatter devices. However, there are optimizations that can increase the performance of such systems; We outline some of them:

(a) Multiple bit-rates: Our current prototypes operate at a specific bit rate (either 100 bps, 1 kbps or 10 kbps). In principle, one can design a single device that has demodulators for different rates and switches between them. Further, one can design rate adaptation algorithms that adapt the rate to the channel conditions and can significantly increase the performance.

(b) Collision Avoidance: Carrier sense enables MAC protocols like CSMA that allow devices to share the medium. One can further reduce the number of collisions by designing collision avoidance mechanisms. Prior work on random number generation on low-power RFIDs [88] can, in principle, be leveraged to achieve this.

---

[3]To further minimize power, the microcontroller can sleep through the entire back-off interval, if we use non-persistent CSMA [121].

(c) Hidden Terminals: The devices can, in principle, use the RTS-CTS mechanism to address the hidden terminal problem. The overhead of RTS-CTS can be reduced by stripping the RTS-CTS messages of the data and header information, and having the transmitter send a unique preamble to denote the RTS message; the receiver sends back another unique preamble as a CTS message. Any nodes that hears these messages will not transmit for a fixed pre-determined amount of time, i.e., the time required to transmit the data packet and receive the ACK.

2.5 Prototype Implementation

We implement our prototype on a 4-layer printed circuit board (PCB) using off-the-shelf circuit components. As shown in Fig. 2.2, the prototype uses a dipole antenna that consists of two 2 sections of 5.08 in long 16 AWG magnetic copper wire. The prototype's harvesting and communication components are tuned to use UHF TV signals in the 50 MHz band centered at 539 MHz[4].

The packets sent by the transmitter follow the format shown in Fig. 2.6. Further, it is capable of transmitting packets at three different rates: 100 bps, 1 kbps, and 10 kbps. We also implement both preamble correlation and energy detection in digital logic to perform carrier sense at the transmitter. Our implementation currently does not use error correction codes and has a fixed 96-bit data payload with a 64-bit preamble.

On the receive side, we feed the output of the receiver circuit (described in §2.3.4) to the MSP430 microcontroller which performs preamble correlation, decodes the header/data and verifies the validity of the packet using CRC. We implement different bit rates by setting the appropriate values for the capacitor and resistor as described in §2.3.4.

Table 2.1 compares the power consumption of the analog portion of our transmitter/receiver with that of the WISP, an RFID-based platform [187]. The table shows that the power consumption numbers for ambient backscatter are better than the WISP platform, and almost negligible given the power budget of our device. This is because ambient backscatter operates at lower rates (10 kbps) when compared to existing backscatter systems like the WISP, which operates at 256 kbps.

[4]To target a wider range of frequencies, one can imagine using a frequency-agile, auto-tuning harvester that autonomously selects locally available channels, with a design similar to the dual-band RFID tag in [189].

Table 2.1: **Power Consumption of Analog Components**

	Tx	Rx
Ambient Backscatter	$0.25\mu W$	$0.54\mu W$
Traditional Backscatter (WISP [187])	$2.32\mu W$	$18\mu W$

So, we were able to optimize the power consumption of our prototype and achieve lower power consumption values.

Our prototype also includes two sensing and I/O capabilities for our proof-of-concept applications that are controlled by the microcontroller: low-power flashing LEDs and capacitive touch buttons implemented on the PCB using a copper layer. However, these sensors as well as the microcontroller that drives them can significantly add to the power drain. In fact, in the smart card application (see §2.7.1), the transmit modulator consumed less than 1% of the total system power, while the demodulator required another 1%; demonstrating that ambient backscatter significantly reduces the communication power consumption. The power management circuitry required an additional 8% of the total power. Flashing the LEDs and polling the touch sensors at the intervals used in §2.7.1 consumed 26% of the total power. The remaining 64% was consumed by the microcontroller.[5]

We note that in scenarios where the TV signal strength is weak, our prototype uses duty cycling to power the sensors and the microcontroller. Specifically, when the prototype is in the sleep mode, it only harvests RF signals and stores it on a storage capacitor. Once enough energy has been accumulated on the capacitor, it goes into active mode and performs the required operations. In hardware, the duty cycle is implemented by a voltage supervisor that outputs a high digital value (indicating active mode) when the voltage on the storage capacitor is greater than 1.8 V.

[5]We note that the high power consumption for the digital circuit (i.e., microcontroller) is an artifact of our prototype implementation. Specifically, the microcontroller is a general-purpose device that is not typically used in commercial ultra-low-power devices. Instead, commercial systems use Application-Specific Integrated Circuits (ASICs) that can consume orders of magnitude less power than general-purpose solutions [169, 187]. In ASIC-based low-power devices, the power consumption of the analog components often dominates that of the digital circuit [74].

2.6 Evaluation

We evaluate our prototype design in the Seattle metropolitan area in the presence of a TV tower broadcasting in the 536-542 MHz range. We ran experiments at six total locations to account for attenuation of the TV signal and multipath effects in different environments. The TV signal power in the 6MHz target band for the given locations ranged between -24 dBm and -8 dBm. These locations consist of:

- *Location 1 (Indoor and near):* Inside an apartment 0.31 miles away from the TV tower. The apartment is on the seventh floor of a large complex with 140 units and is located in a busy neighborhood of a metropolitan area.

- *Location 2 (Indoor and far):* Inside an office building 2.57 miles away from the TV tower. The office tested is on the sixth floor of the building.

- *Location 3 (Outdoor and near):* On the rooftop of the above apartment.

- *Location 4 (Outdoor and far):* On the rooftop of the above office building.

- *Location 5 (Outdoor and farther):* On a street corner 5.16 miles away from the TV tower.

- *Location 6 (Outdoor and farthest):* On the top level of a parking structure 6.50 miles away from the TV tower.

We evaluate the various aspects of our design including our ambient backscattering transmitter and receiver, carrier sense, and interference at TV receivers. Most of our experiments were limited to locations 1-4 due to limited extended access to space in locations 5 and 6. The latter two locations, however, were included to demonstrate that ambient backscatter can operate at longer ranges and were tested using our smart card application.

Those test verified that we were able to get our end-to-end system to operate battery-free up to 6.5 miles away from the TV tower. Note, however, that the operational distance of our prototype

is dependent on the operating voltage of the device. In our prototype, the bottleneck was the microcontroller, which requires 1.8 V. In principle, an ASIC-based design should work with much lower voltage requirements and hence can operate at farther distances.

2.6.1 Effectiveness of Ambient Backscattering

The effectiveness of a backscattering transmitter is determined by the extent to which it affects the received signal. To quantify this, we compute the ratio of the received power, after averaging, between the non-reflecting and reflecting states of the transmitter. Specifically, if P_1 and P_2, $P_1 \geq P_2$, are the two average power levels at the receiver, we compute the ratio, $\frac{P_1}{P_2}$. A ratio close to one means that the receiver cannot distinguish between the two power levels; while a higher ratio increases the ability of the receiver to distinguish between them.

Experiments: We configure our prototype to send an alternating sequence of bits—switching between reflecting and non-reflecting states—at a rate of 100 bps. The results are similar for the other bit rates. Since our receiver prototype does not provide the exact power values, we instead use an USRP-N210 as a receiver to compute the power ratio between the two states. The USRP is connected to the same dipole antenna used by our receiver prototype to ensure that the antenna gains are identical. We configure the USRP to gather raw signals centered at 539 MHz using a bandwidth of 6.25 MHz—the bandwidth of the ambient TV signals. We average the received signal, and compute the ratio between the two average power levels (which corresponds to square root of the voltage). We repeat the experiments for different distances (from 0.5 feet to 3 feet) between the transmitter and the receiver in locations 1-4.

Results: Fig. 2.7 plots the CDF of the observed power ratios at the receiver. The CDF is taken across both indoor/outdoor and near/far locations to provide an overall characterization of ambient backscatter that we delve into next. The figure shows the following:

- The median power ratio is about 1.4, which is in the range targeted by traditional backscatter communication in RFID devices [199] and is a favorable ratio. To get an intuition for why this is the case, consider a hypothetical scenario where the transmitter and a receiver see the

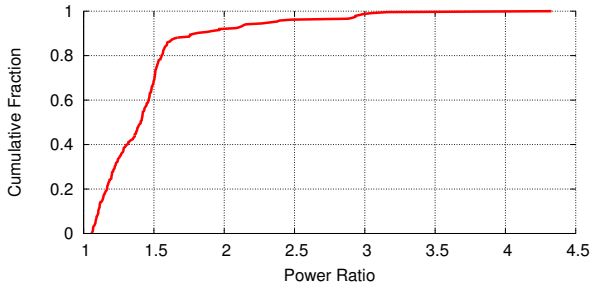

Figure 2.7: **Performance of an ambient backscattering transmitter:** The x-axis plots a CDF of the ratio of the average power received during the reflecting and non-reflecting states of the backscattering transmitter. The CDF is taken across multiple positions in both indoor/outdoor and near/far scenarios.

same ambient TV signal strength and the transmitter backscatters all its incident signals in the direction of the receiver. In this case, even if the transmitter and receiver are placed next to each other, the average received power with backscatter is twice the received power without backscatter, i.e., the power ratio is 2. In practice, however, the ratio is often much lower than this idealized value, as a transmitter reflects only a fraction of its incident signal in the receiver's direction; larger distances further attenuate the signal strength.

- The power ratio can be as high as 4.3. This is due to the wireless multipath property. Specifically, because of multipath, nodes that are located at different locations see different signal strengths from the TV tower. So when the transmitter is in locations where it sees a much higher TV signal strength than the receiver, its backscattered signal can be significantly higher in amplitude than the direct TV signal.

2.6.2 BER at the Ambient Receiver v/s Distance

Next, we evaluate our low-power receiver described in §2.3.4.

Experiments: We repeat the previous experiments, but with our prototype ambient receiver receiving from the backscattering transmitter. We measure the bit error-rate (BER) observed at the receiver as a function of the distance between the transmitter and the receiver. For each distance value, we repeat the experiments at ten different positions to account for multipath effects; the

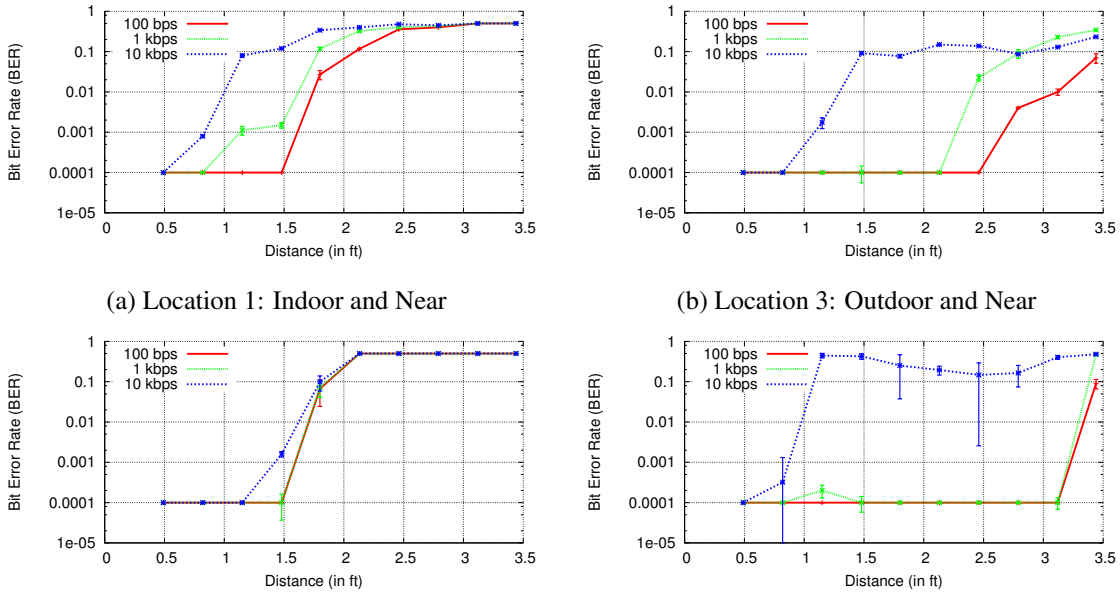

(a) Location 1: Indoor and Near

(b) Location 3: Outdoor and Near

(c) Location 2: Indoor and Far

(d) Location 4: Outdoor and Far

Figure 2.8: **BER v/s Distance.** BER for transmitter-receiver pairs in a range of environments, both outdoor and indoor, close to the TV tower, and far away. We show BER for distances of over three feet and three different rates.

transmitter sends a total of 10^4 bits at each position. The BER is computed by comparing the transmitted bits with the bits output by the prototype's demodulator circuit. Since the total number of bits transmitted at each position is 10^4, we set the BER of experiments that see no errors to 10^{-4} (the upper bound on the BER for these experiments). Finally, since the BER depends on the transmitter's bit rate, we evaluate three different prototypes that are designed to work at 100 bps, 1 kbps, and 10 kbps. We note that, in total, we perform 1020 measurements across bit rates and locations.

Results: We plot the results in Fig. 2.8. The figures show that:

- As the distance between the transmitter and receiver increases, the BER across bit rates and locations increases. Further, the BER is better in outdoor locations than in indoor locations. This is because TV signals are significantly attenuated in indoor locations and hence the ambient signal strength is much lower.

- Locations 1 and 3 perform slightly worse than locations 2 and 4, even though they are closer

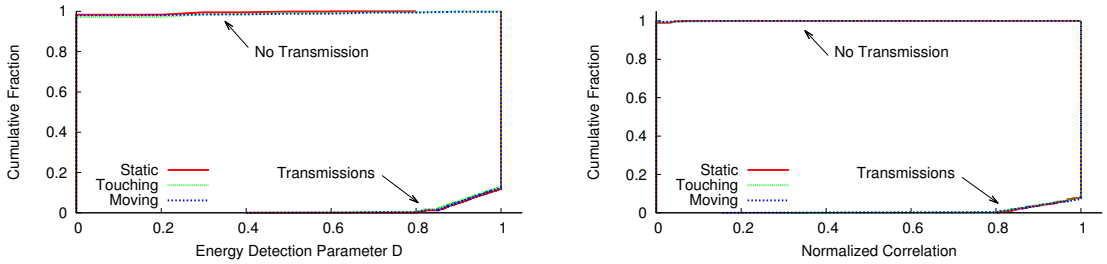

(a) Energy Detection (b) Preamble Correlation

Figure 2.9: **Performance of Carrier Sense:** These figures show that we can effectively perform energy detection and preamble correlation—the two main components of CSMA—on ambient backscattering devices.

to the TV tower. This is due to the fact that the TV tower is not an ideal isotropic antenna: the radiated power is less at low angles, and thus the signal strength is less at the near locations.

- For a target BER[6] of 10^{-2}, the receiver can receive at a rate of 1 kbps at distances up to 2.5 feet in outdoor locations and up to 1.5 feet in indoor locations. Such rates and distances are sufficient to enable ubiquitous communication in multiple scenarios, including our proof-of-concept applications.

2.6.3 Evaluating Carrier Sense

We implement carrier sense using both energy detection and preamble correlation. Energy detection is performed by computing $D = 1 - \frac{|\#ones - \#zeros|}{\#ones + \#zeros}$, where $\#ones$ and $\#zeros$ denote the number of ones and zeros seen at the receiver, within a 10-bit interval. Preamble correlation is performed by correlating with a known 64-bit preamble.

We place a transmitter and receiver, both designed for 1 kbps, in random locations within two feet of each other in both of the indoor locations. These distance are enough to include configurations where a 1 kbps receiver can hear the transmitter, but experiences high bit error rate ($>10\%$). This is corroborated by the fact that the BER observed across the tested locations is in the range of 10^{-4} to 0.17. The experiments are performed both in the presence and absence of backscattering from the transmitter. We repeat the experiments at 300 locations and for three different scenarios:

[6]The packet size is 96 bits and hence can tolerate a 10^{-2} BER with simple repetition coding [172].

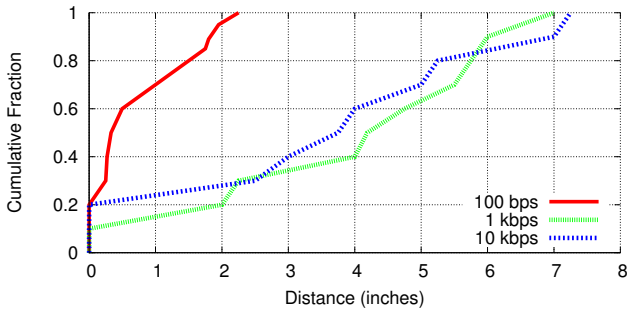

Figure 2.10: **Interference with TV Receivers:** CDF of the minimum distance at which ambient backscatter transmitters of various rates do not interfere with traditional TV receivers.

no motion near the receiver, human motion near the receiver, and a human holding the receiver and waving her hand in front of it.

In Fig. 2.9(a) we plot the CDF of the computed energy detection values (Ds). The plot shows the following: Firstly, in the absence of backscatter, D is *exactly* zero in more than 98% of the experiments. This happens because, as described in §2.4.1, the analog comparator used in the receiver, typically, outputs either a constant sequence of ones or a constant sequence of zeros in the absence of a backscattered signal. Thus, the receiver sees the same bit during a 10-bit interval. Secondly, human mobility does not create statistically significant differences in the computed D values. This is because while motion can change the signal strength at the receiver and the corresponding bits output by the comparator, it is unlikely that it either creates bit changes at the rate of 1 kbps or creates an equal number of bit changes in a 10-bit interval. Finally, the plot shows that in more than 99% of the experiments there is a clear distinction between the presence and absence of a backscattering transmitter.

We also plot in Fig. 2.9(b) the CDF for preamble correlation both in the presence and absence of a packet that starts with a preamble. The correlation values are normalized by the length of the preamble (64). The plot shows a clear distinction between the presence and absence of a preamble, in more than 99.5% of the experiments. This is again because of the property of the comparator which outputs sequences of either constant one bits or constant zero bits in the absence of backscatter, which are unlikely to be confused with a pseudo-random preamble.

2.6.4 Interference with TV Receivers

Since the backscattered signals are reflections of existing TV signals, in theory, one could either synchronize ambient backscatter with the TV transmissions or modulate data at a slow enough rate that TV receivers would be immune to interference. However, even without these constraints, the backscattered signals are weak enough that they do not affect TV receivers except in less favorable conditions. In this section, we stress-test ambient backscatter to get a sense for the upper bound of its effects on TV receivers. To that end, we tested very small antenna-tag distances (less than a foot) and performed the experiments inside the office building of location 2, which has the weakest TV signal power.[7]

We use an off-the-shelf Panasonic Plasma HDTV (Model No: TC-P42G25) connected to a cheap tuner (Coby DTV102) and a basic RCA indoor antenna (Model No: ANT111). We tune the TV channel to the transmissions at 539 MHz. To evaluate the worst case behavior where the transmitter always backscatters information, we connect the transmitter to a power source and set it to continuously transmit random bits. The transmit antenna is placed parallel to the TV antenna to maximize the effects of backscatter on the TV receiver. The transmitter is placed at a random location one foot away from the TV antenna. It is then moved towards the TV antenna until we first notice visual glitches in the video; we measure the distance at which this happens. Note that, in digital television, interference is relatively easy to quantize as errors result in corrupted portions of the image, rather than just noise as is the case in analog television. To quantify visually observable glitches, we had two users simultaneously looking for any momentary, visually observable artifact (including misplaced squares of pixels) on the screen.

Fig. 2.10 plots the CDF of the glitch distance for different bit rates at the transmitter. The CDF is taken across multiple experiments. The plots show the following:

- A 100 bps backscattering transmitter does not create any noticeable glitches at the TV receiver unless it is less than 2.3 inches from the TV antenna. This is because the backscattered signal

[7] Results from locations that have stronger TV signals show that the TV receiver was more resilient to interference. The majority of the time, there were no visual artifacts for distances above 1 in, and we never observed any glitches for any bit rate at distances above 3 in.

effectively creates a new path from the transmitter to the TV receiver. Since TV receivers are designed to compute the multi-path channel parameters, they can estimate the effects of this new path and decode the TV transmissions without interference. However, for small distances (less than 2.3 inches), the near-field effects dominate and hence the linearity model, typically assumed while estimating the multi-path channel, does not hold; resulting in video glitches.

- The distance at which the video glitches are noticeable is larger for higher transmission rates: the median distances is about 4.1 inches and 3.7 inches for 1 kbps and 10 kbps respectively. At high transmission rates, the transmitter changes the multipath channel at a higher rate; hence, making it difficult for the TV receiver to estimate the fast-changing multipath channel.

- Across bit rates, the TV receiver does not see any noticeable glitches for distances greater than 7.2 inches.

2.7 Proof-of-Concept Applications

Ambient backscatter enables devices to communicate using only ambient RF as the source of power. We believe that this opens up a new form of ubiquitous communication where devices can communicate by backscattering ambient RF signals without any additional power infrastructure. In this section, we demonstrate proof-of-concepts for two applications that are enabled by ambient backscatter: a bus card that can transfer money to other cards anywhere and a grocery store application where item tags can tell when an item is placed in a wrong shelf. These proof-of-concepts are similar to existing RFID applications, but differ in ways that were previously impossible—they are able to function anywhere and with no maintenance. They are only a glimpse into the possibilities opened by this technique, and we consider fully exploring the potential uses and addressing issues such as security or usability to be out of the scope of this paper.

2.7.1 Smart Card Application

We use our prototype design to evaluate a smart card application where passive cards can communicate with each other anywhere, any time, without the need for a powered reader. Such an application can be used in multiple scenarios, such as money transfer between credit cards, paying bills in a restaurant by swiping the credit card on the bill or to implement a digital paper technology which can display digital information using e-ink [212] and transfer content to other digital paper using ambient backscatter.

In this section, we implement and evaluate a simple proof-of-concept of the smart card application. We leverage our prototype that comes complete with an ambient backscattering transmitter/receiver, MSP430 microcontroller, capacitive touch sensor, and LEDs. When a user swipes the touch sensors (marked by A, B, C in Fig. 2.2), in the presence of another card, it transmits the phrase "Hello World". The receiver on the other card decodes the transmission, checks the CRC, and confirms a successful packet decoding by flashing the LED. We perform this experiment at three different locations including the two locations farthest from the TV tower.

Experiments: We place the cards 4 inches from each other and have the user perform the swipe. The transmitter and receiver communicate at a bit rate of 1 kbps. The microcontroller is programmed to detect changes at the touch sensors and trigger transmissions. The transmitter backscatters a packet with a 96 bit "Hello World" payload and a 4-bit CRC. The receiver decodes the packet and if the CRC check passes, blinks the LED (for 1 ms) to provide a visual confirmation to the user. The devices are powered completely by harvested TV energy. The user performs the swipe 100 times at an interval of three seconds between the swipes. Since blinking the LED drains the capacitor in the harvester, the three second time interval allows the harvester to duty-cycle and accumulate charge on the capacitor to perform the LED blinking operations again.

Results: Fig. 2.11(a) plots, for locations 5 and 6, the number of retries required by the user to successfully perform the whole operation: the user swiping the touch sensors, the card transmitting the packet, and finally, the LED blinking on the other card. The plot shows that in 94% of cases, the user only had to perform one swipe to see the LED blink on the other card, and even in the

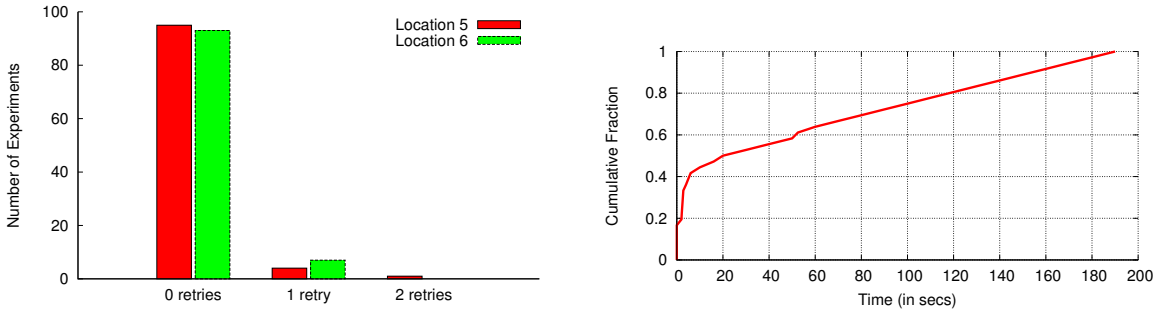

(a) Smart Card Application

(b) Grocery Store Application

Figure 2.11: **(a) Smart Card Application:** The number of retries necessary to successfully communicate between two battery-free smart cards. 94% of tests were successful without any retries. **(b) Grocery Store Application:** The CDF of the time it takes for the out-of-order item to blink its LED.

worst case, the user did not require more than two retries to successfully complete the operation. Furthermore, automation of communication startup (i.e., removing the user and touch sensor from the process) decreased the failure rate to nearly zero, indicating that it was not the communication mechanism that was failing, but rather user input error.

2.7.2 Grocery Store Application

Ambient backscatter can also be used to tell when an item is missing or out of place on a shelf in a grocery store. In this section, we use our prototype to evaluate a proof-of-concept for this application. The algorithm we use is simple: each device broadcasts its ID periodically (every 5 sec). Neighboring tags listen to these transmissions and store the successfully decoded IDs. Each tag determines on its own if it is out-of-place by computing the difference between its ID and that of the overheard IDs. If the tag has at least two different stored IDs that have this distance to be greater than a threshold, it concludes that it is out-of-place and flashes the LED.

Experiments: We attach ten of our prototype tags to ten cereal boxes, and place those boxes next to one another on a shelf. We set the IDs for nine of these tags to be between 201 and 209 and place them in-order. We then set the ID for the tenth tag to be 100 and place it in the 8 locations between the in-order tags, for a total of 40 experiments. A nearby antenna broadcasts an RFID signal, and we measure the time it takes for the out-of-place tag to flash the LED.

Results: We plot the results in Fig. 2.11(b). The plot shows that in about 50% of the experiments, the out-of-place tag requires less than 20 seconds to flash the LED. Further, in the worst case, the out-of-place tag starts blinking within 190 seconds. We note, however, that the results in this section are not optimized and are only presented to demonstrate the application's feasibility.

2.8 Discussion

The proposed architecture has certain implementation-dependent limitations which impact the practicality of the system. We will discuss three of them here: frequency selectivity of the RF harvester, operating distance from an ambient source, and the constraints imposed by the duty cycling behavior of the system.

2.8.1 Frequency selectivity

Like many existing RF energy harvesting systems, the RF front end used by our prototype operates over a relatively narrow frequency range, and is pre-tuned to cover a single, target TV channel. That channel may not be present in all areas, which may hinder the potential for ubiquitous deployment. To address this frequency selectivity, we describe two possible solutions. The first is the use of a wideband harvesting system that efficiently converts RF energy to DC power over a large range of frequencies. For instance, one existing system claims a 160% fractional bandwidth in the microwave band [104]. A second possible solution is the use of a frequency-agile, auto-tuning harvester which autonomously selects and "tunes in" locally available channels. [189] describes such an architecture that could be repurposed for ambient backscatter. Wideband and/or frequency-agile RF front ends will improve the ability of the system to both harvest energy and communicate in more diverse locations.

2.8.2 Operating range

At a minimum, a device's harvester must produce a voltage across the charge storage capacitor that meets or exceeds the device's required operating voltage (in our prototype, 1.8 V). The minimum

ambient power required to achieve that voltage is called the *harvester sensitivity*, and determines how far from a TV tower the device is able to operate. Harvester sensitivity determined by the voltage gain of the harvester and amount of leakage present in the system. Trends in IC manufacturing are producing devices with increasingly low operating power, and thus, improved operating range [191]. Leakage must be addressed both at a system level by shutting off unused components and at a component level in the harvester/charge storage design. When application components are disabled, the primary source of leakage in this system is the storage capacitor itself. Reductions in power consumption of the system will reduce the size requirement for this storage capacitor and thereby reduce the leakage, further increasing operating range.

2.8.3 Duty cycling behavior

If less power is available than the active power consumption of the device, it must duty cycle (i.e., enter an inactive or idle state periodically). Ideally, this allows the device to operate with arbitrarily small amounts of incoming power. In the prototype tests and example applications shown, the duty cycle was typically 100% as sufficient power was available to sustain continuous operation. However, in cases when the device needs to duty cycle, it must be able to complete a minimal task—an atomic workload—prior to suspending or shutting off again. This atomic workload can be made smaller by implementing opportunistic checkpointing as described in [180], allowing delay tolerant applications to be practical even at a very low duty cycle. In short, for many applications it is possible to trade duty cycle for operating range to an arbitrary extent.

2.9 Related Work

Prior work mainly falls in the following two domains:

(a) Wireless Communication: Today, wireless communication is limited to two main approaches: radio communication and backscatter communication. Conventional radio communication requires devices to generate radio signals. This approach is problematic from a power perspective since it requires power-hungry analog components such as digital-to-analog converters (DACs), mixers,

oscillators and power amplifiers at the transmitter [128, 182] and low noise amplifiers, mixers, oscillators and ADCs at the receiver [128, 182]. While prior research has focused on reducing the power consumption of these analog components [21, 74, 98, 162, 177], backscatter communication is two orders of magnitude more power-efficient than state-of-the-art radio communication [77, 169, 199]; and hence is more appropriate for battery-free devices [185, 208].

Traditional backscatter communication (e.g., RFID), however, requires a powered device called the reader to generate a high-power constant signal which battery-free devices backscatter back to the reader. These battery-free devices are rendered unusable in the absence of the powered reader and hence require an infrastructure of powered readers that can be expensive and infeasible.

This paper introduces ambient backscatter, a new approach to communication where devices can communicate without either generating signals (as in radio communication) or backscattering from a dedicated reader (as in traditional backscatter). Ambient backscatter eliminates the need for a power infrastructure and hence can enable new forms of ubiquitous communication at locations and scales that were previously infeasible.

The closest to our work is recent work in [152] that demonstrates direct communication between two RFID tags placed 25 mm away from each other. However, it works only in the presence of a dedicated RFID reader that generates a constant high-power signal. Our work is orthogonal to [152] in that we enable devices to communicate using ambient RF signals. We note, however, that in principle the techniques in this paper can also be used to enable RFID tag-to-tag communication at much larger distances than 25 mm.

(b) Power Harvesting: In this domain, our work is most directly related to wireless power and ambient RF power harvesting. Wireless power aims to wirelessly charge and power devices by transmitting energy from a dedicated power source [187]. Ambient backscatter is complementary to this work. Specifically, it focuses on enabling communication using ambient RF as the only source of power, without requiring any additional power sources.

Recent work on ambient RF power harvesting demonstrated that one can harvest useful amounts of power from ambient TV [186] and cellular signals [166]. Our work is motivated by this work and takes it one step further. Specifically, we introduce a new communication system that enables

devices to communicate with each other using ambient RF. We achieve this by introducing ambient backscatter where devices communicate by backscattering ambient RF signals.

2.10 Conclusion

For the first few decades of their existence, computers were fundamentally limited by the infrastructure on which they rely. Computers were tethered by their power cords and were rendered useless without a nearby power outlet. Wireless communication combined with battery packs liberated these devices for short periods of time so that they could compute and communicate, untethered, as long as their batteries were occasionally recharged or replaced.

In this work, we introduce ambient backscatter, a new form of communication that provides connectivity between computers out of what is essentially thin air. In this technique, TV signals and other source of RF signals serve as both the source of power and the means of communication. Because ambient backscatter avoids the maintenance-heavy batteries and dedicated power infrastructure of other forms of low-power communication (e.g., RFID and NFC), it enables a bevy of new applications that were previously impossible or at least impractical. We believe that ambient backscatter provides a key building block that enables ubiquitous communication (with no restrictions except the existence of ambient RF signals) among pervasive devices which are cheap and have near-zero maintenance.

Chapter 3

ENABLING INSTANTANEOUS FEEDBACK WITH FULL-DUPLEX BACKSCATTER

3.1 Introduction

Backscatter communication is a promising technique for low-power computing devices. RFID, for instance, has enabled a multitude of useful applications [185, 208] that take advantage of the fact that one side of the link can be completely battery-free. More recently, researchers have shown the feasibility of using backscatter to communicate directly between two low-power devices [131]. The power draw of this technique is such that both devices can operate solely off of harvested energy. By creating a network of battery-free devices that can communicate with each other, this can enable the types of applications envisioned by the Internet-of-Things (e.g., smart homes).

Unfortunately, existing link- and network-layer protocols are ill suited for such communication networks. Extreme power constraints not only magnify the effects of collisions and failed transmissions; they also make it difficult for these devices to use multi-round protocols like RTS-CTS as storing increasing amounts of energy becomes exponentially hard. In this chapter, we explore the feasibility of an instantaneous feedback channel where a receiver can simultaneously send feedback information while still receiving the original transmission. Such a feedback channel could address the following problems facing these devices:

- *Packet Collisions:* Instantaneous feedback could be used to terminate transmission as soon as a collision is detected. Because recharge time dominates transmission time [187], enabling early termination increases throughput by orders of magnitude.

- *Rate Adaptation:* Instead of waiting for packet drops to adjust rate, a feedback channel allows us to gather bit error rate statistics and use them to change rate at the level of bits rather than

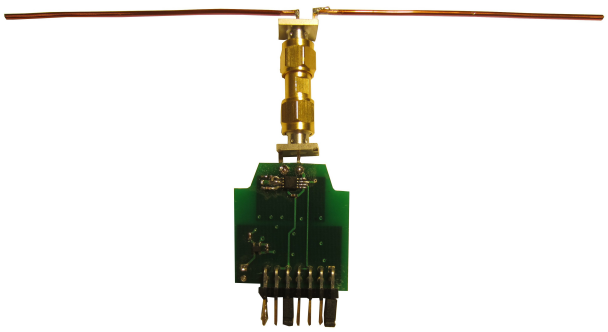

Figure 3.1: **Full-duplex Backscatter Prototype:** A photo of our Full-duplex Backscatter prototype. This top side contains the transmitter switch, while the receiver is implemented on the bottom side.

packets. This is particularly important for backscatter devices, as their transmission rates are orders of magnitude lower than traditional radio communication (e.g., Wi-Fi). Thus, they take significantly longer to transmit the same amount of data—long enough that the channel may change within the span of a single packet.

- *Retransmissions:* In the same way, we can use a feedback channel to inform transmitters of bit-level errors. Thus, the transmitter only needs to retransmit the subset of bits that were incorrect, rather than throwing away the packet entirely and trying again.

Traditional radio communication designs have explored the idea of adding a feedback channel, but none of those existing approaches are applicable to backscatter systems. They mainly fall into two categories: frequency-division duplexing and full-duplex communication. Frequency division is not practical because battery-free backscattering tags use simple analog envelope detectors for reception and switches for transmissions. They therefore cannot be designed to be frequency selective while maintaining low power consumption [187].[1] Similarly, recent work on full-duplex communication [72, 86] does not apply, because these devices do not have the computational or energy resources to implement complex interference cancellation techniques or include space-consuming arrays of antennas. In fact, existing full-duplex designs are prototyped on software radios, each of which consumes more than 1-2 Watts of power [53, 107]. More importantly, these designs use

[1] RFID tags can not decode in the presence of concurrent transmissions in non-overlapping frequency bands.

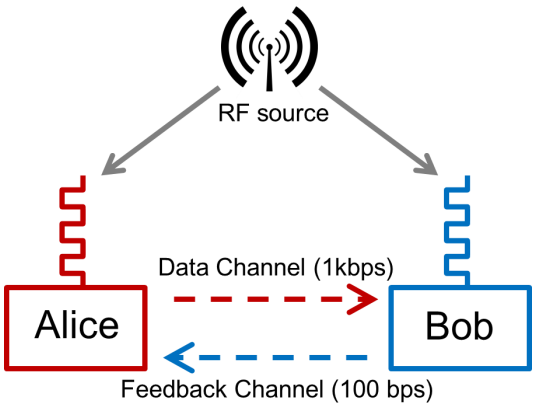

Figure 3.2: **Full-duplex Backscatter:** Two battery-free devices, Alice and Bob, communicate by backscattering signals from the RF source. Alice transmits to Bob at a high rate on the data channel and receives instantaneous feedback from Bob on the feedback channel.

components such as ADCs and oscillators that are power-consuming and hence are not applicable for backscatter devices [77].

We introduce a novel technique called Full-duplex Backscatter with which backscatter devices can obtain instantaneous feedback on the same frequency as that of the transmission, without the need for multiple antennas or power-consuming cancellation hardware. Our technique uses fully-passive analog components in order to enable devices to simultaneously maintain a data channel and a low-rate feedback channel in the opposite direction. By using simple analog components such as diodes and resistors, we incur near-zero power cost.

To understand Full-duplex Backscatter, consider the ambient backscatter setup from Chapter 2 shown in Fig. 3.2 with two battery-free devices, Alice and Bob. Alice communicates with Bob by backscattering the transmissions from the RF source; she does so by reflecting or absorbing the incident signals to convey a '0' bit and '1' bit, respectively. Bob decodes these bits by tracking the changes in the average amplitude of the received signal.

The challenge in creating a simultaneous feedback channel from Bob to Alice is that the act of backscattering at Bob creates large changes in Bob's received signal amplitude that greatly degrades his decoding capabilities. Conceptually, since Bob uses the same antenna to receive and transmit, the amplitude of his received signal will change significantly based on whether he is

reflecting (sending a '0') or absorbing (sending a '1') on the feedback channel.

Our intuition to create a full duplex communication is as follows: *Since Bob's receiver decodes by tracking the amplitude of the received signal, if Bob can backscatter information without changing the amplitude, he will not create any self-interference.*[2] Our approach is therefore to reflect and absorb a *fixed* amount of signal, such that the amplitude of the received signal is constant. To achieve this, Full-duplex Backscatter changes the impedance of the antenna to create phase shifts to the received signal while maintaining the same amplitude. Two reflections with different phases will interfere with the ambient signal at Alice to create two different amplitude levels that Alice can decode using a standard, amplitude-tracking backscatter receiver. We note, however, that while in theory one can pick the impedance values to create equal-amplitude signals with different phase shifts, practical circuits have small impedance mismatches that result in residual interference. In §3.3, we describe how Full-duplex Backscatter leverages the low rate of the feedback channel to eliminate the effect of this residual interference.

To demonstrate the feasibility of our techniques, we built the prototype in Fig. 3.1 using off-the-shelf components. Our prototypes communicate with each other by backscattering continuous wave transmissions from an RF source in the 920 MHz range. We configure the devices to transmit at 100 bps on the feedback channel and 1 kbps on the data channel; the latter is the state-of-the-art for device-to-device backscatter communication [131]. Our evaluation shows that our hardware reduces the self-interference close to the noise floor across the frequency range and for a range of received power levels, while consuming only 0.25 μW and 0.54 μW of transmit and receive power respectively.

We also integrate the feedback channel provided on our prototype into the backscatter network stack in order to perform collision detection, implement rate adaptation at sub-packet granularities, and reduce retransmissions by performing in-frame error correction. Our results are as follows:

- By terminating colliding transmissions early, we reduce transmitter recharge time by two orders of magnitude for small 64-byte packets across the operational range of receive power levels.

[2]Decoding both amplitude and phase information requires power-consuming oscillators [114] that are avoided in backscatter devices. Thus, phase is a free-parameter that is available for our purposes.

The reduction is higher for longer packet sizes.

- By performing in-frame error correction, we reduce the number of retransmitted bits by an order of magnitude, for a range of bit error rates on the data channel.

- Performing in-frame rate adaptation can increase the throughput of a backscatter communication system by about 33% compared to an idealized, SNR-based, packet-level adaptation algorithm. We note that existing backscatter systems do not have access to SNR information and thus the benefits of our system are likely greater.

Contributions. In this chapter, we make the following contributions:

- Introduce the first design that enables an instantaneous feedback channel on battery-free backscatter devices. We do so by leveraging the properties of backscatter communication to perform self-interference cancellation using only passive analog components.

- Develop and integrate the feedback channel with the network stack to address collisions, fine-grained rate adaptation, and error-correction on battery-free devices.

- Design and build hardware prototype that demonstrates the feasibility of our techniques and protocols in practice.

Network designers have traditionally believed that a full-duplex feedback channel is difficult to achieve on battery-free devices. We not only demonstrate the feasibility of such a channel but also show that it can enable link-layer and network-layer mechanisms that can significantly benefit backscatter communication. This is an important step towards creating a practical network of backscatter devices that can communicate with each other.

3.2 Motivation for Feedback in Backscatter Systems

A feedback channel has multiple benefits for both traditional backscatter (between a battery-free device and a powered reader) as well as ambient backscatter communication (between two battery-

free devices). The key operational challenge for these devices is that they may not have enough energy to transmit multiple packets back-to-back due to the fact that they largely rely on harvested power—a resource that is both limited and unpredictable. [3]

Below, we provide three illustrative examples of core issues in backscatter communication, and discuss how an instantaneous feedback channel can help address these issues.

Wireless Collisions. Collisions in power-constrained scenarios are problematic because the power used to transmit the packet is essentially wasted. This problem is aggravated when backscatter communication is used in general scenarios beyond traditional RFID/NFC systems that read and write small bits of information. As the packet sizes increase in general backscatter communication, the amount of power wasted in transmitting an undecodable packet increases proportionally. The problem is even worse in scenarios like ambient backscatter communication [131] where a network of battery-free devices communicate with each other without a central coordinator. In such networks, the presence of hidden terminals can aggravate the collision problem.

Multi-round protocols such as RTS-CTS are not attractive in this domain because the devices may not have enough power to transmit and receive multiple packets in succession.

An instantaneous feedback channel can help alleviate the problem of collisions and hidden terminals. Specifically, the receiver can send a message back to the transmitter informing it of the presence of a collision. Since the feedback arrives at the transmitter while it still is transmitting, the transmitter can terminate transmission immediately and avoid wasting power.

Rate Adaptation. Rate adaptation enables devices to adapt to changing channel conditions to achieve optimal throughput. The problem with backscatter devices is that they generally have low transmission rates. So low, in fact, that traditional rate adaptation reacts too slowly to be effective—channel quality can change significantly within the span of a single packet. This is important because, whereas maximum transmission time for an 802.11g packet is 542 µs [129], backscatter devices such as those proposed in [131] can take twice as much time to transmit a single bit. These devices can take up to 2s to transmit a 256-byte packet. Even an environment

[3] One might think that a solution to this problem would be to simply gather enough energy for multiple packets; in §3.4.1, we explain why storing more energy is not straightforward due to leakage issues.

that changes at a modest rate (e.g., a person walking) can create changes within that time frame. Note that reducing the packet size is not a desirable solution since it increases the overhead of the preambles and the headers significantly.

An instantaneous feedback channel can address this problem by delivering feedback about channel conditions within the duration of a single packet transmission. This is beneficial in scenarios like that of wearable devices [157], which can experience channel changes as the person moves in the environment.

Retransmissions. Finally, we consider the issue of retransmission. While forward error correction, which may include rate adaptation, can effectively decrease the frequency of failures, there are inevitably cases where certain bits are unrecoverable. This traditionally requires retransmission of the entire packet.

The challenge in backscatter communication is that retransmissions are very undesirable—more than in traditional communication. Not only is the penalty of failed/wasted transmissions much higher in these types of devices, so is the probability of errors. As an example, consider the problem presented in the preceding case. In dynamic channels, packets take so long to transmit that it is likely that the channel will change or interference will occur during the course of a single transmission. Even if only a single bit is unrecoverable, retransmitting an entire packet means that the energy spent in the transmission of the original packet was wasted.

An instantaneous feedback channel can be used to address the above issue by allowing the transmitter to only retransmit the set of bits that are incorrectly received. Specifically, the receiver can use the feedback channel to transmit checksums of small subsets of bits [214] during the packet transmission. The transmitter can use these checksums to identify subsets of bits that are incorrectly received. Note that since these checksums are transmitted on the feedback channel at the same time as the packet transmissions, they do not add additional transmission overhead to the system.

3.3 Full-duplex Backscatter Design

We introduce Full-duplex Backscatter, a novel communication technique that creates a low-rate, instantaneous feedback channel for backscatter devices. Our goal is to design a feedback channel that uses a single antenna for transmit and receive and consumes a negligible amount of power.

Full-duplex Backscatter builds on the ambient backscatter communication technique introduced in Chapter 2. However, creating such a feedback channel on backscatter devices is challenging. First, conventional backscatter transmitters actively change the amplitude of reflected signal to encode data. This is bad for the receiver, which relies on changes in the amplitude of the received signal to decode transmissions. Since the transmitter and receiver share a single antenna, a device's transmitter will create a great deal of interference for it's own receiver. Second, backscatter transmitters and receivers must share the same signal. Conservation of energy means that there is a fundamental tradeoff between the strength of the signal scattered by the transmitter and the strength of the signal received by the receiver. Finally, backscatter devices have severe power, area, and cost constraints that rule out traditional interference-cancellation techniques including multiple antennas and complex adaptive cancellation techniques.

3.3.1 Overview

Our approach has two key ideas. (1) The design of a transmitter that can backscatter information with minimal changes in the amplitude of its reflections. We show that such a design provides significant cancellation on our hardware prototypes. However, in practice we observe some residual interference. (2) To deal with the residual interference we exploit the rate difference between the data and feedback channels to design custom receiver circuits for the data and feedback channel receivers respectively. At a high level, the receiver of the feedback channel sees high-frequency residual interference, whereas the data channel receiver sees low-frequency residual interference. We use low-pass and band-pass filters designed using passive components to eliminate this interference. We describe these ideas in more detail below.

Traditional backscatter transmitters as discussed in Chapter 2 modulate the amplitude of re-

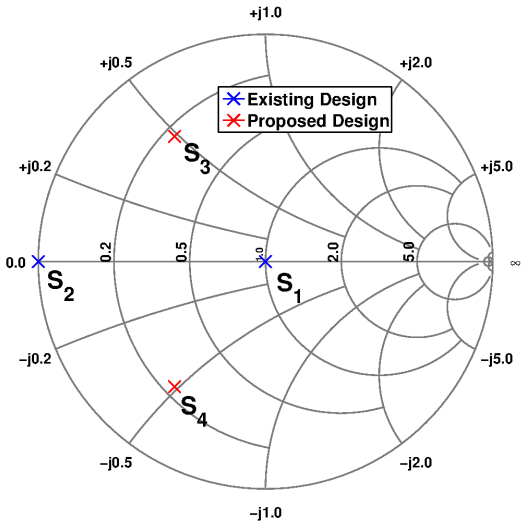

Figure 3.3: **Full duplex transmitter design:** In conventional backscatter transmitter, the antenna impedance switches between matched state (S_1) to a short impedance state (S_2). Instead in full-duplex backscatter design, we switch impedances between states S_3 and S_4.

flections to transmit bits. Instead, a Full-duplex Backscatter transmitter reflects/absorbs a fixed amount of signal, such that the amplitude of the received signal is constant. The device transmits information by switching between two impedances that modulate the *phase* of the reflected signal instead of the amplitude.

To formally, understand this design consider the Smith chart shown in Fig. 3.3. The transmitter switches between two conjugately-matched impedance states corresponding to S_3 and S_4. The two impedance states S_3 and S_4 are chosen such that $\left|\Gamma^*_{S_3}\right| = \left|\Gamma^*_{S_4}\right|$ and $\Gamma^*_{S_3} = \Gamma_{S_4}$. Hence the magnitude of the signal reflected in the two states is equal and the reflected signals are out of phase.

The two reflections with different phases from the transmitter will interfere with the ambient signal at the corresponding receiver to create waves with two different amplitude levels. We revisit the model developed in Chapter 2 to formalize this. When Alice transmits '0' and '1' bits by toggling the antenna impedance between S_3 and S_4 states, the absolute value of the voltage received on the Bob's receiver can be written as.

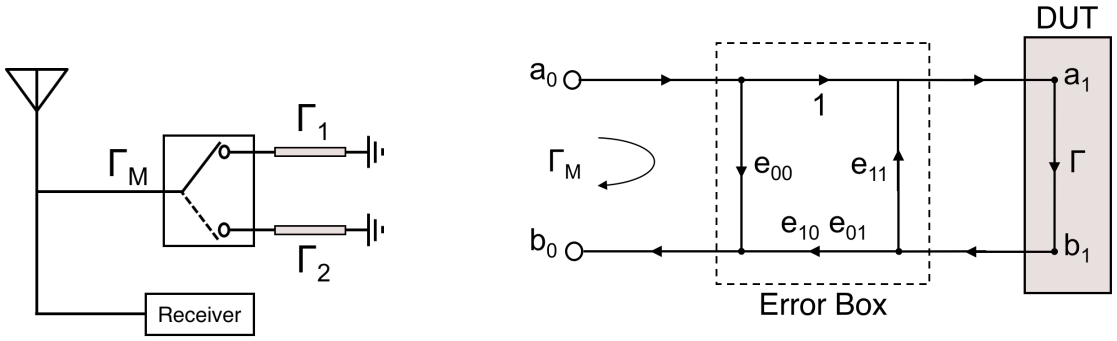

(a) Antenna and single pole dual throw switch model.

(b) One port error model for calibration.

Figure 3.4: **Calibration procedure:** Signal flow graph for the one port calibration model. The error box models the errors introduced by the components which have to be de-embedded using OSM (Open-Short-Matched) calibration procedure.

$$V^0_{Bob} = \frac{l_e}{2}\left|E_{direct,Bob}\right|\left|1 + k\frac{e^{-j\beta D}}{D}\left(1+\Gamma^*_{S3}\right)\right|$$
$$V^1_{Bob} = \frac{l_e}{2}\left|E_{direct,Bob}\right|\left|1 + k\frac{e^{-j\beta D}}{D}\left(1+\Gamma^*_{S4}\right)\right|$$

where l_e is the electrical length of the antenna. The two voltage levels V^0_{Bob} and V^1_{Bob} will be unique and Bob's receiver can decode the bits using an amplitude-tracking receiver. We note that the use of complex impedances to change the phase does not necessarily degrade the quality of transmissions. The quality of the transmitted signal is a function of the distance between the two impedance states on the Smith chart. Thus, the above approach can achieve equal or greater signal quality than traditional backscatter transmitters.

3.3.2 Full-duplex Backscatter Transmitter Design

The Full-duplex Backscatter transmitter consists of a switching network which modulates the impedance seen by the antenna between two conjugately matched impedance states corresponding to $\left|\Gamma^*_{S_3}\right|$ and $\left|\Gamma^*_{S_4}\right|$. Achieving these impedance states however, on a device whose architecture is shown in 3.4(a), is not straightforward. The antenna is connected to the receiver and a single pole double throw (SPDT) switch. The two output ports of the switch are connected to the two different

impedance networks corresponding to Γ_1 and Γ_2. The challenge is that the reflection coefficients at the antenna port, Γ_{M1} and Γ_{M2} are respectively not equal to Γ_1 and Γ_2. This is because the antenna is connected to a receiver in parallel to a non-ideal switch. The switch has losses and parasitics which transforms the impedance network. Secondly, the impedance seen by the antenna is a parallel combination of the receiver and the input impedance of the switch. Finally, the layout of the switch will add additional electrical length. As a result of all these effects, the impedance network is transformed and the reflection coefficient at the antenna port is not equal to Γ_1 and Γ_2 in the two states.

In order to achieve the desired impedance states $\left|\Gamma^*_{S_3}\right|$ and $\left|\Gamma^*_{S_4}\right|$, we have to first model this transformation. We use a one port error model which consists of three components: directivity (eel_{00}), port match (e_{11}) and tracking ($\Delta_e = e_{10}e_{01}$) and is shown in Fig. 3.4(b). Given these three parameters, the following transformation can be used to compute Γ_M from Γ and vice-versa.

$$\Gamma_M = \frac{e_{00} - \Delta_e\Gamma}{1 - e_{11}\Gamma} \qquad \Gamma = \frac{\Gamma_M - e_{00}}{\Gamma_M e_{11} - \Delta_e} \tag{3.1}$$

To estimate the parameters of the error box, we use the OSM (open, short and matched) calibration method used in VNA calibration. We connect open (Γ_1), short (Γ_2) and load (Γ_3) impedances at the switching network and measure the impedance states seen by the antenna $\Gamma_{M1}, \Gamma_{M2}, \Gamma_{M3}$ respectively. Then we use the following transformation to compute the error box parameters.

$$\begin{bmatrix} e_{00} \\ e_{11} \\ \Delta_e \end{bmatrix} = \begin{bmatrix} 1 & \Gamma_1\Gamma_{M1} & -\Gamma_1 \\ 1 & \Gamma_2\Gamma_{M2} & -\Gamma_2 \\ 1 & \Gamma_3\Gamma_{M3} & -\Gamma_3 \end{bmatrix}^{-1} \begin{bmatrix} \Gamma_{M1} \\ \Gamma_{M2} \\ \Gamma_{M3} \end{bmatrix}$$

Now, that we have the model for the impedance transformation, we can use (3.1) to determine the impedances that are required at the output port of the switch. Specifically, we set $\Gamma_{M1} = \Gamma_{S3}$ and $\Gamma_{M2} = \Gamma_{S4}$ and use it to compute Γ_1 and Γ_2 and set the corresponding impedances at the output ports of the switch. We use the ADG919 RF switch [5] as the SPDT switch in our transmitter design. To show how the impedance transformation modeled by the error box affects the design,

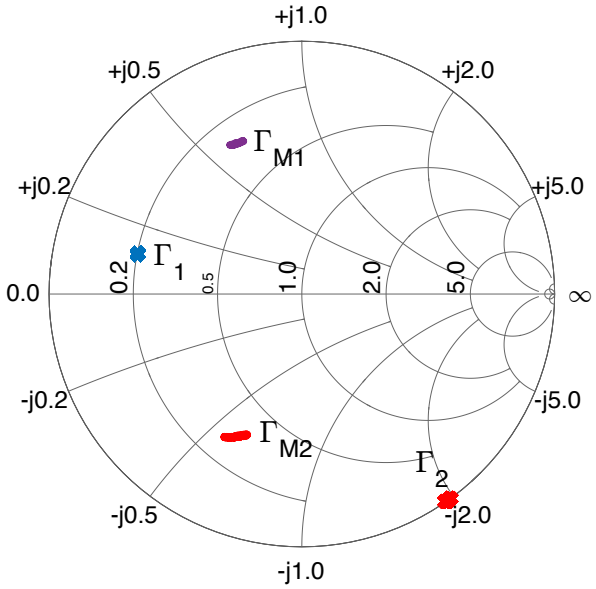

Figure 3.5: **Impedance states at antenna and switch port.** Smith chart with impedance states for Alice's transmitter measured at the antenna terminal and switch port.

we consider the data channel, i.e. transmitter at Alice. Γ_1 impedance states are implemented using a series combination of a 10 Ω resistor and a 1 nH inductor, while the Γ_2 state is implemented using a 1.8 pF capacitor. All the four impedance states are shown on the smith chart in Fig. 3.5. We can see that the two impedance states have undergone a linear and rotational transformation at the antenna port. Similarly, for the feedback channel on Bob, the transmitter's Γ_1 impedance state is implemented using a parallel combination of a 6.8 nH inductor and a 50 Ω resistor, while the Γ_2 state is implemented using a series combination of a 15 Ω resistor and a 0.8 pF capacitor. Note that these impedance values are specific to the implementation of the switch, receiver and PCB design. The static power consumption of the analog transmitter design (i.e. the switch) is about 0.25 μW.

In theory, the above technique should result in perfect cancellation; however, in practice, impedance mismatches occur due to component tolerances, and temperature variations which change the impedance of the receiver. These mismatches result in residual interference. We describe how our receiver addresses such residual interference in the next section.

3.3.3 Full-duplex Backscatter Receiver Design

The Full-duplex Backscatter receiver is based on the architecture described in Chapter 2. It consists of an envelope detector followed by a comparator, which performs either a low-pass filter or a band-pass filter operation. The envelope detector is used to remove the RF carrier frequency and extract amplitude information. The impedances at the output of the envelope detector and the comparator implement the passive filter and output a digital value. The goal of the receiver in Full-duplex Backscatter system is to remove the effect of the residual interference. We do so by leveraging the fact that the two devices are transmitting at different rates. As an example, consider Bob from Fig. 3.2 whose received signal is illustrated in Fig. 3.6(a).

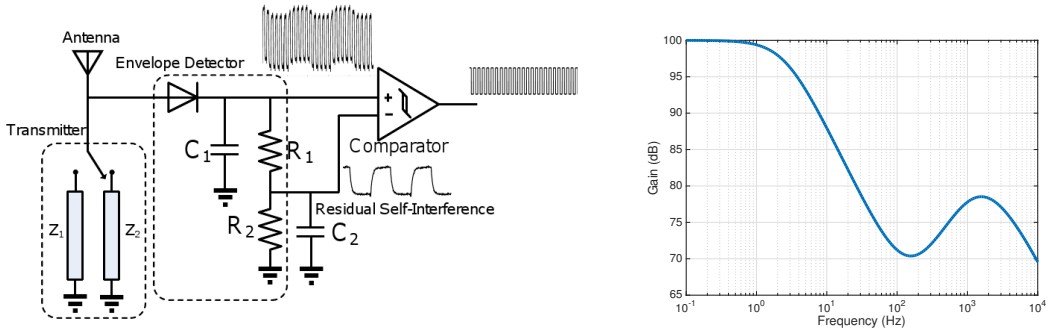

(a) Bob's transmitter and receiver architecture. (b) Gain of passive filter on Bob's receiver.

Figure 3.6: **Bob's Transmitter and Receiver:** Bob switches between Z_1 and Z_2 at a low rate to transmit data to Alice. Bob's receiver consists of three main components: an envelope detector to remove the carrier frequency, a low pass filter to isolate the low frequency residual self-interference and a comparator to cancel the residual self-interference from the output of the envelope detector and decode the received bits. In effect, by doing so, Bob is implementing a high pass filter by using a low pass filter to track the residual interference and subtracting it from the envelope signal using a low power comparator. The high pass operation cancels the low rate self-interference from the desired high rate signal.

We can use a low-pass filter consisting of resistors R_1, R_2, capacitors C_1 and C_2 and a comparator to track the slowly-varying self-interference and compute a time-varying threshold as shown in the Fig. 3.6(a). We feed the received signal to the positive terminal and the computed slowly varying threshold to the negative terminal of the comparator. This essentially performs a band-pass

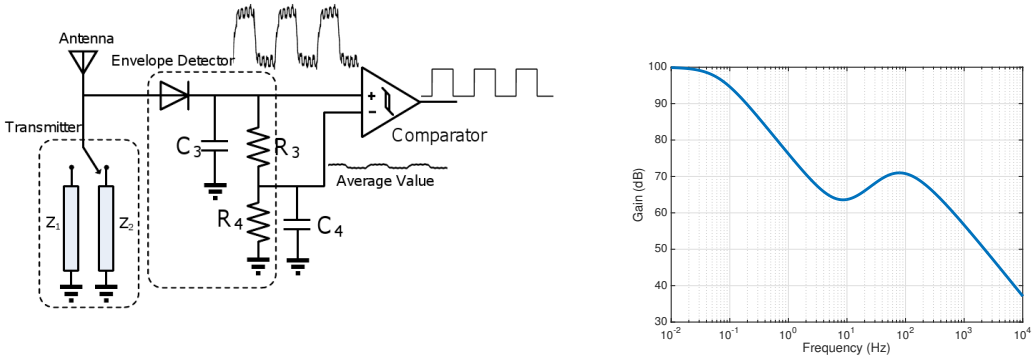

(a) Alice's transmitter and receiver architecture. (b) Gain of passive filter on Alice's receiver.

Figure 3.7: **Alice's Transmitter and Receiver:** Alice switches between Z_1 and Z_2 impedances at a high rate to transmit data to Bob. The receiver on Alice consists of two main components: an envelope detector/low pass filter to remove the carrier frequency and self-interference and another low pass filter to track the average value. These two signals are fed to a comparator to threshold the received signal and decode the digital bits sent by Bob.

filter operation that subtracts the low rate interference from the received signal and enables Bob to efficiently decode Alice's transmissions. The bode plot of the band-pass filter implemented on Bob's receiver is shown in Fig. 3.6(b). We use a TS881 [59] as the ultra-low power comparator. The capacitor and resistor values R_1, R_2, C_1, and C_2 on Bob are set to 100 $k\Omega$, 10 $M\Omega$, 4.7 nF and 1.47 nF, respectively.

The operations at Alice, as shown in Fig. 3.7(a), are analogous but different. Specifically, we implement an envelope detector using different C_3, C_4, R_3, and R_4 that suppresses both the RF carrier and the high-frequency self-interference. We then use a low pass filter consisting of R_4 and C_4 to compute the long-term average of the envelope and subtract this average from the output of the envelope detector to decode the low-rate feedback transmissions from Bob. From the perspective of receiving data at 100 bps, this is a low pass filtering operation and the corresponding bode plot is shown in Fig. 3.7(b). The capacitor and resistor values R_3, R_4, C_3, and C_4 on Alice are set to 100 $k\Omega$, 10 $M\Omega$, 27 nF, and 220 nF respectively.

The analog prototype receiver (including the comparator) on both Alice and Bob consume 0.54 μW. We note that in principle, it is possible to use a conventional amplitude modulating transmitter in combination with the presented residual interference cancellation techniques to en-

able a simultaneous feedback channel. However, in practice, use of such conventional transmitters results in a residual interference of about 100 dB (relative to the noise floor). In order to eliminate such residual interference, we need to implement high-order filters. Given the losses associated with passive components such as capacitors and resistors and loading effects of cascaded passive filters, a practical implementation of such filters using only passive components is not a feasible in practice. Hence, we believe that our proposed phase modulating transmitter architecture followed by the residual cancellation technique is essential to the design of an instantaneous backscatter feedback system.

3.3.4 Integrating Transmitter and Receiver on the device

Thus far we have described Alice and Bob's designs separately. In reality, every device must take on multiple roles depending on whether it has a packet to send, wants to receive a packet, or is simply stay idle. So, we need to consolidate the two design illustrated in Fig. 3.7 and Fig. 3.6 and engineer the system to maximize receive and harvesting efficiency.

A straightforward way to consolidate the designs is to include hardware for both forward and feedback channels and the device can opportunistically switch based on its operating state. In particular, the transmitter needs to have four impedance values—two for the forward channel and two for the feedback channel. Similarly, devices also need two receive chains for each target transmission rate in order to implement the low-pass and band-pass filters. However, for the receiver in an ASIC or COTS implementation, tunable resistors and capacitors can be used to eliminate hardware redundancy between the different states.

For both reception and harvesting, efficiency is highest when the matched impedance is close to S_1 in Fig. 3.3 (where it reflects the least amount of energy back on the antenna). Since the transmitter design on the feedback channel picks impedance values S_3 and S_4 that are necessarily distinct from S_1, the receiver efficiency of the data channel is lower than existing backscatter receivers. To address this issue, we jointly optimize the impedance values at the transmitters of both the data channel (Alice) and feedback channel (Bob) to ensure that the efficiency (receive BER) of the data channel remains the same. This translates into impedance values that are closer to the

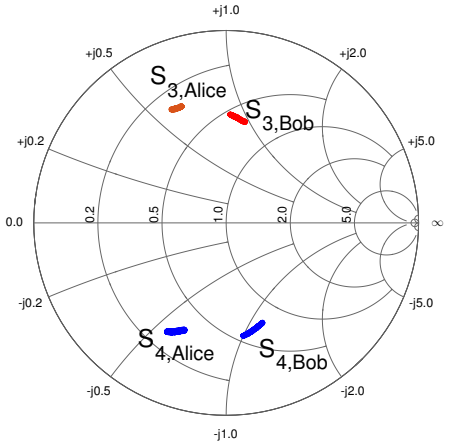

Figure 3.8: **Alice and Bob's transmitter impedance states:** The conjugately matched impedance states for Alice and Bob's transmitter represented on a smith chart. Alice's impedance states are optimized for a higher data rate transmitter whereas Bob's impedance states are optimized for higher data receiver.

origin for Bob's feedback transmissions and those that are further from the origin for Alice's data transmissions a shown in Fig. 3.8. In our prototype implementation, this results in a feedback rate that is one-tenth of the data rate.

Additionally, the design can have an additional impedance value for the case when there are no ongoing transmissions. During this state, the device does not need to transmit any signals and therefore can fully optimize for receive/harvesting efficiency by staying in the S_1 state. This is important because during transmission, the forward channel has a harvesting efficiency of 44.83% compared to 50% in a traditional backscatter system. Because communication periods are typically very short in energy-constrained systems, staying in a matched state during idle periods makes these losses negligible.

The Full-duplex Backscatter prototype shown in Fig. 3.1 was implemented on four-layer printed circuit boards (PCB) using off the shelf components. We uses a dipole antenna consisting of two sections of 3-inch long 16 AWG magnetic copper wire. We implemented prototypes with 1 kbps transmission on the data channel and 100 bps on the feedback channel. The hardware is tuned to operate at frequencies in the 3 MHz band centered at 920 MHz. Since the bit rates used by backscatter data communication are significantly lower than traditional radio communication (e.g.,

Wi-Fi), a 3 MHz data bandwidth does not limit the underlying RF source bandwidth. Specifically, as long as the data bandwidth is less than 3 MHz, the system would work even with an underlying RF source with higher bandwidth (e.g., TV signals). This is because the circuit would low-pass filter out the higher frequency components.

3.4 A Link-layer Design for Full-duplex Backscatter

The hardware design in the preceding section enables an entirely new set of protocols for general-purpose backscatter devices. In this paper, we explore one example instantiation of a link-layer protocol for backscatter communication. In particular, our protocol tackles the three problems described in §3.2: wireless collisions, error correction, and rate adaptation.

In this section, we first explain how RF power harvesting can affect protocol design. We then describe our link-layer protocol that uses the feedback channel to address the above problems.

3.4.1 Design Principle

Our protocols strive to minimize the amount of energy required to transmit a single packet successfully. To see why this is desirable, we need to look at the workings of an RF energy harvester.

The purpose of a device's harvesting circuitry is to charge up a capacitor that can be discharged whenever the device needs to perform computation, sensing, or communication.

At a high level, harvesting happens in two stages: rectification and energy storage. When a signal arrives through the antenna, it translates into a voltage value, say, $V(t)$. The rectification stage, in addition to converting the time-varying AC signal to a DC voltage, also acts as a voltage multiplier, i.e., it raises the voltage of the signal by trading off current. The resulting DC voltage is applied to the storage capacitor in order to charge it up.

In an ideal world, one could convert arbitrarily weak voltage signals into arbitrary amounts of energy in two ways:

- *Increasing the voltage value.* Adding more stages to the rectification stage will increase the voltage multiplication effect [210]. This can potentially be used to increase the voltage applied

to the capacitor and therefore the total energy stored by the capacitor.

- *Increasing the capacitor size.* Increasing the capacitance of the capacitor increases the amount of charge it can hold for a given voltage according to $Q_{max} = CV$.

Unfortunately, in practice, both of the above approaches come with significant tradeoffs that limit their applicability. First, adding more stages to the voltage multiplier decreases harvesting efficiency and also significantly increases power losses due to parasitic capacitances and leakage effects [130]. Second, increasing the size of the capacitor increases the charge time disproportionately. For instance, doubling the size of the capacitor would increase the time to fully charge by more than $8x$. Furthermore, larger capacitors generally have greater leakage currents which decreases the sensitivity of the harvester and increases the charge time even further.

The above discussion implies that we need to leverage our feedback channel to minimize the need for multiple packet transmissions that requires significant amount of energy.

3.4.2 Link-layer Protocol

An instantaneous feedback channel allows us to design a MAC protocol that minimizes the amount of energy required to send a packet. Our protocol reduces the penalty associated with collisions and aids in detection of hidden terminals, all with minimal energy requirements. It is inspired by a variety of related work including CSMA/CD [144], busy tones [109], μACKs [214], and others [146], but the underlying mechanism is simple: use the feedback channel to acknowledge data transmissions at a fine granularity.

For ease of exposition, we describe the feedback and data channels separately. We begin by describing the feedback channel operation in the common case; we later describe how the data channel takes advantage of feedback.

Feedback Channel

Our system uses the feedback channel to acknowledge ongoing data transmissions. We consider a source sending a packet using the format shown in Fig. 3.9. The destination uses the feedback

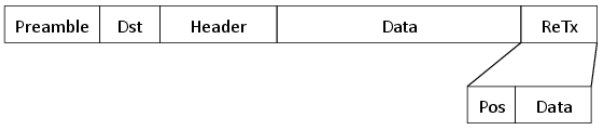

Figure 3.9: **Packet format.** The packet format used on the data channel. At the end of every packet, the transmitter can optionally append retransmissions of bit chunks and their positions.

channel to perform the following operations:

- The destination first decodes the transmitted preamble and the first field of the header (i.e., the destination address). As soon as the destination realizes that it is the intended recipient, it begins transmitting a preamble on the feedback channel.

- The destination will then divide the packet (including header) into chunks of b bits. For each group of b bits, it computes a c-bit checksum. The destination transmits the checksum back to the source on the feedback channel.

The ratio of b to c is determined by the difference in transmission rate between the data and the feedback channels. For instance, a 1 kbps data channel with a 100 bps feedback channel, must satisfy the following condition: $\frac{b}{c} = 10$. The time for both the data and feedback transmissions are therefore approximately equal, with the feedback channel lagging behind the data channel slightly. Our prototype implementation uses 40 and 4 bits for b and c respectively.

Data Channel

The source uses the above feedback channel to adapt its own transmissions to errors and collisions. In particular, it performs the following protocol:

- The source first listens on the wireless medium to ensure there are no existing transmissions.[4] If the channel is empty, it begins to transmit the preamble, header, and payload. Otherwise, it exponentially backs off before retrying. See §3.4.3 for details.

[4] In backscatter-style communication, this can be effectively implemented by simply checking the output of the receive chain for bit transitions [131].

- While transmitting, the source continues to listen on the feedback channel for errors. It uses the incoming checksums to change rates at the level of bits rather than entire packets. If it detects a collision, it will terminate the transmission and back off accordingly. See §3.4.4 and §3.4.3 for details.

- At the end of the packet, any failed groups of bits are retransmitted along with a short header denoting their position in the original stream. See §3.4.5 for details.

3.4.3 Dealing With Collisions

There are two places where collisions can occur: (a) at the beginning of the packet, during the preamble or destination address and (b) in the middle of the packet. Collisions at the beginning of the packet can occur when two nodes start to transmit simultaneously or when the new transmitter cannot hear the existing connection. These collisions can either interfere with the forward channel or the feedback channel, but in both cases, no feedback will return and the data channel transmitter will assume a collision has happened. If the collision happens in the middle of a packet, we leverage our bit-level rate adaptation to account for the resulting bit errors. Further, devices assume that a collision has occurred after multiple consecutive drops happen in the data rate during our bit-level rate adaptation.

Back off is implemented similarly to existing protocols: the detecting transmitter will wait for a random number of time slots between 0 and 2^r, where r is the number of retries attempted by the device. Our link-layer protocol detects collisions, even from hidden terminals, and minimizes their effects. We note the following about the effect of collisions in our system.

Firstly, by terminating the packet as soon as a collision is detected, our design can minimize the energy penalty associated with collisions. Specifically, recharging the energy spent during the short amount of time before collision detection at the transmitter takes significantly less time than recovering the energy required for an entire packet.

Secondly, the transmitter that detects the collision does not need to jam/inform the any other transmitters because bit-level error correction can correct from any collision-related errors. As

long as the preamble and header are decoded correctly, any subsequent bit errors can be corrected as long as the network ends up with a single active transmitter-receiver pair.

Thirdly, the feedback channel is compatible with our collision detection technique because transmitters should see a checksum of their own transmissions. Any competing transmission will cause the bits on the feedback channel to differ from the expected bits significantly, allowing the transmitter to detect a collision.

Finally, we note that false positives in detecting collisions can occur due to bit errors; however, a conservative collision detection estimate is preferable to retransmission of the entire packet.

3.4.4 Adapting Rate at a Fine Granularity

Rate adaptation proceeds in a fashion similar to existing techniques, except at the level of b-bit chunks rather than entire packets. The exact algorithm used is orthogonal to our work, but we take [192] as a baseline. At a high level, the goal of the rate control algorithm in [192] is to maximize throughput while occasionally testing alternative rates ($\sim$10% of the time). Throughput for a particular rate is defined by $(\textit{Probability of success} * \textit{Transmission rate})$, where the probability of success is based upon an exponentially-weighted, moving average of chunk-level success rates. If through the occasional probes, we find an alternative rate with a higher throughput, we will switch to that rate. If the line code is self-clocking, the receiver will adjust to rate changes automatically.

The above protocol allows us to adjust rates at the level of chunks of b bits, rather than entire packets. Adjusting the rate at a fine granularity gives devices the ability to react very quickly to changes in channel state. More importantly, we can react within the span of a single packet—a useful benefit since bit rates are often very low in backscatter systems. It also allows us to react after fewer packet drops, as looking at small chunks of bits gives us a much greater sample size for rate adaptation rather than just coarse packet drops.

3.4.5 In-Packet Error Correction

Finally, the source can use the feedback channel to decrease the energy penalty of bit errors. Specifically, error correction in our system proceeds in a manner similar to that of [214]: the transmitter resends failed bit sequences at the end of each packet along with their position in the original stream.

This type of error correction is necessary to make the above bit-level rate adaption useful—without it, errors in the packet would make the entire packet useless. The additional benefit to being able to tolerate bit errors in this fashion is (as mentioned above) that collisions no longer necessarily cause all transmissions to be wasted.

3.5 Evaluation

Next, we evaluate various practical aspects of our Full-duplex Backscatter prototype. In particular, we first measure the effect of our low-power backscatter cancellation technique described in §3.3.2 as a function of the input power level. Next, we evaluate the effect of the feedback channel on the bit error rate (BER) of the data channel at the receiver. We also evaluate the bit error rate (BER) achieved on the feedback channel as a function of the distance between the backscatter devices. Finally, we evaluate the feedback channel with collisions, retransmissions, and rate adaptation.

3.5.1 Full-duplex Backscatter Cancellation Effectiveness

In this section, we evaluate how well our design in §3.3.2 reduces the self-interference from the feedback channel.

Experiments: To do this, we examine the degree to which we decrease self-interference at the destination (i.e., the device sending the feedback) using the transmitter cancellation (described in §3.3.2) and the residual cancellation techniques (described in §3.3.3). We place the Full-duplex Backscatter prototype device in the presence of a continuous wave transmission from an RFID reader at 920 MHz and configure it to continually transmit an alternating sequence of bits at 100 bps. Note that this is one feedback bit for every 10 data bits on an ambient backscatter commu-

nication system, which, as we show later, is sufficient to address the networking issues described in §3.2. We then tap into the outputs of the envelope detector and the low pass filter on our prototype board and connect it to an ADC to get direct access to the voltage values.

Results: Fig. 3.10 shows the magnitude of the voltage values (in dB) across the frequency spectrum. Each plot shows the frequency-domain representation of the received voltage values with just the transmitter cancellation technique, with both transmitter and receiver cancellation techniques (to cancel the residual interference as described in §3.3.3) and without our cancellation technique. They also depict noise values when there is no transmission on the feedback channel. We plot the results for three different received power levels of the RFID continuous wave signals at the receiver. The graphs show the following:

- In the absence of our cancellation circuit, the plots show spikes at frequencies corresponding to the odd multiples of 100 Hz. This is expected because the bit rate on the feedback channel is 100 bps which results in a spike at 100 Hz. Further, since the transmitted signal on the feedback channel is approximately a square wave, we also see spikes at frequencies that are odd multiples of 100 Hz.

- Our technique reduces the interference from the feedback channel to the noise floor of our device across the frequency range. This is very significant and impressive since, this approach requires near-zero power compared to a conventional device.

- The reduction in self-interference is similar for different power levels of the received signal. This results in higher power levels having higher residual interference (in comparison to noise). In the next section, we will demonstrate how this cancellation technique results in comparable receive bit error rates for systems with and without the feedback channel.

- The self-interference reduction varies by about 5-8 dB across time and also the 3 MHz operational bandwidth. This variation is expected because the conjugate impedance network used for the cancellation is implemented using off the shelf resistors, capacitors and inductors that have tolerances and variations that change with environmental conditions.

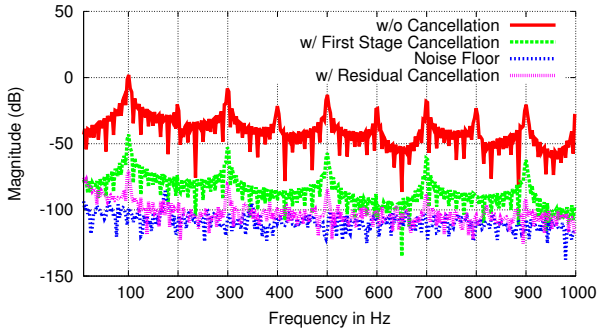

Power level at receiver: -1.7 dBm

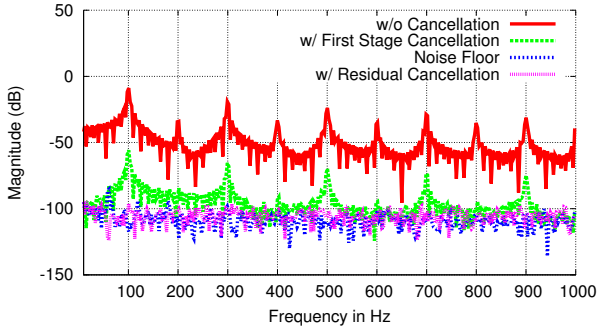

Power level at receiver: -20 dBm

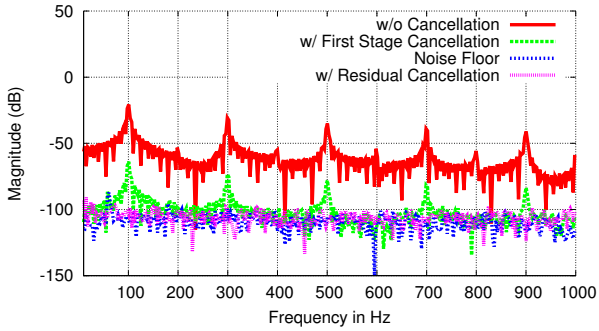

Power level at receiver: -30 dBm

Figure 3.10: **Self-interference Cancellation in Full-duplex Backscatter:** The graphs show the strength of the voltage signal received over a frequency spectrum due to the device's own 100 bps transmissions. Our technique reduces the self-interference from the feedback channel down to the noise floor of the device across the frequency range. We note that the typical power level at the receiver is less than -15 dBm.

3.5.2 *Effect on Data Channel BER*

Experiments: We place two prototype devices at different distances from each other. The devices are both configured to transmit an alternating sequence of bits at the same time—the sender transmits data bits at a rate of 1 kbps, and the receiver of the data packet transmits bits at a rate of 100 bps on the feedback channel. We place the two devices in the presence of an RF signal source that broadcasts continuous wave RFID signals centered at 920 MHz. The source has a dipole antenna and is placed equidistant from both the tags. To analyze the BER, we connect an NI myDAQ to the output of the receiver's receive chain. We capture 100 seconds of data at each distance value, which corresponds to a total of 10^5 bits; when no bit errors occur, we set the BER to 10^{-5}. We consider distance values of up to 1.5 meters, which spans the communication range of the ambient backscatter devices in [131]. We also vary the RF power to span the receive power levels from -18 dBm to -4 dBm; the lower value is the minimum power level at which the harvester [166] works.

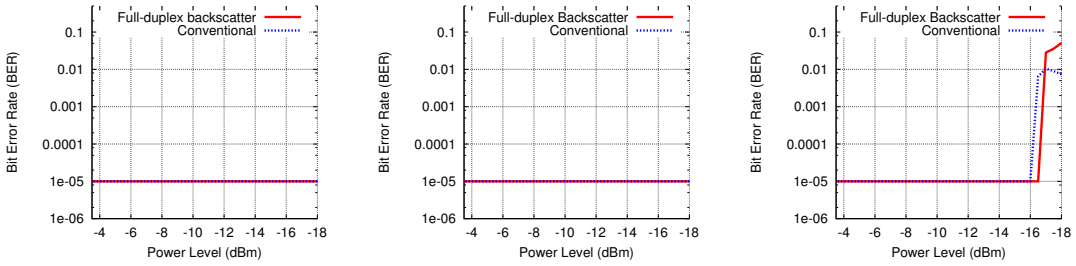

Distance: 0.75 meters Distance: 1 meters Distance: 1.5 meters

Figure 3.11: **Data BER versus power:** BER as observed by the receiver of the forward data channel. We show the variation in BER versus power at the tag for three different tag-to-tag distances. The plots show that the feedback channel does not significantly affect the data BER.

Results: Fig. 3.11 shows BER versus the average received power at our prototype devices for three different tag-to-tag distances. We also plot as a baseline the bit error rate results for the ambient backscatter prototype that we replicate from [131] that does not have a feedback channel. The results show the following:

- For both Full-duplex Backscatter and existing backscatter systems, as the distance between the devices increases the BER increases. Similarly, BER reduces as the RF power level at the prototype devices increases.

- We do not observe any bit errors for both the systems when the distance between the tags is less than or equal to 1 m. As the distance increases to 1.5 m, we start observing bit errors for lower power values. We note that in these cases, the observed bit error rate for our system across all three distances is comparable to conventional backscatter.

- In some cases, the bit error rate for our system is worse and in other cases better than existing backscatter systems. This is because existing backscatter systems modulate information by changing the amplitude of the transmission (i.e., using ASK modulation). In contrast, we encode information in phase values (i.e., effectively use PSK modulation). Thus, the signals in our design interfere differently than those of a traditional backscatter system and hence exhibit the above noted behavior. The key point however is that, the feedback channel in our results does not significantly affect the performance of the forward data channel.

3.5.3 BER of the Feedback Channel versus Distance

Next, we evaluate the performance of the feedback channel. Specifically, we measure the bit error rate (BER) on the feedback channel.

Experiments: As before, we place two prototype devices in the presence of an RF source broadcasting an RFID continuous wave transmission at 920 MHz. The data rate and feedback rate are set to 1 kbps and 100 bps, respectively. We vary the distance between the two devices and measure the bit error rate on the feedback channel. We again measure the BER with the assistance of an NI myDAQ and compare the received feedback bits with those transmitted. We transmit a total of 10^{-4} bits in each experiment. Note that this value is less than that of the previous set of experiments since the feedback channel has a lower bit rate than the data channel. We set the bit error rate to 10^{-4} in experiments which do not see any bit errors. We run the experiments at a fixed

transmit power of 7.5 dBm at the RF source. The BER trends are similar at other power levels.

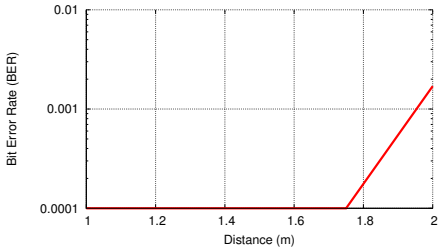

Figure 3.12: **Feedback BER versus distance:** BER observed on the feedback channel. We show the variation in BER versus distance between the two devices.

Results: Fig. 3.12 shows the BER of our feedback channel at different distances between the two prototype devices. The figure shows that the feedback channel sees bit errors as we approach 2 meters between the transmitter and receiver. There are no bit errors even at distances greater than 1.7 meters. Further, the observed feedback BERs match very well with the observed BERs on the data channel. This implies that the feedback channel is reliable enough to not be a limiting factor.

3.6 Evaluating Full-duplex Backscatter's Network Stack

Finally, we evaluate how our feedback-channel-enabled network protocol addresses collisions, retransmissions, and rate adaptation.

3.6.1 Recharge-time Reduction during Collisions

As described in §3.4.3, a feedback channel allows our prototype device to terminate its transmission as soon as it detects a collision. We will use the metric of recharge time to evaluate the benefit of this feature. Terminating wasteful transmissions can result in a reduction in the recharge time at the transmitter since it can conserve power that is otherwise wasted in transmitting an undecodable packet. Furthermore, the same principle applies on the receiver and the receiver can also conserve power as it does not to spend energy on decoding collided packets. This metric is very useful because recharging and duty cycling tend to be the bottleneck in virtually all energy-constrained systems.

Experiments: To evaluate the benefits of Full-duplex Backscatter in this context, we measure the time it takes to recover from a collision. Our prototype devices use a preamble length of eight bits and transmit at a rate of 1 kbps on the data channel (similar to the design in [131]). The feedback channel also uses a preamble length of eight bits and has a bit rate of 100 bps. We measure the time it takes for the transmitter to detect an existing transmission and then terminate its own. We then calculate the total delay from the beginning of the canceled transmission to the time at which the device has enough power to retry transmitting the packet, given a packet size of 64 bytes. We repeat the measurements for both scenarios, with and without the feedback channel. In the absence of the feedback channel the transmitter continues transmitting the whole packet even in the presence of the collision.

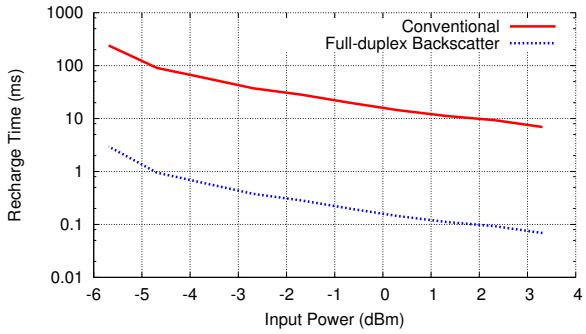

Figure 3.13: **Recharge time reduction:** The time needed to recover from a packet collision in Full-duplex Backscatter versus conventional transmitters. The x-axis plots the available power at the device. Our system reduces the recharge time by two orders of magnitude.

Results: Fig. 3.13 shows results for the above experiments. We plot the recharge time it takes to collect enough power to retransmit the collided packet as a function of the power level received at the prototype device. The figure shows the graphs for both our system and conventional systems. The results show the following:

- Our system decreases the time between retries by two orders of magnitude. The savings are expected to be higher for packet sizes longer than 64 bytes. This is because recharge times are nonlinear—as the number of wasted bits grows, the recharge time non-linearly increases.

Thus terminating the transmission earlier minimizes the recharge time it takes to collect enough energy for packet retransmissions.

- The recharge time is higher at lower power levels. This is expected because in power-constrained scenarios, the capacitor takes longer to collect the required energy. This is further complicated by leakage issues as explained in §3.4.1. We note that Full-duplex Backscatter can provide orders of magnitude reduction across all power levels, specifically in power-constrained scenarios.

3.6.2 In-frame Error Correction

We next look at the effectiveness of our feedback channel in achieving in-frame error correction. The error correction mechanism described in §3.4.5 decreases latency and increases throughput. This is because it allows us to recover from bit errors by sending a few extra bits at the end of a packet, rather than re-sending the entire packet. We evaluate these benefits in practice.

Experiments: The backscatter transmitter sends bits at a bit rate of 1 kbps using a packet size of 64 bytes. The receiver computes a 4-bit checksum for every chunk of 40 received bits. It then transmits these checksum bits on the feedback channel to the transmitter. The transmitter uses the checksum to detect when chunks of bits are incorrectly decoded at the receiver. The transmitter then retransmits these chunks at the end of the packet as described in §3.4.5. We change the distance between the transmitter and the receiver to span a range of bit error rates. We calculate the additional bytes of data that are necessary to deliver a fully correct packet using our prototype. We also repeat the experiments for existing backscatter systems that have to retransmit complete packets until at least one is successfully decoded.

Results: Fig. 3.14 shows the results of our experiments. The figure plots the overhead, i.e., the number of extra bytes necessary to deliver a single, correct packet as a function of the bit error rate observed on the data channel. The average overhead is plotted in log-scale. The figure shows the following:

- The average overhead incurred by our system is at least an order of magnitude less than in

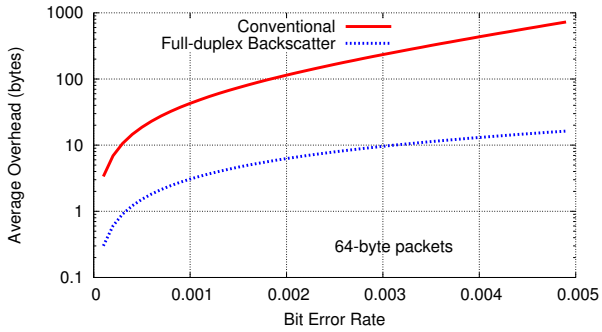

Figure 3.14: **Overhead of transmitting a single packet:** This graph shows the average number of extra bits transmitted for a given error rate. The extra bits are either in the form of a retransmitted packet in the case of conventional receivers or retransmitted bit sequences in the case of Full-duplex Backscatter. Our system reduces the overhead by at least an order of magnitude.

conventional systems, even for low error rates. This is because, for a single error, our system would need to transmit an additional 47 bits. On the other hand, a conventional system would need to retransmit the entire packet (i.e., 64 bytes), assuming the second packet arrives error-free.

- Further, as the average bit error rate on the data channel increases, the overhead for existing backscatter systems increases much faster than the overhead incurred on our system. In Fig. 3.14, the Full-duplex Backscatter system has an order of magnitude less overhead for low error rates, but as the BER increases to 0.005, the overhead increases two orders of magnitude. This is due to the fact that the probability that a 64-byte packet is correctly decoded becomes increasingly small at higher BERs. On the other hand, the probability that a chunk of 40 bits is decoded without error grows at a much slower rate.

- We note that Full-duplex Backscatter's advantage over conventional backscatter systems in both of the above points increases with larger packet size. We also note that the actual cost of this overhead in terms of time, energy, and throughput is much higher than shown here. The underlying issue is recharge time. As we saw in the previous section, charge time magnifies the penalty associated with retransmission of extra bits. Full-duplex Backscatter can reduce this

charge time and hence alleviate these issues.

3.6.3 Rate Adaptation

Finally, we evaluate our intra-packet rate adaptation.

Experiments: We place a transmitter and a receiver in the presence of our 920 MHz signal source. We adapt the rate of the transmitter between three rates of 100 bps, 1 kbps, and 10 kbps. We compare three rate adaptation strategies: "slowest", which always chooses 100 bps; "packet-level", which models an idealized algorithm that chooses the best rate for the SNR at the beginning of the packet; and "bit-level", which implements the algorithm described in §3.4.4. Since we would like to compare the three algorithms in the same scenarios, we connect our receiver to a myDAQ that continually takes measurements of the received voltage values. The transmitter is set to continuously transmit bits as we move the receiver in the area around the transmitter at an average speed of about 3 m/s. From the captured traces, we compute the SNR at every time instance. Since the SNR determines the achievable bit rate, we use it to compute the achievable throughput of the three algorithms given the captured voltage traces.

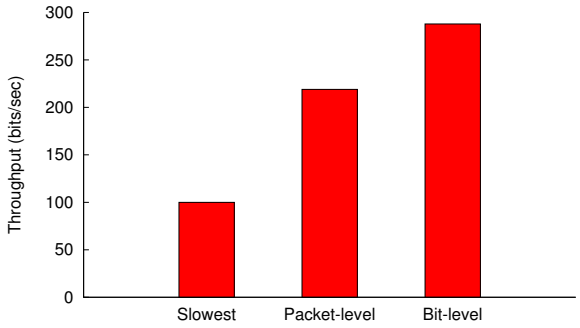

Figure 3.15: **Effect of Rate Adaptation:** A graph of the throughput for different rate adaptation algorithms. The throughput is calculated using real channel traces with an average mobility of 3 m/s.

Results: In Fig. 3.15, we plot the throughput of each rate adaptation strategy. Our results show the following:

- Bit-level rate adaptation increases performance by $\sim$33% compared to an idealized SNR-based packet-level algorithm. Further, when compared to the slowest-rate strategy, the speedup is

almost 3x. This is because devices can take advantage of and protect against channel changes. We also note that since backscatter devices do not have ADCs, they do not have SNR information. Thus, we can expect higher throughput gains in practice.

- We note that about 30% of packets in the packet-level strategy failed due to changes in the channel within the span of a single packet. This is due to the fact that transmission rates are extremely low in these types of systems.

3.7 Related Work

There has been recent interest in improving the performance of backscatter systems [79, 122, 180]. Prior work has also proposed coding techniques to address the problem of collisions and increase the efficiency of RFID networks [77, 148, 209]. More recently, work on ambient backscatter [131] enabled two battery-free devices to communicate by backscattering signals from an RF source. Our work builds on these foundational works and achieves the first instantaneous feedback channel for backscatter communication systems.

Our work is also related to recent research on full-duplex communication that uses passive cancellation techniques [109], active cancellation (both in the analog and digital domain) [72, 190], or a combination of the above [86, 87]. While these approaches are effective for traditional radio communication, their reliance on space-consuming antenna arrays and/or power-hungry cancellation techniques are not applicable to ultra-low power backscatter communication, that has a power budget which is orders of magnitude lower. We note that our goal is not to design a full-duplex radio. Instead, we set out to create a low-rate feedback channel for backscatter devices that can address many of the higher-layer issues facing these devices.

Similarly, there has also been previous work on implementing full-duplex communication through clever combinations of signal processing, modulation and coding techniques [71,102,136]. Like the above approaches, these assume that at least one of the participants are powered and can therefore perform relatively complex synchronization and decoding that are not applicable to the types of devices that we investigate in this paper.

Full-duplex Backscatter is also related to prior work that implements QAM transmissions on RFID tags by using complex impedances [199]. Our work is similar in that we also use complex impedances to adjust the phase of the reflected signal; however, the purpose of our system is complementary and is to enable a feedback channel between the backscatter devices and demonstrate the benefits of such a channel for the link- and the network-layers.

Finally, there is a large body of work related to network-layer protocols for wireless sensor network or other types of mesh networks. Some of these works deal with half-duplex mesh networks and how to implement MAC protocols [194], rate adaptation [117], and hidden terminal detection [193]. Others are designed for duplex communication [109]. In contrast, the focus of this paper is to enable an instantaneous feedback channel for battery-free backscatter devices, which is a goal that is complementary to prior efforts.

3.8 Conclusion

Energy, and specifically the energy required for communication, is a key bottleneck in the design of computing devices. Recent developments in backscatter communication promise to remove this bottleneck. Specifically, they promise to provide a way for devices to send bits to one another using orders of magnitude less power than is required today. However, an essential question remains: how do we design link- and network-layer protocols for these networks? Existing protocols are ill-suited to these devices where even the transmission of a single packet can exhaust all available energy.

In this work, we introduce a novel technique called Full-duplex Backscatter that enables almost-zero-power, instantaneous feedback and use it to implement a network stack tailored for battery-free devices. Our technique uses passive, analog circuitry in order to allow for a low-rate feedback channel that operates alongside any data transfer. In addition, we presented a network stack that uses Full-duplex Backscatter to minimize the energy wastage associated with MAC protocol, rate adaptation, and error correction. We believe that our technique improves the practicality of battery-free devices and brings us closer to having a practical, generalized backscatter communication system.

Chapter 4

POWERING THE NEXT BILLION DEVICES WITH WI-FI

In the late 19th century, Nikola Tesla dreamed of eliminating wires for both power and communication [198]. As of the early 21st century, wireless communication is extremely well established—billions of people rely on it every day. Wireless power, however, has not been as successful. In recent years, near-field short range schemes have gained traction for certain range-limited applications, like powering implanted medical devices [207] and recharging cars [83] and phones from power delivery mats [82, 108, 133]. More recently researchers have demonstrated the feasibility of powering sensors and devices in the far field using RF signals from TV [131, 186] and cellular [166, 205] base stations. This is exciting, because in addition to enabling power delivery at farther distances, RF signals can simultaneously charge multiple sensors and devices due to their broadcast nature.

In this work, we show that a ubiquitous part of wireless infrastructure, the Wi-Fi router, can provide far-field wireless power without significantly compromising network performance. This is attractive for three key reasons:

- In contrast to TV and cellular transmissions, Wi-Fi is ubiquitous in indoor environments and operates in the unlicensed ISM band, where transmissions can be legally optimized for power delivery. Repurposing Wi-Fi networks for power delivery can ease the deployment of RF-powered devices without additional power infrastructure.

- Wi-Fi uses OFDM, an efficient waveform for power delivery because of its high peak-to-average ratio [201, 202]. Given Wi-Fi's economies of scale, Wi-Fi chipsets provide a cheap platform for sending these power-optimized waveforms, enabling efficient power delivery.

- Sensors and mobile devices are increasingly equipped with 2.4 GHz antennas for communi-

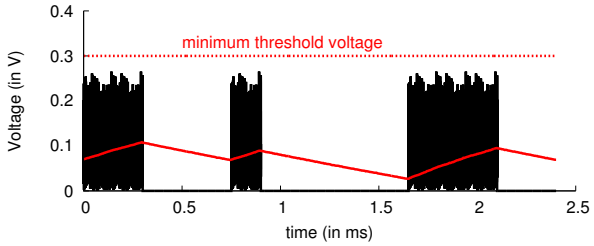

Figure 4.1: **Key challenge with Wi-Fi power delivery.** While the harvester can gather power during Wi-Fi transmissions, the stored energy leaks during silent periods, limiting Wi-Fi's ability to meet the minimum voltage requirements of the PoWiFi.

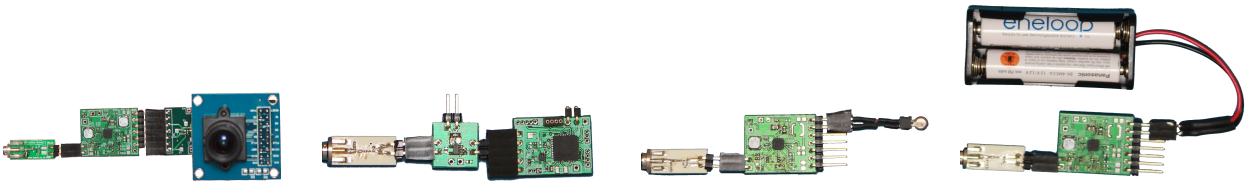

(a) Battery Free Camera (b) Temperature Sensor (c) Li-Ion Battery Charger (d) NiMH Battery Charger

Figure 4.2: **Prototype hardware demonstrating PoWiFi's potential.** The prototypes harvest energy from Wi-Fi signals through a standard 2 dBi Wi-Fi antenna (not shown). The low gain antenna ensures that the device is agnostic to the antenna orientation and placement. The prototypes use the harvested energy to (a) capture pictures, (b) measure temperature, and (c)/(d) recharge batteries.

cation via Wi-Fi, Bluetooth or ZigBee. We can, in principle, use the same antenna for both communication and Wi-Fi power harvesting with negligible increase in the size of the device.

The key challenge for power delivery over Wi-Fi is the fundamental mismatch between the requirements for power delivery and the Wi-Fi protocol. To illustrate, Fig. 4.1 plots the voltage at a tuned harvester in the presence of Wi-Fi transmissions. While the harvester can gather energy during Wi-Fi transmissions, the energy leaks during silent periods. In this case, Wi-Fi transmissions cannot satisfy the platform's minimum voltage requirement. Unfortunately for power delivery, silent periods are inherent to a distributed medium access protocol such as Wi-Fi, in which multiple devices share the same wireless medium. Continuous transmission from the router, while optimal for power delivery, would significantly deteriorate the performance of Wi-Fi clients and other nearby Wi-Fi networks.

We introduce *PoWiFi*, the first power over Wi-Fi system that delivers power to energy-harvesting

sensors and devices while preserving network performance. We achieve this by co-designing harvesting hardware circuits and Wi-Fi router transmissions. At a high level, a router running PoWiFi imitates a continuous transmission while minimizing the impact on its Wi-Fi clients and other Wi-Fi networks. The key intuition is that it is unlikely that all Wi-Fi channels are simultaneously occupied at the same instant. Thus, PoWiFi opportunistically injects superfluous broadcast traffic (which we call *power packets*) on non-overlapping Wi-Fi channels to maximize the *cumulative* occupancy across the channels. To harvest this energy, we introduce the first multi-channel harvester that efficiently harvests power across multiple Wi-Fi channels and generates the 1.8–2.4 V necessary to run microcontrollers and sensor systems (see §4.2.2).

To be practical, PoWiFi must not significantly degrade network performance. So our second component is a transmission mechanism that minimizes the impact on Wi-Fi performance while effectively providing continuous power delivery to harvesters. Specifically, to minimize the impact on associated Wi-Fi clients, PoWiFi injects power packets on a channel only when the number of data packets queued at the Wi-Fi interface is below a threshold. Further, the router transmits power packets at the highest Wi-Fi bit rate. Since higher-rate transmissions occupy the channel for a smaller duration, PoWiFi achieves per-channel occupancies that are fair to other Wi-Fi networks.

To further minimize its impact on neighboring Wi-Fi networks, PoWiFi uses two key techniques.

- *Rectifier-aware transmissions.* The intuition is that when there are packets on the air, a harvester's rectifier charges exponentially, but it also discharges exponentially during silent periods. To balance power delivery and channel occupancy, PoWiFi must minimize the energy loss due to leakage. We achieve this by designing an occupancy modulation scheme that jointly optimizes the rectifier's voltage behavior and the Wi-Fi network's throughput to ensure that harvesting sensors can meet their duty-cycling requirements (see §4.2.1).

- *Scalable concurrent transmissions.* The goal is to maintain good network performance when there are multiple PoWiFi routers in an area. Because PoWiFi's power packets do not contain useful data, our insight is that transmissions from multiple PoWiFi routers can safely collide.

Further, by making each PoWiFi router transmit power packets with random data payload, we ensure that concurrent packet transmissions do not destructively interfere to reduce available power at sensors. Fig. 4.5 shows our transmission structure that enables multiple PoWiFi routers to co-exist.

We build PoWiFi prototypes using Atheros chipsets and build our multi-channel harvester with off-the-shelf components. Our results show the following:

- Power packets from the PoWiFi router do not noticeably affect TCP or UDP throughput or webpage load times [7] at an associated client. Meanwhile, PoWiFi achieves an average cumulative channel occupancy of 95.4% across the three non-overlapping 2.4 GHz Wi-Fi channels.

- PoWiFi's unobtrusive transmission strategy allows neighboring Wi-Fi networks to achieve better-than-equal-share fairness, because a PoWiFi router transmits power packets at the highest bit rate to minimize its channel occupancy.

- Our rectifier-aware transmission mechanism further reduces the effect on the neighboring network — it reduces the required average per-channel occupancy from 40% to 4.4%, while delivering power to a sensor 16 feet away that reads temperature values once every minute.

- We perform a proof-of-concept evaluation of our concurrent transmission mechanism with 1, 3 and 6 PoWiFi routers. While the variance of neighboring Wi-Fi networks' throughput increases slightly, their mean throughput does not differ statistically. This shows the feasibility of scaling our design with multiple PoWiFi routers.

To demonstrate the potential of our design, we use our harvester to build two battery-free, Wi-Fi–powered sensing systems shown in Fig. 4.2: a temperature sensor and a camera. The devices use Wi-Fi power to run their sensors and a programmable microcontroller that collects the data and sends it over a UART interface. The camera and temperature-sensor prototypes can operate battery-free at distances of up to 17 and 20 feet, respectively, from a PoWiFi router. As expected,

the duty cycle at which these sensors can operate decreases with distance. Further, the sensors can operate in through-the-wall scenarios when separated from the router by various wall materials.

We also integrate our harvester with 2.4 V nickel–metal hydride (NiMH) and 3.0 V lithium-ion (Li-Ion) coin-cell batteries. We then build battery-recharging versions of the above sensors wherein PoWiFi trickle charges the batteries using Wi-Fi. The battery-recharging sensors can run energy-neutral operations at distances of up to 28 feet.

Finally, we deploy PoWiFi routers in six homes in a metropolitan area. Each home's occupants used the PoWiFi router for their Internet access for 24 hours. Even under real-world network conditions, PoWiFi efficiently delivers power while having a minimal impact on user experience.

Contributions. We make the following contributions:

- Introduce PoWiFi, a novel system for power delivery using existing Wi-Fi chipsets. We do so without compromising the Wi-Fi network's communication performance.

- Co-design router transmissions and harvesting hardware circuits to balance power delivery and network performance. Our novel multi-channel harvester can efficiently harvest power from multiple 2.4 GHz Wi-Fi channels.

- Prototype the first battery-free temperature and camera sensors that are powered using Wi-Fi chipsets. We also demonstrate the feasibility of recharging NiMH and Li-Ion coin-cell batteries using Wi-Fi signals.

- Deploy our system in six homes in a metropolitan area and demonstrate its real-world practicality.

Limitations. Given today's FCC limits in the ISM band (1 W), power over Wi-Fi is limited to low-power sensors and devices and can not recharge smartphones (5 W). Further, the range of our system is determined by the sensitivity of our harvester hardware, which is built with off-the-shelf components. We believe that an ASIC design would be able to improve the sensitivity and double

PoWiFi's power-delivery range. Finally, while our current design does not account for MIMO, in principle, we can use multiple antennas to focus more power toward a sensor and increase the range, but such optimizations are beyond the scope of this work.

4.1 Understanding Wi-Fi Power Delivery

To understand the ability of a Wi-Fi router to deliver power, we run experiments with our organization's router and a temperature sensor. The router is an Asus RT-AC68U access point operating at 2.437 GHz with a transmit power of 23 dBm on each of its three 4.04 dBi gain antennas. The temperature sensor is battery free and uses our RF harvester to draw power from Wi-Fi signals. A typical RF harvester has to provide a minimum voltage at the sensor or microcontroller to run meaningful operations. This is typically done using a rectifier that converts the carrier signal to DC and a DC–DC converter that increases the voltage level of the DC signal to match the requirements of the sensor or microcontroller. The key limitation in harvesting power is that every DC–DC converter has a minimum input voltage threshold below which it cannot operate. We use a DC–DC converter with the lowest threshold of 300 mV [50].

We place the sensor ten feet from the router for 24 hours and measure the voltage at the rectifier's output throughout our experiments. We also capture the packet transmissions from the router using a high frequency oscilloscope connected through a splitter. Over the tested period, the sensor could not reach the 300 mV threshold. Fig. 4.1 plots both the packet transmissions and the rectifier voltage. It shows that while the sensor can harvest energy during the Wi-Fi packet transmission, there is no input power during the silent slots. The hardware power leakages during these durations ensure that it does not cross the 300 mV threshold.

Fig. 4.1 is a snapshot of router transmissions during peak network utilization. More generally, the router's channel occupancy was in the 10–40% range, mostly at the lower end of this range. Note that clients such as smartphones typically transmit at lower power than the router. Our measurements show that, to save energy, smartphones such as Nexus S, Nexus 4 and iPhone 5 reduce their per-packet transmission power to between 0–2 dBm. Thus, efficient power delivery specifically requires high channel occupancies at the router.

4.2 PoWiFi

PoWiFi is a novel system, which provides power over Wi-Fi using existing Wi-Fi chipsets. It combines two elements: (1) a PoWiFi router injects small amounts of unobtrusive power traffic on multiple Wi-Fi channels to increase channel occupancy with minimal impact on network performance and (2) energy-harvesting hardware that can efficiently harvest from multiple Wi-Fi channels simultaneously.

4.2.1 PoWiFi Router Design

Our goal is to maximize power-delivery efficiency that requires maximizing channel occupancy. A naïve solution is to continuously transmit packets at the lowest Wi-Fi bit rate, i.e., 1 Mbps. Since such transmissions occupy the Wi-Fi channel for the longest duration, they maximize the channel occupancy. However, such an approach would significantly deteriorate the performance of the Wi-Fi network, as our evaluation confirms (see §4.3.1).

Our key insight is that, at any moment, it is unlikely that all Wi-Fi channels will be occupied. Thus, PoWiFi opportunistically injects power packets across multiple Wi-Fi channels with a goal of maximizing *cumulative* occupancy. Our idea is to inject small amounts of traffic on multiple Wi-Fi channels at the router to ensure that cumulative occupancy is high. The rest of this section first describes how we can inject additional packets while minimizing the effect on Wi-Fi clients and then describes design choices that ensure fairness with other Wi-Fi networks.

Our harvesting hardware does not decode Wi-Fi signals and as a result, from its perspective, all router transmissions look identical. Thus, it can harvest similar amounts of power from artificial packets as well as traffic to Wi-Fi clients and beacon transmissions. We leverage this property to design a system that balances client traffic and additional power traffic.

At a high level, our design injects UDP broadcast packets[1] at the highest Wi-Fi bit rate to transmit power on each of the Wi-Fi channels. However, PoWiFi drops these broadcast packets when the number of packets in the wireless interface's transmit queue is above a threshold. This

[1]UDP broadcast packets do not require acknowledgments from clients, either at the PHY or the higher layers.

ensures that when the router queue has client traffic, we do not add additional packets and hence can minimize the effect on the client delay and throughput.

Specifically, we implement a user-space program that injects 1500-byte UDP broadcast datagrams with a constant inter-packet delay. We use a *selective transmission* mechanism that hoists information from the MAC layer to the IP layer. Our mechanism has three main components:

- `Power_Socket`: A standard UDP broadcast socket with the addition of a custom IP option, `IP_Power`, to distinguish its outgoing IP datagrams from other traffic.

- `Power_MACshim`: A shim interface between the IP stack and the `mac80211` subsystem that enables the IP stack to query the Wi-Fi subsystem for the queue status of individual channels. On socket creation, the user-space program sets an additional IP option with an integer that uniquely identifies the corresponding wireless interface at the router.

- `IP_Power`: A mechanism in the IP stack that checks for our power packets on the outgoing IP datagrams and uses our shim interface to decide when to drop the packets.

The decision to drop packets is performed on a per-packet basis in the packet transmission logic of the IP stack, i.e., `ip_local_out_sk()`. We check whether the pending queue depth is above a threshold value. This check is channel specific; it is applied after the kernel has determined a route and therefore an interface for the packet. If the queue depth is indeed at or above a threshold value, then there are already enough power and Wi-Fi client packets in the queue to maximize channel occupancy. In this case the router drops the packet before transmitting it and returns the corresponding error code to user space. On the other hand, if the queue depth is below the threshold value, then `IP_Power` queues the packet for transmission at the MAC layer. We note that in our evaluation, the router is configured to provide Internet connectivity on only one 2.4 GHz Wi-Fi channel. Thus, on other Wi-Fi channels, there are no client packets in the queue and hence, we do not drop any UDP broadcast packets.

Finally, we summarize some of our design decisions.

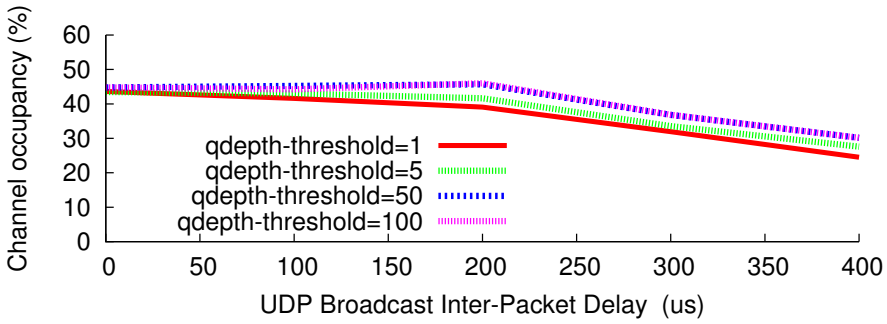

Figure 4.3: **Effect of inter-packet delay on occupancy.** Results in the absence of client traffic for different queue depth thresholds.

i) Queue threshold value. After extensive testing, we set a fixed queue depth threshold of five frames. Specifically, our tests showed that for thresholds less than five, occupancy decreases since the queue is repeatedly drained and the user-space program that sends UDP broadcast packets was unable to keep the queue full. Larger threshold values, on the other hand, required more frequent transmissions, resulting in increased slowdown for client traffic.

ii) UDP broadcast data rate. If the UDP broadcast rate is high, then frames pile up in the queues and affect the kernel's responsiveness. On the other hand, a low rate significantly reduces the occupancy across Wi-Fi channels. Fig. 4.3 shows the occupancy on a single Wi-Fi channel for different inter-packet delays as well as queue thresholds, in the absence of client traffic. The figure shows that varying the queue-depth threshold does not significantly affect occupancy in the absence of client traffic as long as inter-packet timing is less than the length of the corresponding frames on the air. Our implementation uses 1500 byte packets transmitted at the highest 802.11g bit rate of 54 Mbps. These packets occupy around 160 us on the wireless channel, and so we pick an inter-packet delay of 100 us to balance occupancy and kernel responsiveness.

iii) Fairness with other Wi-Fi networks. PoWiFi is compliant with the 802.11 MAC protocol to ensure that active Wi-Fi devices get equal access to the wireless channel. In practice, PoWiFi provides better than equal-share fairness to transmissions from other Wi-Fi devices. Specifically, the UDP broadcast packets are transmitted at the highest Wi-Fi bit rate. These transmissions occupy the channel for a shorter duration than transmissions at lower Wi-Fi bit rates. Thus, for the average

transmitter bit rate in the network, we achieve better than equal-share fairness. This is validated in our experiments in §4.3.1.

The rest of this section describes two techniques that further reduce PoWiFi's effect on neighboring Wi-Fi networks.

Rectifier-aware PoWiFi transmissions

When PoWiFi knows a harvester's electrical characteristics, it can tune its transmission strategy to precisely fit the device's power requirements. For example, suppose we need to read a temperature sensor once per minute. PoWiFi can modulate its occupancy to deliver energy to the harvester so that the sensor reaches its required voltage of 2.4 V just in time, minimizing the total channel occupancy subject to this goal and thereby minimizing its effect on other networks.

Empirically modeling rectifier voltage. A rectifier converts incoming Wi-Fi transmissions into DC voltage to charge a storage capacitor. Once the voltage on the capacitor reaches the required threshold ($V_{th} = 2.4, \mathrm{V}$ for the temperature sensor), a reading occurs. Suppose the average power at the harvester after multi-path reflections and attenuation is P_{in} and the channel occupancy of the PoWiFi router packets is C. To a first approximation, the harvester's behavior can be modeled as a DC voltage source charging a capacitor through a resistor. The difference, however, is that the approximated resistance value depends on the impedance of the harvester's diodes, which is a function of P_{in} and C. We can write the voltage as a function of time as

$$V(t) = V_0 * e^{-t/\tau(P_{in},C)} + V_{max}(P_{in}, C) * \left(1 - e^{-t/\tau(P_{in},C)}\right),$$

where V_0 is the initial voltage, τ is the time constant, and V_{max} is the maximum achievable voltage. Note that both τ and V_{max} are functions of P_{in} and the channel occupancy.

Given the non-linearitys of diodes, it is difficult to obtain closed-form solutions for $\tau(P_{in}, C)$ and $V_{max}(P_{in}, C)$. We instead connected the harvester through a cabled setup to a Wi-Fi source with variable input power and channel occupancy and measured the output voltage. We fitted the resulting data with the proposed exponential model to estimate how τ and V_{max} vary with input power and channel occupancy. The properties of our model fitting are: 1) V_{max} is non-linearly

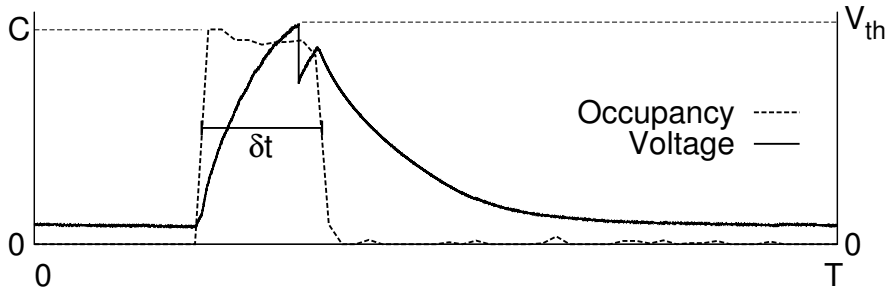

Figure 4.4: **Rectifier-aware power Wi-Fi transmissions and corresponding rectifier voltages.** The plot shows the optimized rectifier aware power Wi-Fi transmission and the corresponding voltage at a temperature sensor's storage capacitor (dotted line).

proportional to the input power and channel occupancy; 2) the time constant τ is exponentially proportional to the input power and/or the channel occupancy; and 3) it takes exponentially more time for the same increment in the voltage at a higher voltage value than at a lower one.

We next describe how PoWiFi can modulate its channel occupancy using this empirical model, while minimizing its effect on neighboring Wi-Fi networks.

Joint optimization for efficient power delivery. To reduce the impact of power packets on neighboring Wi-Fi networks, PoWiFi must minimize the total number of power packets required to collect a sensor reading. Our key intuition is that when there are packets on the air, the capacitor charges exponentially. However, when there are no packets, the voltage on the capacitor discharges exponentially. To maximize the effectiveness of power delivery, PoWiFi must minimize capacitor leakage. We achieve this by using the channel-occupancy modulation scheme described shown in Fig. 4.4. In every sensor update time window (T), the router transmits no power packets for a period ($T - \delta t$), then transmits power packets for a period of δt, targeting a channel occupancy of $0 < C \leq 1$. When the channel occupancy is zero, the voltage on the capacitor is very low and there is no leakage. However, when a sensor update is required, a high channel occupancy continuously charges the capacitor (minimizing leakage) to maximizes the effectiveness of power delivery. Our goal is to find δt and C to minimize the mean of the power packet occupancy given by $C * \frac{\delta t}{T}$.

We find these values by substituting different C and δt in our empirical model and finding the minimum value. We reduce the search space by noting that for a given P_{in}, there is a minimum

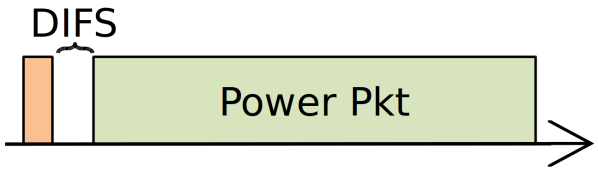

Figure 4.5: **Energy pattern for concurrent power packet transmissions.** PoWiFi transmissions consists of the short packet with a 1-byte payload transmitted at 54 Mbps followed by DIFS period and then followed by the power packet transmission.

value of C below which the threshold voltage is not achievable. Further, given a channel occupancy, we know the time constant that limits the value of δt to a maximum value of $\tau(P_{in}, O)$. Finally, we limit the granularity by which channel occupancy can be modulated to 10%. Using these values, we were able to reduce the search space to 75 points.

We note two main points. First, the above description assumes that the router can estimate the available power, P_{in}, at the sensor. To bootstrap this value, PoWiFi initially transmits power packets at a high occupancy of around 90% and notes the times when the sensor outputs a reading. PoWiFi uses our empirical model to estimate P_{in} for the next cycle. At the end of every cycle it re-estimates P_{in} to account for wireless channel changes. Second, in the presence of multiple sensors, we can optimize the parameters to satisfy the minimum duty cycle requirement across all the sensors, but we omit this simple extension for brevity.

Scaling with concurrent PoWiFi transmissions

A practical issue with each PoWiFi router independently introducing power packets is that such a system would not preserve network performance in the presence of many PoWiFi routers. Useful Wi-Fi capacity would degrade at least linearly with the number of PoWiFi routers.

To address this scaling problem, we enable concurrent transmissions from PoWiFi routers that are in decoding range of one another. Our key insight is that since power packets do not contain useful data, transmissions from multiple PoWiFi routers can safely collide. Further, if each PoWiFi router transmits a random power packet, we can ensure that concurrent packet transmissions do not destructively interfere to reduce the power available to harvesters.

Specifically, in our system, we have a leader PoWiFi router that transmits the energy pattern shown in Fig. 4.5. The pattern consists of a short packet with a 1-byte payload transmitted at 54 Mbps, followed by a DIFS period and then a power packet. Other PoWiFi routers decode this short packet and join the packet transmission of the leader router within the DIFS period. This strategy ensures that all nearby PoWiFi routers transmit power packets concurrently and hence do not reduce the Wi-Fi network's capacity.

Similar to [97], we enable concurrent transmissions from the follower routers in software by setting CW_{min} and CW_{max} to 1, preventing carrier sense back off by setting the noise floor registers to "high" and placing their power packets in the high-priority queue. PoWiFi could not turn around and begin transmission with the current software access within a DIFS duration. However, we believe that with better access to the router's hardware queues, PoWiFi can turn around within a DIFS period. Finally, one can design distributed algorithms to find the leader router whose transmissions can be decoded by all other PoWiFi routers, but we consider this to be outside the scope of this paper.

4.2.2 Multi-Channel Harvester Design

The first goal of our harvester design is to efficiently harvest across multiple 2.4 GHz Wi-Fi channels. A related goal is to achieve good sensitivities across these channels. Sensitivity is the lowest power at which the harvester can boot up and power the sensors and the microcontroller. In theory, one can wait for a long time and harvest enough power to boot up the sensors. In practice, however, because of power leakage, a harvester cannot operate below a minimum power threshold. This is important because the power available at the sensor decreases with the distance from the Wi-Fi router; thus, the harvester's sensitivity determines its operational range.

itChallenge: The key challenge is the impedance mismatch between the Wi-Fi antenna and the harvester. To understand this, consider a wave entering a boundary between two different mediums. If the impedance of the two mediums differs, a fraction of the incident energy is reflected. Similarly, when the antenna and the harvester have different impedance values, a fraction of the

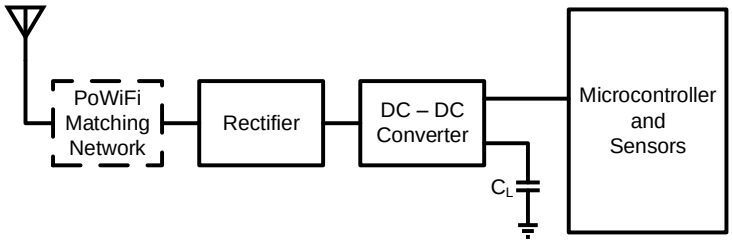

Figure 4.6: **RF Harvester Architecture.** An antenna receives RF signals, which a rectifier converts into DC power and feeds into a DC–DC converter that increases the voltage to match the sensor and microcontroller's requirements.

RF signal is reflected back, reducing the available RF power.

Fig. 4.6 shows the architecture of a typical RF harvester. A receiving antenna is followed by a rectifier that converts the 2.4 GHz signal into DC power. This power is fed into a DC–DC converter that increases the voltage of the DC signal to match the voltage requirements of the sensor and microcontroller. The problem is that the rectifier hardware is extremely non-linear with input power, operational frequency and the parameters of the DC–DC converter, making it challenging to achieve good harvester sensitivity and efficiency across the 72 MHz band that spans the three Wi-Fi channels.

Our Approach: As shown in Fig. 4.6, we design a matching network to transform the rectifier's impedance to match that of the antenna. This is, however, not straightforward because the rectifier's impedance varies significantly with frequency and is dependent on the DC–DC converter. Our approach is to co-design all the components in the harvester—the matching network, rectifier, and DC–DC converter—to achieve good impedance matching across the 72 MHz Wi-Fi band. Our intuition is that the input of the DC–DC converter affects the input impedance of the rectifier. Thus, if we can co-design the rectifier with the DC–DC converter, we can relax the constraints on the matching network.

Design Details: The rest of the section describes each of the above components—rectifier, DC–DC converter, and matching network—in detail.

1) Rectifier Design. The key design consideration for rectifiers is that DC–DC converters cannot operate below a minimum input voltage. Thus, the rectifier must be designed to maximize its

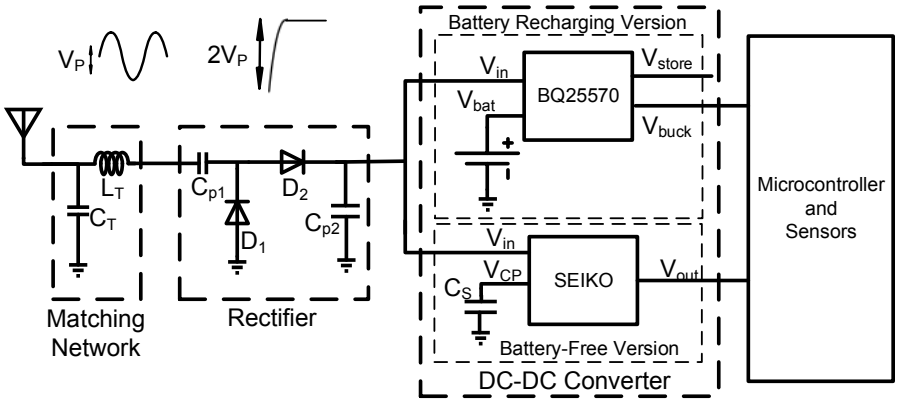

Figure 4.7: **PoWiFi harvester schematic.** PoWiFi co-designs the matching network, rectifier, and DC–DC converter to achieve good impedance matching across Wi-Fi bands. The figure shows the optimized DC–DC converters for both battery-free and battery-recharging versions of our harvester.

output voltage. Fig. 4.7 shows the various components used in our rectifier design. At a high level, our rectifier tracks twice the envelope of the incoming signal and converts it into power. Specifically, it adds the positive and negative cycles of the incoming sinusoidal carrier signal to double the amplitude. To do this, it uses a specific configuration of diodes and capacitors as shown in Fig. 4.7. However, in practice, diodes and capacitors have losses that limit the output voltage of the rectifier. We use SMS7630-061 diodes by Skyworks [52] in ultra-miniature 0201 SMT packages since they have low losses, i.e., loss threshold voltage, low junction capacitance and minimal package parasitics. We also use high–quality-factor, low-loss UHF-rated 10 pF capacitors that minimize losses and maximize the rectifier's efficiency and sensitivity.

2) DC–DC converter design. In our design, a DC–DC converter serves two purposes: i) boost the voltage output of the rectifier to the levels required by the microcontroller and sensors, and ii) make the input impedance of the rectifier less variable across the three Wi-Fi channels. The key challenge is the cold-start problem: in a battery-free design, all the hardware components must boot up from 0 V. Practical DC–DC converters, however, have a nonzero minimum voltage threshold. We use the SZ882 DC–DC converter from Seiko [50], which is the best in its class: it can start from input voltages as low as 300 mV, which our rectifier can provide, and boost the output on a storage capacitor to 2.4V. Once the 2.4 V threshold is reached, the Seiko charge pump connects the

storage capacitor to the output, powering the microcontroller and sensors.

A DC–DC converter can be further optimized while recharging a battery. Specifically, the battery can provide a minimum voltage level and hence the hardware components need not boot up from 0 V. We use the TI bq25570 energy-harvesting chip [13] that contains a boost converter, a battery charger, voltage monitoring solutions and a buck converter. We connect the rechargeable battery to the battery charging node, V_{bat}, of the bq25570. We use the boost as our DC–DC converter to achieve the voltage required to charge the battery. Finally, we leverage the maximum power point tracking (MPPT) mode of the TI chip to tune the input impedance of the DC–DC converter so as to minimize the variation of the rectifier's impedance across Wi-Fi channels. Specifically, we set the buck converter's MPPT reference voltage to 200 mV.

3) Matching Network Design: With our rectifier and DC–DC converter designs, we have relaxed the constraints on the impedance-matching network. The resulting circuit can match impedances between the rectifier and a 50 Ω antenna across Wi-Fi channels, using a single-stage LC matching network. In LC matching networks, inductors are the primary source of losses. To mitigate this, we use high-frequency inductors in 0402 footprint which have minimal parasitics and a quality factor of 100 at 2.45 GHz [1]. The resulting matching network consumes less board area than traditional transmission lines and distributed-element–based matching networks and can be modified to meet different system parameters without any loss. We use 6.8 nH and 1.5 pF as the LC matching network for our battery-free harvester, and 6.8 nH and 1.3 pF for our battery-recharging harvester.

4.3 Evaluation

We build the rectifiers for our harvester prototypes using 2-layer 20 mils Rodgers 4350 substrate printed circuit boards (PCBs). We use the Rodgers substrate because unlike FR4 [49], it has low losses at 2.4 GHz and does not degrade the sensitivity and efficiency of our harvester. The DC–DC converter and sensor applications however were built using a 4-layer FR4 substrate PCBs and connected to the harvester using 10 mil headers. The PCBs were designed using Altium design software and were manufactured by Sunstone Circuits. A total of 40 PCBs were ordered at a

total cost of $2500. The off-the-shelf circuit components were hand-soldered on the PCBs and individually tested, requiring a total of 200 person-hours.

We implement a PoWiFi router using three Atheros AR9580 chipsets that independently run the algorithm in §4.2.1 on channels 1, 6, and 11 respectively. The chipsets are connected to 6 dBi Wi-Fi antennas via amplifiers; the antennas are separated by 6.5 cm, which is approximately half a wavelength at 2.4 GHz. Our prototype router provides Internet access to its associated clients on channel 1 via NAT and transmits at 30 dBm, which is within the FCC limit for communication in the ISM band. Since the Atheros chipsets operate independently, the cumulative occupancy across the three Wi-Fi channels can be greater than 100% in under-utilized networks. One can implement simple algorithms that would scale back the transmission rate for power packets to ensure that the cumulative occupancy remains less that 100%. We do not currently implement this feature. Note however that all our sensor and harvester benchmark evaluations were performed in a busy office network where the average cumulative channel occupancy was around 90%.

Measuring the router's channel occupancy. One of our key metrics is the router's channel occupancy that includes both the power packets and packets to its clients. To measure this, we use `aircrack-ng`'s `airmon-ng` tool to add a monitor interface to each of the router's active wireless interfaces. To measure the router's channel occupancy on a specific interface, we start `tcpdump` on the monitor interface to record the radio–tap headers for all frames and their retransmissions. At the end of the duration, we stop `tcpdump` and use `tshark` to extract frames sent by the router, recording the corresponding bitrate and frame size (in bytes). We then compute the average channel occupancy as $\sum_{i \in frames} \frac{size_i}{rate_i \times total\_duration}$.

4.3.1 Effect on Wi-Fi clients

Our system is designed to provide high cumulative channel occupancies for power delivery while minimizing the effect on Wi-Fi traffic. To evaluate this, we deploy a PoWiFi router and evaluate its effect on Wi-Fi traffic. We use a Dell Inspiron 1525 laptop with an Atheros chipset as a client associated with our router on channel 1.

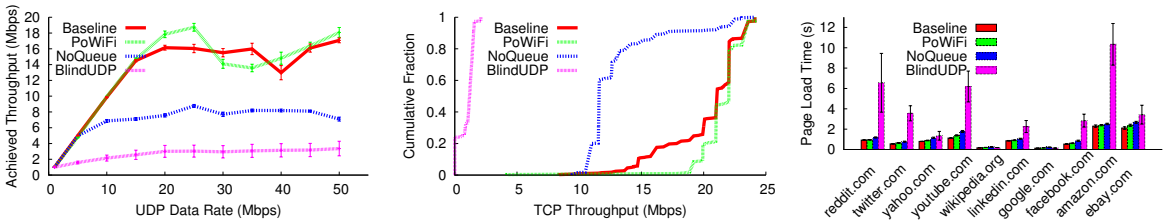

Figure 4.8: **Effect on Wi-Fi traffic.** The figures show the effect of various schemes on TCP and UDP throughput as well as the page load times of the top ten websites in the United States [7]. The plots show that PoWiFi minimizes its effect on the Wi-Fi traffic.

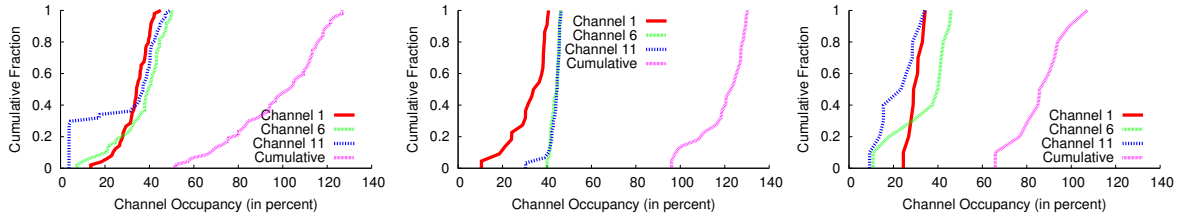

(a) UDP experiments *(b) TCP experiments* *(c) PLT experiments*

Figure 4.9: **PoWiFi channel occupancies.** The plots show the occupancies with PoWiFi for the above UDP, TCP, and PLT experiments.

We compare four different schemes:

- *Baseline.* PoWiFi is disabled on the router, i.e., the router introduces no extra traffic on any of its interfaces.

- *BlindUDP.* The router transmits UDP broadcast traffic at 1 Mbps so as to maximize its channel occupancy.

- *PoWiFi.* The router sends UDP broadcast traffic at 54 Mbps and uses the queue threshold check in §4.2.1.

- *NoQueue.* The router sends UDP broadcast traffic at 54 Mbps but disables the queue threshold check.

We evaluate PoWiFi with various Wi-Fi traffic patterns and metrics: the throughput of UDP and TCP download traffic, the page load time (PLT) of the ten most popular websites in the United States [7], and traffic on other Wi-Fi networks in the vicinity of our benchmarking network.

(a) Effect on UDP traffic. UDP is a common transport protocol used in media applications such as video streaming. We run iperf with UDP traffic to a client seven feet from the router. The client sets its Wi-Fi bitrate to 54 Mbps and runs five sequential copies of iperf, three seconds apart. We repeat the experiments with target UDP data rates between 1 and 50 Mbps, and measure the achieved throughput computed over 500 ms intervals. All the experiments are run during a busy weekday at UW CSE, with multiple other clients and 43 other Wi-Fi networks operating at 2.4 GHz.

Fig. 4.8(a) plots the average UDP throughput as a function of the eleven tested UDP data rates. The figure shows that *BlindUDP* significantly reduces throughput. With *NoQueue*, the router's kernel does not prioritize the client's iperf traffic over the power traffic. This results in roughly a halving of the iperf traffic's data rate as the wireless interface is equally shared between the two flows. With PoWiFi, however, the client's iperf traffic achieves roughly the same rate as the baseline. This result demonstrates that PoWiFi effectively prioritizes client traffic above its power traffic.

For the PoWiFi experiments above, Fig. 4.9(a) plots the CDFs of individual channel occupancies on the three Wi-Fi channels. The figure shows that the individual channel occupancies are around 5–50% across the channels. The mean cumulative occupancy, on the other hand is 97.6%, demonstrating that PoWiFi can efficiently deliver power even in the presence of UDP download traffic.

(b) Effect on TCP traffic. Next we run experiments with TCP traffic using iperf at the client. The router is configured to run the default Wi-Fi rate adaptation algorithm. We run experiments over a duration of three hours with a total of 30 runs. In each run, we run five sequential copies of iperf, three seconds apart, and compute the achievable throughput over 500 ms intervals, with all the schemes described above.

Fig. 4.8(b) plots CDFs of the measured throughput values across all the experiments. The plot

shows that *BlindUDP* significantly degrades TCP throughput. As before, since *NoQueue* does not prioritize the client traffic over the power packets, it roughly halves the achievable throughput. PoWiFi sometimes achieves higher throughput than the baseline. This is because of channel changes that occur during the three-hour experiment duration. The general trend however points to the conclusion that PoWiFi does not have a noticeable effect on TCP throughput at the client.

Fig. 4.9(b) plots the CDFs of the channel occupancies for PoWiFi during the above experiments. The figure shows that PoWiFi has a mean cumulative occupancy of 100.9% and hence can efficiently deliver power.

(c) Effect on PLT. We develop a test harness that uses the PhantomJS headless browser [45] to download the front pages of the ten most popular websites in the US [7] 100 times each. We clear the cache and pause for one second in between page loads. The traffic is recorded with `tcpdump` and analyzed offline to determine page load time and channel occupancy. The router uses the default rate adaptation to modify its Wi-Fi bit rate. The experiments were performed during a busy weekday at UW CSE over a two-hour duration.

Fig. 4.8(c) shows that *BlindUDP* significantly deteriorates the PLT. This is expected because the 1 Mbps power traffic occupies a much larger fraction of the medium and hence increases packet delays to Wi-Fi clients. *NoQueue* improves PLT over *BlindUDP*, with an average delay of 294 ms over the baseline. PoWiFi further minimizes the delay to 101 ms, averaged across websites. This residual delay is due to the computational overhead of PoWiFi from the per-packet checks performed by the kernel. This slows down all the processes in the OS and hence results in additional delays. However, increasing processing power and moving these checks to hardware can help further reduce these delays. In our home deployments (§4.5), the users did not perceive any noticeable effects on their web performance.

For completeness, we plot the CDFs of channel occupancies for PoWiFi in Fig. 4.9(c). The plot shows the same trend as before, with a mean cumulative occupancy of 87.6%.

4.3.2 Effect on neighboring Wi-Fi networks

(a) High cumulative channel occupancy transmissions. PoWiFi leverages the inherent fairness of the Wi-Fi MAC to ensure that it is fair to other Wi-Fi networks. As a worst-case evaluation, we consider a situation where PoWiFi always tries to achieve high cumulative channel occupancies at all times. To do this, we place our PoWiFi router in the vicinity of a neighboring Wi-Fi router–client pair operating on channel 1. We configure the PoWiFi router to transmit power aware packets at the highest achievable channel occupancies using our algorithm on all three channels. We run iperf with UDP traffic on the neighboring router–client pair at the highest data rate and measure the achievable throughput as before. We repeat the experiments for different Wi-Fi bit rates at the neighboring Wi-Fi router–client pair. We compare three schemes: *BlindUDP* where our router transmits UDP packets at 1 Mbps, *EqualShare* where we set our router to transmit the UDP packets at the same Wi-Fi bit rate as the neighboring router–client pair, and finally PoWiFi. *EqualShare* provides a baseline when every router in the network gets an equal share of the wireless medium.

Figure 4.10(a) shows the throughput for the three schemes, averaged across five runs. As expected, *BlindUDP* significantly deteriorates the neighboring router–client performance. Further, this deterioration is more pronounced at the higher Wi-Fi bit rates. With PoWiFi, however, the throughput achieved at the neighboring router–client pair is higher than *EqualShare*. This is because PoWiFi transmits power packets at 54 Mbps; transmissions at such high Wi-Fi bit rates occupy the channel for a smaller duration than, say, a neighboring router transmitting at 16 Mbps. This property means that PoWiFi provides better than equal-share fairness to other Wi-Fi networks. We note that while our experiments are with 802.11g, PoWiFi's power packets use the highest bit rate available for Wi-Fi. Thus, the above fairness property would hold true even with advanced protocols such as 802.11n/ac.

(b) Rectifier-aware power transmissions. Next we evaluate the potential of our rectifier-aware technique to significantly reduce the average channel occupancy of the power transmissions, while efficiently delivering power to the sensors. To do this, we place our battery-free temperature sensor close to its maximum operational range at 16 feet from a PoWiFi router; the sensor is set to transmit

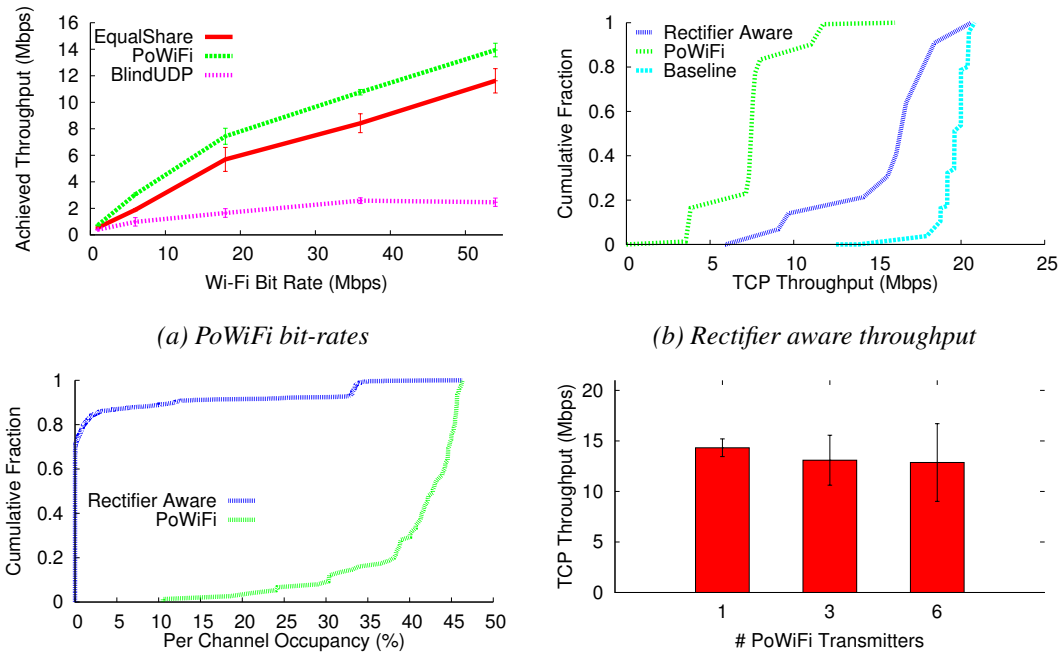

(a) PoWiFi bit-rates

(b) Rectifier aware throughput

(c) Rectifier aware occupancies

(d) Concurrent transmissions

Figure 4.10: **Effect of PoWiFi, rectifier aware and concurrent power transmissions on neighboring Wi-Fi networks.** (a) show that PoWiFi power transmissions provide better than *EqualShare* throughput performance. Rectifier aware power transmissions further improve the throughput by reducing the per channel occupancy by a factor of 10. Additionally, increasing the number of concurrently transmitting PoWiFi devices does not degrade the performance of neighboring Wi-Fi devices.

a temperature value over a UART interface once every 60 seconds. The router implements the joint-optimization algorithm from §4.2.1.

We ran the experiments for a total of ten minutes and observed that the temperature sensor achieves a mean update rate of 59.93s with a 0.43s variance. More importantly, in contrast to transmitting at high channel occupancies ($> 90\%$) all the time, our algorithm estimated that the router should transmit for a duration of 9s every 60s with a 80% cumulative occupancy and stay quite for the remaining time. Fig. 4.10(b) shows the throughput of a ongoing TCP flow in a neighboring Wi-Fi router-client pair, which shows that the average throughput for rectifier-aware power transmissions significantly improves over high-occupancy PoWiFi and is much closer to the baseline throughput without any power packets. Fig. 4.10(c) shows that rectifier aware transmissions

have an average per-channel occupancy of 3.3%, compared to 40% per-channel occupancy for PoWiFi transmissions — a 10x reduction in average occupancy.

(c) Scalable concurrent power transmissions. Finally, we provide a proof-of-concept evaluation of our concurrent transmission mechanism. Wi-Fi hardware is designed to turn around between decoding a packet and transmitting within a SIFS duration and hence can in principle, easily achieve the timing requirement in Fig. 4.5(d). Since we currently only have software access to the router, we are limited to using high-speed timers and high-priority queue. Our current software system has 36.15 μs mean turn around time with 4.61 μs variance.

Using the above mean turn around time as the silence period, we do a proof-of-concept evaluation. To simplify implementation, we setup a USRP N210 to transmit the pattern in Fig. 4.5 at 30% channel occupancy. The PoWiFi routers join this USRP transmission and concurrently transmit power packets. We evaluate the impact on the TCP throughput of a neighboring Wi-Fi router-client pair as we increase the number of PoWiFi routers. Fig. 4.10(d) shows that as the number of devices increases, the throughput variance increases slightly. This is because as the number of devices increases, the variance in the turn-around time between Wi-Fi power transmissions increases. However, the figure shows that; the mean throughput is statistically unaffected as the number of PoWiFi devices increases from 1 to 6. This demonstrates the feasibility of scaling with multiple PoWiFi routers.

4.3.3 Evaluating the Harvesting Hardware

The harvester's performance is determined by: 1) impedance matching at the antenna interface to maximize the RF energy delivered to the rectifier, and 2) the rectifier's ability to convert RF energy into useful DC power.

(a) Impedance matching versus frequency. If the antenna's impedance differs from the harvester's, a portion of the incident RF signal will be reflected back and cannot be converted into DC power. The amount of reflection is determined by the impedance difference, which our matching network aims to minimize across all three Wi-Fi channels. Impedance matching performance is measured

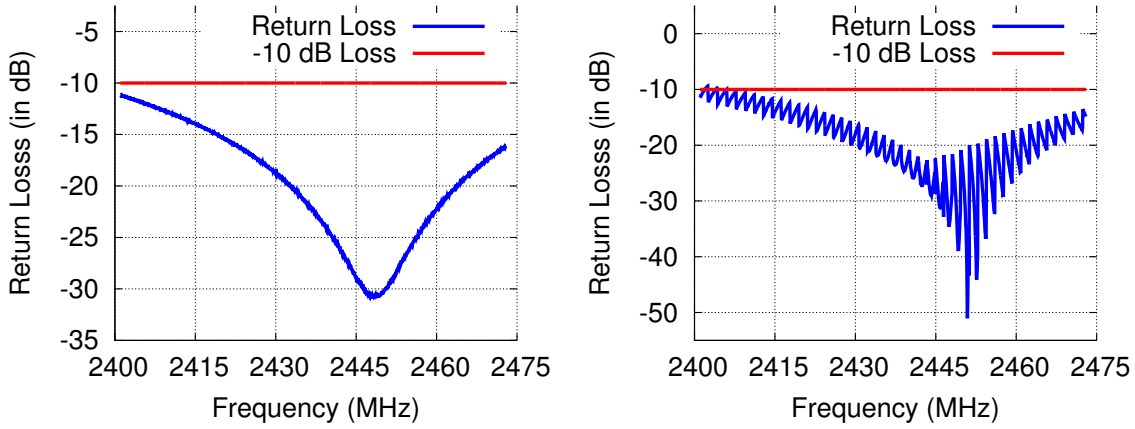

(a) Battery-free harvester *(b) Battery-charging harvester*

Figure 4.11: **Harvester return loss.** This is the ratio of reflected power to the incident power. Across the 2.4 GHz Wi-Fi band, the return loss is less than -10 dB. This translates to less than 0.5 dB of lost power, which is negligible.

using return loss: ratio of reflected power to the incident power.

We compute the return loss by connecting the harvester to a vector network analyzer that transmits RF signals across the entire Wi-Fi band. We analyze the power reflected at each frequency to compute the return loss. Fig. 4.11 plots the return loss of the battery-free and battery-charging versions of our harvester. Across 2.401–2.473 GHz, both of our harvesters achieve a return loss of less than -10 dB, which is acceptable for most RF circuits and systems [171]. This translates to less than 0.5 dB of lost power, which is negligible.

(b) Available power at the rectifier output. The rectifier converts the RF signals at the harvester into DC output voltage. This conversion is typically inefficient due to the inherent nonlinearities and threshold voltage drop of diodes. To measure the available power, we use a cable to connect our hardware to the output of a Wi-Fi transmitter and a continuous wave transmitter. We found that compared to continuous wave, Wi-Fi transmissions have 0.5 dB higher sensitivity which increases the operating range by 6%. Next we vary the output power and the operational frequency of the Wi-Fi transmitter and measure the power available at the rectifier's output.

Fig. 4.12 shows the output power at the rectifier as a function of input RF power. The results are plotted for both our battery-free as well as battery-charging harvesters, across the three Wi-Fi

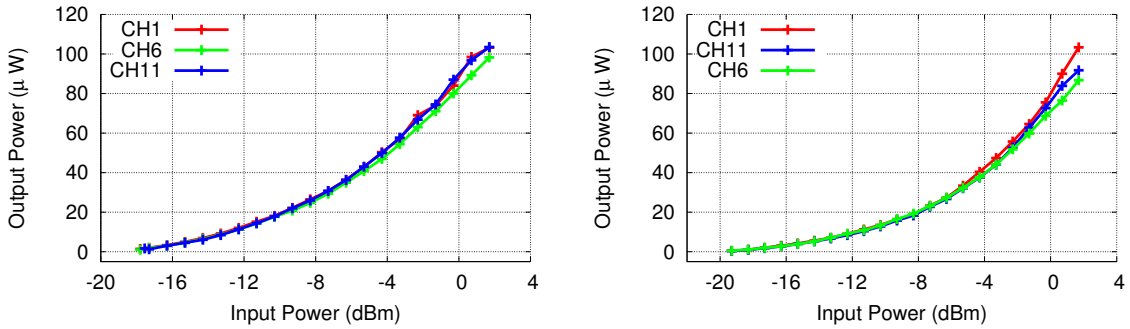

(a) Battery-free harvester *(b) Battery-charging harvester*

Figure 4.12: **Available output power at the harvester.** The battery charging harvester operates at -19.3 dBm compared to -17.8 dBm for the battery free harvester which results in a higher operating range for the battery charging harvester.

channels. The plots show the following:

- The harvester's output power scales with the input power. For instance, at a distance of 2 feet the battery charging system has 100 μW available, compared to 10 μW at 10 feet. This means that as a harvesting sensor moves farther from the router, it operates at a lower duty cycle.

- The battery-charging harvester operates down to -19.3 dBm, compared to -17.8 dBm for the battery-free harvester. This is because the battery-charging harvester does not have the cold start limitation. Specifically, a battery-free harvester has to start all its hardware components from cold start (0 V). In contrast, a battery-charging harvester can use the connected battery to provide a non-zero voltage value, allowing for greater sensitivities.

- Our harvesters perform efficiently across Wi-Fi channels 1, 6 and 11. This is a result of our optimized multi-channel harvester design that ensures efficient power harvesting.

4.4 Sensor Applications

We integrate our harvesters with sensors at two ends of the energy-consumption spectrum: a temperature sensor and a camera. We build both battery-free and battery-recharging versions of each sensor.

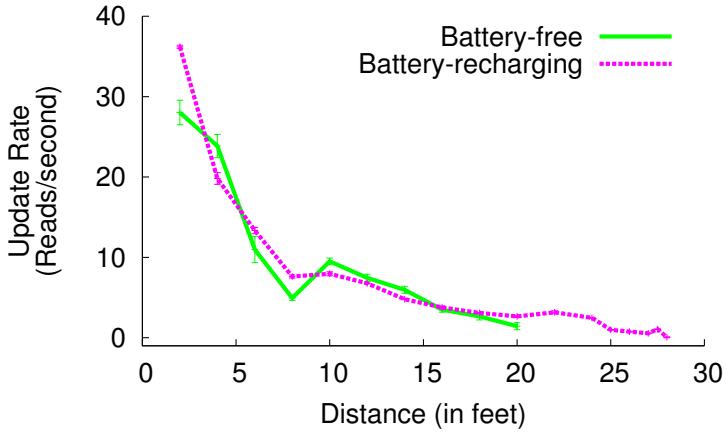

Figure 4.13: **Update rate of temperature sensors.** The battery-free sensor can operate up to 20 feet and the battery-recharging sensor can operate in an energy-neutral manner up to 28 feet.

4.4.1 Wi-Fi powered Temperature Sensor

The battery-free temperature sensor uses our harvester to power an LMT84 temperature sensor [32] and an MSP430FR5969 microcontroller to read and transmit sensor data [36]. The MSP430FR5969 requires at least 1.9 V to run at 1 MHz and boots in less than 2 ms. When the storage capacitor's voltage reaches 2.4 V, the microcontroller boots, samples the temperature sensor, and transmits the reading through a UART port. The microcontroller's firmware is optimized for power: the entire measurement and data-transmission operation uses only 2.77 μJ.

The battery-recharging sensor, on the other hand, consists of our rectifier followed by the TI bq25570 power-management chip [13] to wirelessly recharge two AAA 750 mAh low discharge current NiMH batteries at 2.4 V [44]. We connect the batteries to the TI chip's V_{bat} node. The temperature sensor and microcontroller are powered from the V_{store} node of the chip, which is internally connected to the NiMH battery. The energy per operation is 2.77 μJ as above.

Experiments. We evaluate the effect of distance on the update rate of the temperature sensor. Specifically, we use a PoWiFi router and place both the battery-recharging and battery-free sensor at increasing distances. In the battery-free case, we measure the update rate by computing the time between successive sensor readings. In the battery-operated case, we measure the battery voltage and the charge current flowing into it from the harvester. Since, each temperature sensor

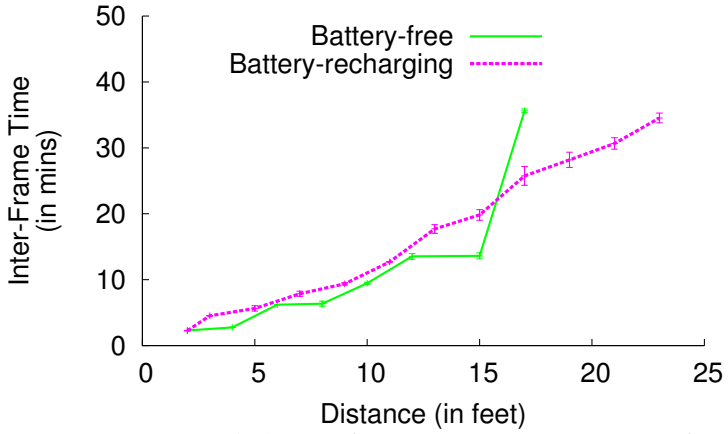

Figure 4.14: **Camera prototype results.** The battery-free camera operates at up to 17 feet and the battery-recharging camera has a range of 23 feet for energy-neutral operations. This enables applications where low-rate cameras can be left in hard-to-reach places, such as walls, attics, and sewers for leakage and structural integrity detection, without the need to replace batteries.

measurement and data transmission takes 2.77 μJ, we compute the ratio of the incoming power to this value to ascertain the update rate of the sensor for energy-neutral operation. The average occupancy across the Wi-Fi channels in our experiments was 91.3%.

Results. Fig. 4.13 plots the results for both our sensors. The update rates decrease with distance from the router. This is a result of less power being harvested and agrees with the harvester benchmarks in §4.3.3. At closer distances, both harvesters have similar update rates. Beyond 15 feet, however, the battery-powered sensor, optimized for lower input power, has a better update rate and extended operational range: it can operate up to 20 feet from the router. The battery-recharging sensor can operate in an energy-neutral manner to greater distances of up to 28 feet.

4.4.2 Wi-Fi powered Camera

We use OV7670, a low-power VGA image sensor from Omnivision [42], and interface it with an MSP430FR5969 microcontroller. The image sensor requires a minimum voltage of 2.4 V and consumes 60 *mW* in active mode operation. We program the sensor to operate in gray-scale QCIF image capture mode with (176×144) pixel resolution. The microcontroller initializes and provides

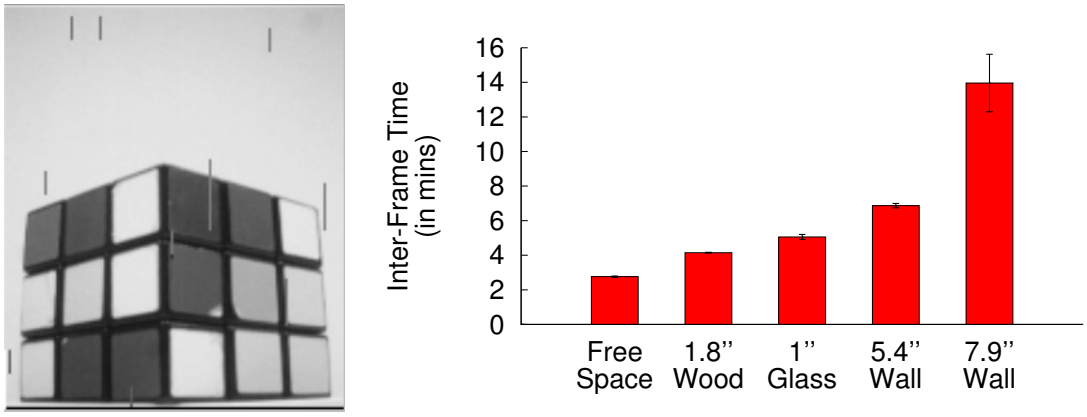

Figure 4.15: **Battery-free camera in through-the-wall scenarios.** The figure on the left is a picture of a Rubik's cube taken with our camera prototype. The plot shows the inter-frame time with different wall materials at a five feet distance from the router.

timing signals to the image sensor. We transfer the sensor data at 48Mbps and store it on the 64 KB non-volatile FRAM on the microcontroller. We optimize our firmware code for power and achieve a per-image capture energy of 10.4 mJ.

On our battery-free camera, we use an ultra-low leakage AVX BestCap 6.8 mF super-capacitor as the storage element [12]. The image sensor and microcontroller are powered by the buck converter of the TI bq25570 chip, which provides 2.55 V regulated output voltage. The TI chip activates the buck converter when the super-capacitor voltage reaches 3.1 V and is active until it discharges to 2.4 V. Our battery-recharging camera consists of the same hardware as before, but uses our wirelessly rechargeable 1 mAh lithium-ion coin-cell battery at 3.0 V [35].

Experiment 1. We evaluate the time between frames as a function of distance for both our camera prototypes. As before, we use a PoWiFi router with an observed average cumulative occupancy of 90.9% across experiments. At each distance from the router, we wait for the camera to take at least six frames and measure the time interval between consecutive frames. For the battery-recharging camera, as before, we ascertain the inter-frame duration for an energy-neutral image capture.

Result 1. Fig. 4.14 shows that the battery-free camera can operate at up to 17 feet from the router, with an image capture every 35 minutes. On the other hand, the battery-recharging camera has an

extended range of 23 feet with an image capture every 34.5 minutes in an energy-neutral manner. Both the sensors have a similar image capture rate at up to 15 feet from the router. We also note that Fig. 4.14 limits the range to 23 feet to better display smaller values. Our experiments, however, show that the battery-recharging camera can operate up to 26.5 feet with an image capture every 2.6 hours.

A key question the reader should ask is: *would cameras with such low image-capture rates be useful in practice?* Taking a picture periodically, as above, is an artificial construct of our experiment. In practice, we could integrate our camera with motion-detection sensors that consume orders of magnitude lower power [114] and turn on the camera only when motion is detected. Another application is to use these cameras in hard-to-reach places such as walls, attics, pipes and sewers for leakage and structural-integrity detection. In these scenarios, replacing batteries can be cumbersome, and our low rate camera sensor would be an effective solution.

Experiment 2. Motivated by the above applications, we next evaluate our camera in through-the-wall scenarios. We place our PoWiFi router next to a wall and place our battery-free camera prototype 5 feet away on the other side of the wall. We experiment with walls of four different materials: a double-pane glass wall of one inch thickness, a wooden door with thickness 1.8 inches, a hollow wall with thickness 5.4 inches, and finally a double sheet-rock (plus insulation) wall with a thickness of 7.9 inches.

Result 2. Fig. 4.15 shows the mean time between frames, averaged over five frames, as a function of the material. The plot shows that as the material absorbs more signals (e.g., double sheet-rock versus glass), the time between frames increases. However, the key conclusion is that PoWiFi can power battery-free cameras through walls and hence enables applications where the cameras can be left in hard-to-reach places such as walls, attics, and sewers, without the need for replacing batteries.

Table 4.1: **Summary of our home deployment**

Home #	1	2	3	4	5	6
Users	2	1	3	2	1	3
Devices	6	1	6	4	2	6
Neighboring APs	17	4	10	15	24	16

4.5 Home Deployment Study

In §4.3.2 we showed that the channel occupancy of PoWiFi can be optimized for different sensor applications and minimize impact on neighboring Wi-Fi devices. However, PoWiFi's ability to efficiently deliver power depends on the traffic patterns of other Wi-Fi networks in the vicinity, as well as the router's own client traffic, both of which can be unpredictable. So we deploy our system in six homes in a metropolitan area and measure PoWiFi's ability to continuously achieve high channel occupancies.

Table 4.1 summarizes the number of users, devices and other 2.4 GHz routers nearby in each of our deployments. We replace the router in each home with a PoWiFi router, and the occupants use it for normal Internet access for 24 hours. Our router uses the same SSID and authentication information as the original router, which we disconnect. We place our router within a few feet of the original router, with the exact location determined by user preferences. In all six deployments, we set our router to provide Internet connectivity on channel 1 and to transmit power packets on channels 1, 6, and 11 using the algorithm in §4.2.1. We stage our deployment over the period of a week— the first two homes in Table 4.1 over a weekend and the rest on weekdays.

We log the router's channel occupancy on each of the three Wi-Fi channels at a resolution of 60 s. Fig. 4.16 plots the occupancy values for each Wi-Fi channel over the 24-hour deployment duration. We also plot the cumulative occupancy across the channels. The figures show that:

- We see significant variation in per-channel occupancy across homes. This is because when the load is high on neighboring networks, our router scales back its transmissions on that channel

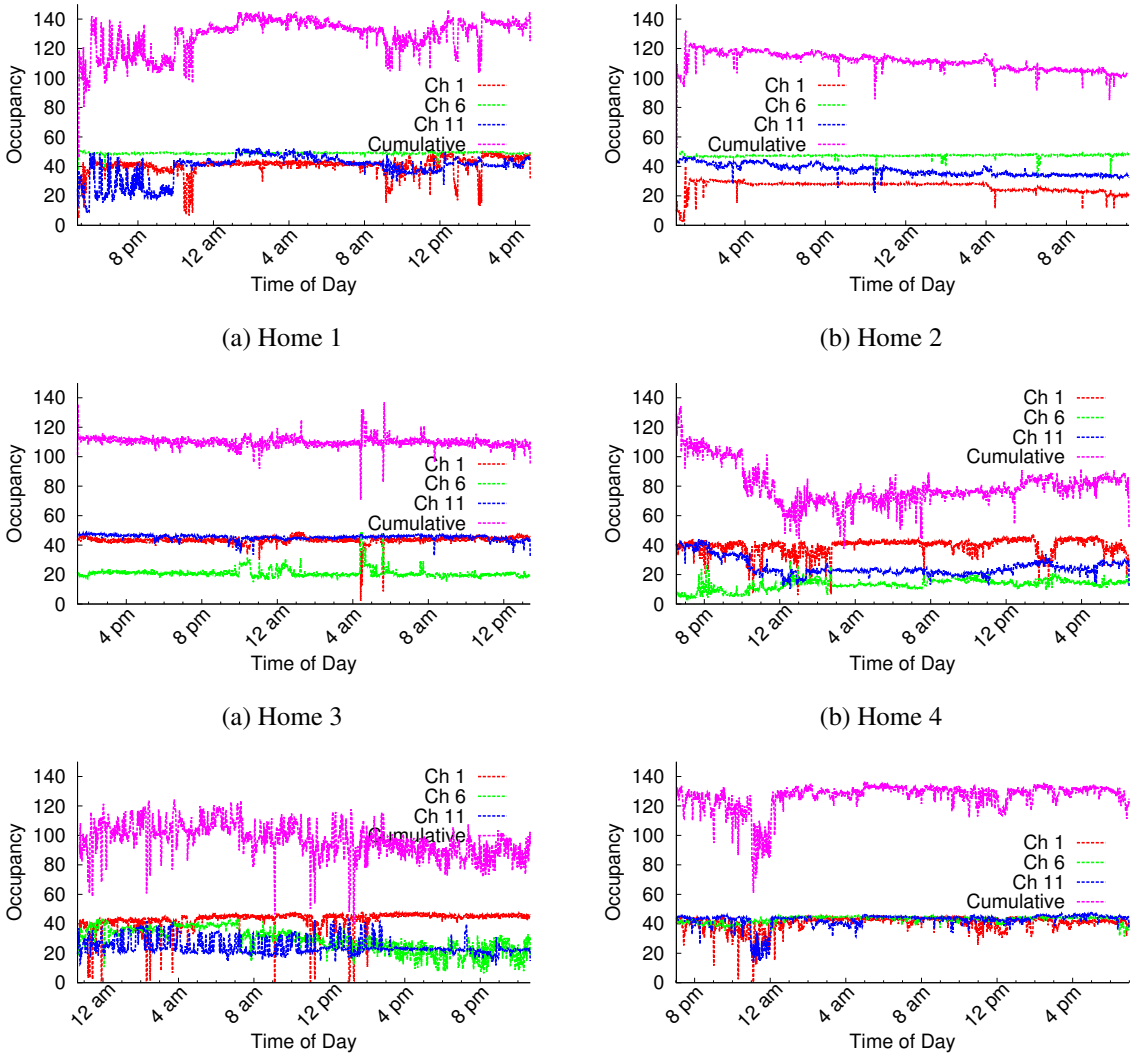

(a) Home 1

(b) Home 2

(a) Home 3

(b) Home 4

(e) Home 5

(f) Home 6

Figure 4.16: **PoWiFi channel occupancies in home deployments.** We see significant variation in per-channel occupancy values across homes. This is because PoWiFi uses carrier sense that reduces its occupancy when the neighboring networks are loaded. The cumulative occupancy, however, is high across time in all home deployments. We note that, in principle, one can modify PoWiFi's algorithm to reduce the per-channel occupancy of the power traffic and keep the cumulative occupancy less than 100%, which is sufficient for harvesting purposes.

and has lower channel occupancy. However, when the load on neighboring networks is low, the router occupies a larger fraction of the wireless channel. This is because PoWiFi uses carrier sense to enforce fairness with other Wi-Fi networks.

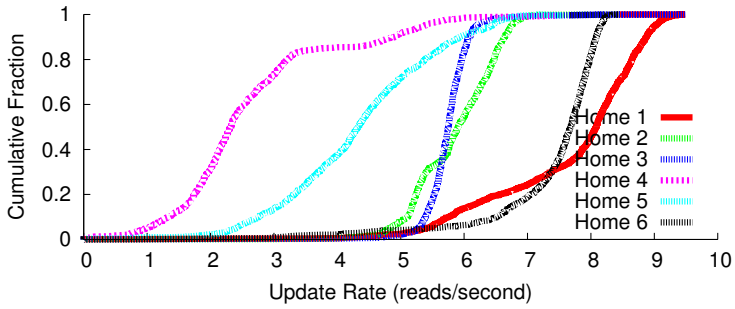

Figure 4.17: **Battery-free temperature sensor across homes.** The computed update rates ten feet away from our router, shows that we can deliver power via Wi-Fi with real-world network conditions.

- The cumulative occupancy is high over time in all our home deployments. Specifically, the mean cumulative occupancies for the six home deployments are in the 78-127% range. We note that some of these occupancies are much greater than 100%, which might not be necessary for power delivery. One can however reduce the per-channel rate of the power traffic based on the cumulative occupancy value to ensure that it is below 100%. Our current system does not implement this feature.

- The users in homes 1–4 did not perceive any noticeable difference in their user experience. The user in home 5, however, noted a significant improvement in page load times and better experience on streaming sites including Hulu, Amazon Prime and YouTube. This was primarily because home 5 originally was using a cheap low-grade router with worse specifications. A user in home 6 noted a slight deterioration in YouTube viewing experience for a 30-minute duration. Our analysis showed that our router occupancy, including both client and power traffic, dipped during this duration. This points to external causes including interference from other devices in the environment.

Finally, Fig. 4.17 plots the CDFs of the computed update rates for our battery-free temperature sensor placed ten feet from the router in the homes. The plots show that we can successfully deliver power via Wi-Fi in real-world Wi-Fi network conditions.

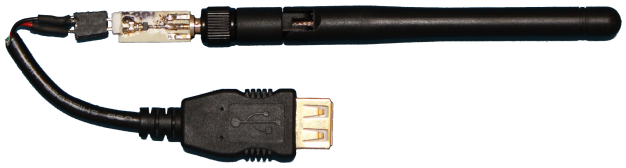

Figure 4.18: **Wi-Fi power via USB.** The charger consists of a 2 dBi Wi-Fi antenna attached to our harvester. Using this, we charge a Jawbone UP24 device in the vicinity of the PoWiFi router from a no-charge state to 41% charged state in 2.5 hours.

4.6 Router as a charging hotspot

In addition to powering custom temperature and camera sensors, PoWiFi can transform the vicinity of a Wi-Fi router into a wireless charging hotspot for devices such as wearable activity trackers. To demonstrate feasibility, we design the general-purpose USB charger shown in Fig. 4.18. It consists of a 2 dBi Wi-Fi antenna attached to a custom harvester that we optimize for higher input power. We then connect our USB charger to a Jawbone UP24 device and place it 5-7 cm away from the PoWiFi router. We observe that the charger supplied an average current of 2.3 mA and charge the Jawbone UP24 battery from a no-charge state to 41% charged-state in 2.5 hours. This demonstrates the potential of our approach. We are currently working on designs that would directly integrate our harvester with the antenna of the wearable device. Further, we are exploring the use of a custom battery charging solution, similar to those demonstrated in this paper, to achieve higher efficiencies and longer-distance wireless charging for these devices.

4.7 Related Work

Wireless power delivery techniques can be primarily divided into two categories: near-field magnetic resonance/inductive coupling [82, 124] and RF power transmission systems. Of the two, RF power delivery is the truly long-range mechanism and hence we focus on the latter category.

Early RF power delivery systems were developed as part of RFID systems to harvest small amounts of power from dedicated 900 MHz UHF RFID readers [187]. The power harvested from RFID signals has been used to operate accelerometers [187], temperature sensors [187], and re-

Table 4.2: **Comparison of our harvester with the state of the art**

	RF Source	Sensitivity	Bandwidth	Startup	Application\ Evaluation
PoWiFi	Wi-Fi	-17.8 dBm @ 2.4 V	100 MHz	Cold start	Temperature/camera
PoWiFi	Wi-Fi	-19.3 dBm @ 2.4/3.0 V	100 MHz	Self start	Recharge battery
[101]	CW	-25 dBm @ 2.4V	N/A	Push button	Recharge battery
[66]	CW	-20 dBm @ 100 mV	75 MHz	Cold start	Rectifier loaded by 8.2 $k\Omega$
[159]	CW	-20 dBm @ 125 mV	100 MHz	Cold start	Rectifier loaded by 10 $k\Omega$
[111]	Microwave oven	-10 dBm @ 150 mV	N/A	Cold start	Rectifier loaded by 10 $k\Omega$

cently cameras [149]. Our efforts on power delivery over Wi-Fi are complimentary to RFID systems. In principle, one can combine multiple ISM bands including 900 MHz, 2.4 GHz, and 5 GHz to design an optimal power delivery system. This paper takes a significant step towards this goal.

Recently, researchers have demonstrated the feasibility of harvesting small amounts of power from ambient TV [112, 131] and cellular base station signals [166, 205] in the environment. While TV and cellular signals are stronger in outdoor environments, they are significantly attenuated indoors, limiting the corresponding harvesting opportunities. The ability to power devices using Wi-Fi can augment the above capabilities and enable power harvesting indoors.

Researchers have also explored the feasibility of harvesting power in the 2.4 GHz ISM band [66, 84, 91, 103, 106, 111, 158–160, 204]. These efforts have demonstrated power harvesting from continuous wave (CW) transmissions[2] and none have powered devices with existing Wi-Fi chipsets. Further, [103, 106, 111] harvest from incoming signals in excess of -5 dBm and can operate only in close proximity of the power source. [66, 160] design a rectifier that outputs voltages around 100 mV for continuous wave transmissions at specific frequency tones. It is unclear how one may transform this into 1.8–2.4 V required by microcontrollers, sensors and batteries. [84] discusses an IC implementation of a 2.45 GHz continuous-wave RFID tag. [91] has recently analyzed the impact of the bursty nature of Wi-Fi traffic on the rectifier and optimizes the size of the rectifier's output capacitor based on Wi-Fi burstiness. However, similar to [66, 160], this work is focused

[2]Continuous wave transmissions are special signals that have a constant amplitude and a single frequency tone.

on rectifier design and does not power sensors and microcontrollers or recharge batteries. We also note that our work takes a different approach to the problem: we mask the burstiness in Wi-Fi traffic and instead create high cumulative channel occupancy at the router. [101] designs an efficient 2.4 GHz rectenna patch and battery charging solution which requires a mechanical push button for startup. The system is evaluated with continuous wave transmissions in an anechoic chamber, and not Wi-Fi signals. In contrast, PoWiFi is the first power over Wi-Fi system that works with existing Wi-Fi chipsets and minimizes its impact on Wi-Fi performance. Table 4.2 shows a summary comparison of our harvester with the state of the art 2.4 GHz harvesters.

Our work is also related to efforts from startups such as Ossia [18] and Wattup [60]. These efforts claim to deliver around 1 W of power at ranges of 15 feet and charge a mobile phone [22]. Back-of-the-envelope calculations however show that this requires continuous transmissions with an EIRP (equivalent isotropic radiated power) of 83.3 dBm (213 kW). This not only jams the Wi-Fi channel but is also 50,000 times higher power than allowed by FCC regulations part 15 for point to multi-point links. In contrast, our system is designed to operate within the FCC limits and has minimal impact on Wi-Fi traffic. We note that in the event of an FCC exception to these startups, our multi-channel design can be used to deliver such high power without having significant impact on Wi-Fi performance.

Finally, recent work on Wi-Fi backscatter [113] enables low-power connectivity with existing Wi-Fi devices. Backscatter communication is order of magnitude more power-efficient than traditional radio communication and hence enables Wi-Fi connectivity without incurring Wi-Fi's power consumption. However, [113] is focused on the communication mechanism and to the best of our knowledge, does not evaluate the feasibility of delivering power using Wi-Fi. Our work is complementary to [113] and can in principle be combined to achieve both power delivery and low-power connectivity using Wi-Fi devices.

4.8 Conclusion

There is increasing interest in the Internet-of-Things where small computing sensors and mobile devices are embedded in everyday objects and environments. A key issue is how to power these

devices as they become smaller and more numerous; plugging them in to provide power is inconvenient and is difficult at large scale. We introduce a novel far-field power delivery system using existing Wi-Fi chipsets. We do so while minimizing the impact on Wi-Fi network performance. While this is a first step towards using Wi-Fi chipsets for power delivery, we believe that with subsequent iterations of the harvester design we can significantly increase the capabilities of our system.

Chapter 5

BRINGING LOW POWER TO WI-FI COMMUNICATION

5.1 Introduction

Wireless Fidelity popularly known as Wi-Fi, is one of the most successful and widely used wireless protocols. Currently, there are about 6.8 billion active Wi-Fi devices including but not limited to enterprise and home routers, PCs, smartphones, tablets, cellphones, making Wi-Fi one of the most prolific technologies around the world [61]. However, Wi-Fi suffers from one key downside, it is extremely power hungry. A typical Wi-Fi radio consumes around 600-700 mW, making it impractical for most energy constrained applications such as IoT and wearable devices [24, 48]. Over the past few years, researchers have explored backscatter to reduce the power consumption of Wi-Fi communication. Wi-Fi backscatter [113] and BackFi [168] create additional narrowband data stream to ride on top of existing Wi-Fi signals. Wi-Fi backscatter decodes this narrow band data on existing Wi-Fi receivers by leveraging CSI and RSSI values. BackFi takes a different approach, it uses a full-duplex radio to suppress self-interference and decode backscattered bits using a custom receiver. While promising, these designs either achieve very low data rates (100s of bps) at close by distances (2-4 feet) [113] or use custom full-duplex hardware that are not available on any existing Wi-Fi devices [168].

In this work, we take a different approach — instead of backscattering existing Wi-Fi signals to send an additional data stream, we use backscatter communication to *directly generate Wi-Fi transmissions* that can be decoded on any of the billion existing devices with a Wi-Fi chipset. To this end, we introduce Passive Wi-Fi that demonstrates for the first time that one can generate 1-11 Mbps 802.11b transmissions using backscatter communication, while consuming 4–5 orders of magnitude lower power than existing Wi-Fi chipsets. Our work is related to [90] where authors generate 1 Mbps BLE packets using backscatter and receive them on Bluetooth radios.

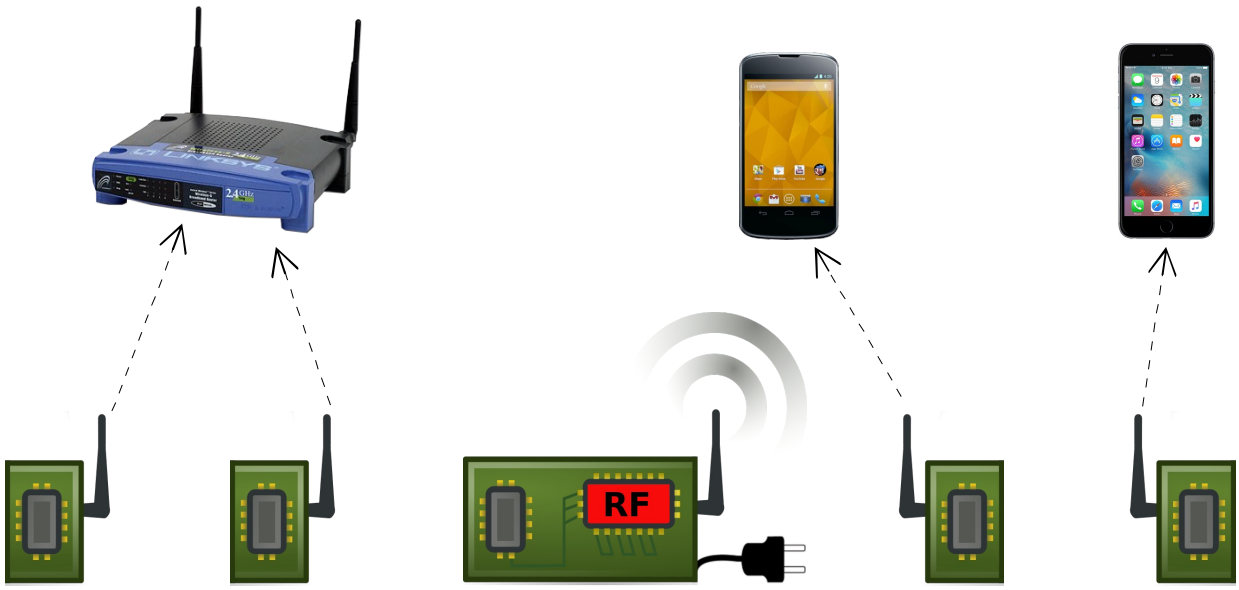

Figure 5.1: **Passive Wi-Fi architecture.** The passive Wi-Fi devices perform digital baseband operations like coding, while the power-consuming RF functions are delegated to a plugged-in device in the network.

Traditional Wi-Fi radios consists of two main components: digital baseband and the analog RF front end. Due to CMOS technology scaling, the digital baseband has experienced exponential reduction in size and power consumption. However, the analog RF front end, which is necessary for generating and receiving RF signal for Wi-Fi communication, has not seen similar power scaling. As a result, Wi-Fi transmissions on sensors and mobile devices still consume hundreds of mill watts of power [135, 138, 147]. To get around this problem, passive Wi-Fi uses backscatter to decouple the baseband Wi-Fi digital logic from the power-consuming RF components, as shown in Fig. 5.1.

In our architecture, the passive Wi-Fi devices perform digital baseband operations like coding and modulation, while the power-consuming RF components such as frequency synthesizers and power amplifiers are delegated to a single plugged-in device in the network. This device provides the RF functions for all the passive Wi-Fi devices in the vicinity by transmitting a single-frequency tone. The passive Wi-Fi devices create 802.11b transmissions by reflecting or absorbing this tone using a digital switch running at baseband frequency of 10's of MHz. Since the passive Wi-Fi devices have no analog RF components, they consume less silicon area and would be smaller and

cheaper than existing Wi-Fi chipsets. More importantly, their power consumption would be orders of magnitude lower since they only perform digital baseband operations. To realize this, however, we need to address three main challenges.

(a) How can Wi-Fi receivers decode in the presence of interference from the plugged-in device? The Wi-Fi receiver receives the backscattered signal in the presence of a strong interference from the tone transmitted by the plugged-in device. Traditional backscatter systems [153, 168] use a full-duplex radio to cancel this strong interfering signal, which is not possible on existing Wi-Fi devices. Bistatic backscatter radios use band-pass filters to filter the data packet from the single tone interference [90, 119, 120, 183]. We leverage the key observation that Wi-Fi receivers also have out of band interference rejection components such as band pass filters to ensure functionality even in the presence of interference in the adjacent band that is 35 dB stronger [28]. Further, as Wi-Fi and Bluetooth radios are being integrated onto a single chipset [16], Wi-Fi hardware is being designed with stringent band pass filtering to make it work in the presence of out-of-band Bluetooth interference. Thus, we set the plugged-in device to transmit its tone at a frequency that lies outside the desired Wi-Fi channel; this ensures that existing Wi-Fi chipsets can suppress the resulting out-of-band interference.

(b) How can we create 802.11b transmissions using backscatter? Our design leverages two key ideas. 1) 802.11b Wi-Fi packets can be generated using only digital logic. 2) We can use sub-carrier modulated backscatter [85, 90, 92, 119, 120] to create 802.11b Wi-Fi packets at a frequency offset from the single tone. 802.11b uses DSSS and CCK encoding with DBPSK and DQPSK modulation. The encoding operation is digital in nature and to create the phase changes required for DBPSK and DQPSK, we approximate a digital square wave as a sinusoid and modulate its phase by changing the timing of the square wave (see §5.2.3). Thus, we can generate baseband Wi-Fi packet using only digital logic. Next, to up convert baseband Wi-Fi packet to RF frequencies, we leverage sub-carrier modulation, a standard technique in backscatter systems [85, 90, 92, 119, 120]. First we multiply the Wi-Fi packet with a square wave of Δf frequency which creates two mirror

copies of baseband Wi-Fi packet at Δf frequency offset on either side of DC (see Fig. 5.2)[1]. Then we input this data to the backscatter switch which mixes the baseband data with the RF carrier to up convert the mirror copies to RF. In §5.2.4, we show how we can eliminate the mirror copy and create a single side band Wi-Fi packet using a novel sub-carrier modulation technique and complex impedance backscatter switch.

(c) How do passive Wi-Fi devices share the Wi-Fi network? Traditional Wi-Fi devices shares the wireless medium using carrier sense. However, this requires a Wi-Fi receiver that is ON before every transmission. Since, Wi-Fi receivers require power-consuming RF components such as ADCs and frequency synthesizers, this would eliminate the power savings from our design. Instead, we delegate the power-consuming task of carrier sense to the plugged-in device. At a high level, the plugged-in device performs carrier sense and signals the passive Wi-Fi device to transmit. §5.3 describes how such a signaling mechanism can also be used to arbitrate the channel between multiple passive Wi-Fi devices and address other link-layer issues including ACKs and retransmissions.

To show the feasibility of our design, we build backscatter hardware prototypes and implement all four 802.11b bit rates on an FPGA platform. Our experimental evaluation shows that passive Wi-Fi transmissions can be decoded on off-the-shelf smartphones and Wi-Fi chipsets over distances of 30–100 feet in various line-of-sight and through-the-wall scenarios. We also design a passive Wi-Fi IC that performs 1 Mbps and 11 Mbps 802.11b transmissions and estimate the power consumption using Cadence and Synopsis toolkits [14, 54]. Our results show that 1 and 11 Mbps passive Wi-Fi single side band transmissions consume 14.5 and 59.2 μW respectively.

Contributions. We make the following contributions:

- We demonstrate for the first time that one can generate 802.11b transmissions using backscatter communication. We present backscatter techniques that synthesize 22 MHz DSSS and CCK spread spectrum transmissions that can be decoded on existing Wi-Fi devices.

- We introduce single side band backscatter technique to eliminate the mirror copy generated

[1] $2 sin ft sin \Delta ft = cos(f - \Delta f)t - cos(f + \Delta f)t$.

using traditional sub-carrier modulation. Using single side band backscatter we synthesize 802.11b 22 MHz DSSS and CCK spread spectrum transmissions that have the same spectrum efficiency as traditional Wi-Fi radios.

- We design a network stack for the passive Wi-Fi transmitters to coexist with other devices in the ISM band. Further, we present a detailed analytical model to understand the operational range of passive Wi-Fi transmissions in different deployment scenarios.

- We build a hardware prototype on an FPGA platform and evaluate it in various scenarios. We also design a passive Wi-Fi IC and present its power numbers.

5.2 Passive Wi-Fi Design

Our design has two main actors: a plugged-in device and passive Wi-Fi devices. The former contains power consuming RF components including frequency synthesizer and power amplifier and emits a single tone RF carrier. It also performs carrier sense on behalf of the passive Wi-Fi device and helps coordinate medium access control across multiple passive Wi-Fi devices. The passive Wi-Fi device backscatters the tone emitted by the plugged-in device to synthesize 802.11b transmissions that can be decoded on any device that has a Wi-Fi chipset.

In the rest of this section, we first provide a quick primer for 802.11b physical layer and backscatter communication. We then explain how the passive Wi-Fi devices generate 802.11b packets using backscatter communication. We then theoretically analyze the range of our transmissions in various deployments scenarios.

5.2.1 Primer for 802.11b Transmissions

802.11b is a set of Wi-Fi physical layer specifications that use spread spectrum modulation. 802.11b uses DBPSK/DQPSK at the physical layer and achieves four bit rates using different spreading codes. The lower two bit rates of 1 and 2 Mbps use direct-sequence spread spectrum (DSSS) while 5.5 and 11 Mbps use complementary code keying (CCK). DSSS uses a single code to spread the

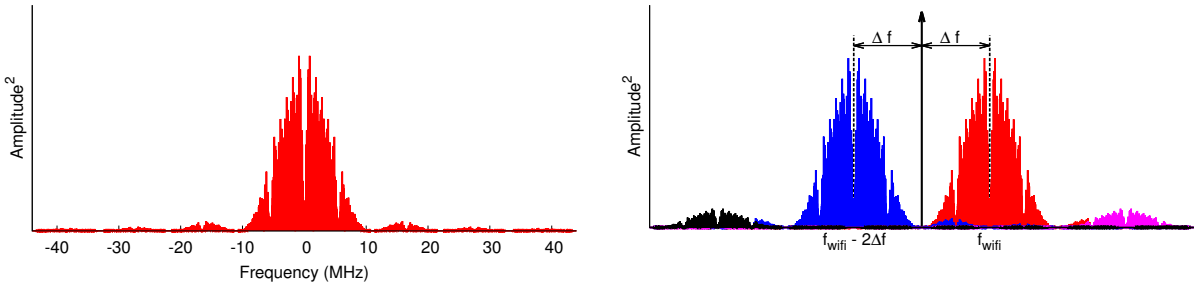

Figure 5.2: **Generation of Wi-Fi packets using backscatter.** The plot on the left shows the 22 MHz main lobe and the side lobes of the baseband 802.11b packet in the frequency domain. The plot on the right illustrates the backscatter operation at the passive Wi-Fi device. The two main lobes are shifted by Δf with respect to the constant tone emitted by the plugged-in device to generate the Wi-Fi packet (in red) at f_{wifi} and a mirror image (in blue) at $f_{wifi} - 2\Delta f$.

information over 22 MHz, while CCK uses a set of multiple code words to both encode bits and also achieve a 22 MHz spread spectrum signal. We outline how each of the 802.11b bit rates are encoded.

1 and 2 Mbps DSSS transmissions. To generate this, 802.11b first creates coded bits from the incoming data using a 11-bit barker code [173]. Specifically, 802.11b uses a single barker sequence, 10110111000, that is generated at a baseband frequency of 11 MHz to spread the spectrum over 22 MHz. To create the coded bits, 802.11b XORs each of the data bits with the barker sequence. Thus, the coded bits for a '1 data bit are 10110111000 and that for the '0 data bit are 01001000111. Each of these coded bits is encoded using DBPSK and DQPSK modulation to achieve 1 and 2 Mbps transmissions respectively. At a high level, this is achieved by setting the phase of the carrier, $sin\theta$. DBPSK modulation encodes a 0 and 1 bit by setting θ to either 0 or π, while DQPSK encodes pairs of bits by modulating the phase between 0, $\pi/2$, π and $3\pi/2$.

5.5 and 11 Mbps CCK transmissions. Instead of using a single barker code, CCK uses a set of 8-bit code words. At a high level, to generate 5.5 Mbps transmissions, the incoming data bit stream is divided into blocks of 4 bits. The first two bits are used to pick the DQPSK phase and the last two bits are used to pick a spreading code amongst four 8-bit code words. To generate 11 Mbps 802.11b transmissions, the incoming data bits are instead divided into 8 bit blocks where the first two bits determine the DQPSK phase shift and the last 6 bits are used to pick a spreading code

amongst 64 8-bit code words.

To summarize, 802.11b requires generation of the coded bits using either DSSS or CCK and then modulating these bits with DBPSK or DQPSK. The first operation is typically implemented in digital baseband logic while the second requires changing the phase of I and Q components. Finally, we also note that since the RF energy is spread across a wide band, spread spectrum transmissions are resilient to narrowband interference both within and outside the Wi-Fi channel [173].

5.2.2 Backscatter Communication Primer

Unlike traditional active radio communication that requires generating RF signals, devices using backscatter communicate by modulating the radar cross-section of their antenna to change the reflected signal. To understand how backscatter works, consider a device that can switch the impedance of its antenna between two states. The effect of changing the antenna impedance is that the radar cross-section, i.e., the signal reflected by the antenna, also changes between the two different states. Now, given an incident signal with power $P_{incident}$, the power in the differential backscattered signal can be written as,

$$P_{backscatter} = P_{incident} \frac{|\Gamma_1^* - \Gamma_2^*|^2}{4} \tag{5.1}$$

Here Γ_1^* and Γ_2^* are the complex conjugates of the reflection coefficients corresponding to the two impedance states. Thus to maximize the power in the backscattered signal we need to maximize the difference in the power of the two impedance states which is given by $|\Delta\Gamma|^2 = \frac{\left|\Gamma_1^* - \Gamma_2^*\right|^2}{4}$. Ideally, to maximize the power in the backscattered signal, we set $|\Delta\Gamma|^2$ to 4 which can be achieved by modulating the reflection coefficients between $+1$ and -1. In practice, however, backscatter hardware deviates from this ideal behavior and incurs losses; our hardware implementation has a loss of around 1.1 dB for double side band backscatter and 2.2 dB for single side band backscatter design.

5.2.3 Generating 802.11b Wi-Fi packets using backscatter

Generating a Wi-Fi packet using backscatter is challenging for two main reasons. First, the backscattered signal is much weaker than the tone transmitted by the plugged-in device. A Wi-Fi receiver would suffer significant in-band interference from this tone, thereby preventing it from decoding. Second, the passive Wi-Fi device has a single digital switch that toggles between two impedance states, resulting in a binary signal. It is unclear how one may generate Wi-Fi transmissions using such a binary system.

We outline how to address these challenges. We first describe the signal transmissions from the plugged-in device and then the operations at the passive Wi-Fi device that allow us to synthesize 802.11b transmissions.

Transmissions at the plugged-in device. It transmits a tone outside the desired Wi-Fi channel. Our key intuition is that Wi-Fi receivers are designed to function in the presence of out-of-band interference: 802.11b receivers are required to ensure that the sensitivity is reduced by no more than 6 dB in the presence of interference in the adjacent band that is 35 dB greater than the in-band signal [28]. Further, as Wi-Fi and Bluetooth radios are being integrated onto the same chipset [16], Wi-Fi frontends are being designed to function in the presence of out-of-band interference from Bluetooth devices. Since the tone from the plugged-in device is narrower in bandwidth than Bluetooth, this would further help suppress the tone if it is outside the desired Wi-Fi channel.

We note however that excessive out-of-band interference, which occurs when the Wi-Fi receiver is right next to the plugged-in device, can saturate and/or compress the RF front end resulting in significant degradation of Wi-Fi performance. This is called the input 1 dB compression point which is around 0 dBm for commercial Wi-Fi devices [33]. Passive Wi-Fi inherently avoids this issue by ensuring that the Wi-Fi receiver (e.g., smartphone or router) is not next to the plugged-in device.

Backscatter operations at passive Wi-Fi devices. At a high level, passive Wi-Fi uses sub-carrier modulation technique [85, 90, 92, 119, 120] from backscatter communication to create DBPSK and DQPSK modulated 802.11b Wi-Fi packets. Sub-carrier modulation is a commonly used technique

and has been demonstrated for amplitude shift keying (ASK) [85,92], phase shift keying (PSK) [85, 92] and frequency shift keying (FSK) [90,119,120] modulation in backscatter communication. The key idea behind sub-carrier modulation is to shift the baseband data to a frequency offset from the single tone. By doing so, we ensure that the single tone transmission is out of the frequency band of the transmitted data and can be easily filtered out by the receiver. In this work, we use sub-carrier modulation to shift the 802.11b Wi-Fi packet from the tone transmitted from the plugged-in device to lie at the center of the desired Wi-Fi channel. To do this, we leverage three key facts: (1) Using sub-carrier modulation i.e. multiplying the baseband data with a sinusoidal signal, we can create a frequency shift. From basic trigonometry, $2sin(ft)sin(\Delta ft) = cos(f-\Delta f)t - cos(f+\Delta f)t$. (2) All bit rates in 802.11b are differentially phase modulated using DBPSK or DQPSK. We implement phase modulation by simply modulating the phase of the sub-carrier sinusoidal signal $sin(\Delta ft)$. (3) Backscatter switch implements mixing operation i.e. the switch multiplies the baseband data with the incoming RF signal. Thus, by just modulating the antenna with the baseband data, we can up convert the baseband data to RF. Since the baseband Wi-Fi packet is at a frequency offset of Δf from DC, the up convert RF signal has the Wi-Fi packet at a frequency offset of Δf from the single tone transmitted by the plugged-in device.

Step 1. Baseband sub-carrier modulation. Let's say that the plugged-in device transmits a tone $cos\left(2\pi(f_{wifi}-\Delta f)t\right)$ outside the Wi-Fi channel. Passive Wi-Fi devices use a square wave at a frequency of Δf for sub-carrier modulation. From Fourier analysis, we can write the square wave as,

$$Square\left(\Delta ft\right) = \frac{4}{\pi}\sum_{n=1,3,5,..}^{\infty}\frac{1}{n}sin\left(2\pi n\Delta ft\right)$$

Here the first harmonic is a sinusoidal signal at the desired frequency Δf. Note that the power in each of these harmonic scales as $\frac{1}{n^2}$. So the third and the fifth harmonic are around 9.5 dB and 14 dB lower than the first harmonic. Thus, we can approximate a square wave as just the sinusoidal signal, $\frac{4}{\pi}sin\left(2\pi\Delta ft\right)$. Now, in digital logic, we multiply the baseband data with the square wave and

considering only the first harmonic of the square wave, we can write the output of this operation as

$$S_{BB}(t) * Square(\Delta ft) \approx S_{BB}(t) * \frac{4}{\pi} sin(2\pi\Delta ft)$$
$$S_{sub-carrier}(t) = S_{BB}(t) * \frac{2}{\pi j} e^{2\pi\Delta ft} - S_{BB}(t) * \frac{2}{\pi j} e^{-2\pi\Delta ft}$$

In the frequency domain we can write this as

$$S_{sub-carrier}(f) = \frac{2}{\pi j} S_{BB}(f - \Delta f) - \frac{2}{\pi j} S_{BB}(f + \Delta f)$$

So using sub-carrier modulation, we create two copies of the baseband Wi-Fi packet, one centered at Δf and the other at $-\Delta f$.

Step 2. Synthesizing baseband 802.11b packets. Now that we know that we can shift the baseband packet using sub-carrier modulation, the next step is to create 802.11b transmissions. 802.11b uses DSSS and CCK encoding which are both digital operations and hence can be performed using digital logic at the passive Wi-Fi device. To implement DBPSK and DQPSK modulation, Passive Wi-Fi notes that DBPSK and DQPSK use a sine wave with four distinct phases: $0, \pi/2, \pi, 3\pi/2$. Since the square wave $Square(\Delta ft)$ generated by our digital switch can be approximated as a sine wave, we can generate the required four phases by changing the timing of our square wave. Specifically, shifting the square wave by half of a symbol time, effectively creates a phase change of π. Phase changes of $\pi/2$ and $3\pi/2$ can be achieved by shifting the square wave by one-fourth and three-fourth of a symbol time respectively. Thus, passive Wi-Fi devices can fully operate in the digital domain while running at a baseband frequency of a few tens of MHz and synthesize 802.11b transmissions.

Step 3. Mixing operation at the backscatter switch. Finally, the sub-carrier modulated baseband Wi-Fi packet is fed to the backscatter switch. The backscatter switch mixes the single tone $cos(2\pi(f_{wifi} - \Delta f)t)$ transmission from the plugged-in device with the baseband data to generate

$$\begin{aligned} S_{RF}(t) &= S_{sub-carrier}(t) * cos(2\pi(f_{wifi} - \Delta f)t) \\ &= \left(S_{BB}(t) * \frac{2}{\pi j} e^{2\pi\Delta ft} - S_{BB}(t) * \frac{2}{\pi j} e^{-2\pi\Delta ft}\right) * \left(\frac{1}{2} e^{2\pi\left(f_{wifi} - \Delta f\right)t} + \frac{1}{2} e^{-2\pi\left(f_{wifi} - \Delta f\right)t}\right) \\ &= S_{BB}(t)\left(\frac{1}{\pi j} e^{2\pi f_{wifi} t} + \frac{1}{\pi j} e^{-2\pi f_{wifi} t}\right) - S_{BB}(t)\left(\frac{1}{\pi j} e^{2\pi\left(f_{wifi} - 2\Delta f\right)t} + \frac{1}{\pi j} e^{-2\pi\left(f_{wifi} - 2\Delta f\right)t}\right) \end{aligned}$$

In the frequency domain, this can be written as

$$S_{RF}(f) = \frac{1}{\pi j}\{S_{BB}(f - f_{wifi}) + S_{BB}(f + f_{wifi})\} - \frac{1}{\pi j}\{S_{BB}(f - f_{wifi} + 2\Delta f) + S_{BB}(f + f_{wifi} - 2\Delta f)\}$$

We note the following about our design.

In addition to creating an 802.11b transmission centered at f_{wifi}, as shown in Fig. 5.2, our backscatter mechanism also creates a mirror copy centered at $f_{wifi} - 2\Delta f$ on the other side of the tone. Thus, we use twice the bandwidth of a traditional 802.11b transmission. This is the tradeoff we make to achieve orders of magnitude lower power consumption. We note that such a tradeoff is common in 802.11n systems which use channel bonding of adjacent Wi-Fi channels to double the throughput.

- 802.11b transmissions have side lobes (Fig. 5.2); the side lobes of the mirror copy create interference for the desired Wi-Fi signal. We plot the signal to interference ratio for different frequency shifts, Δf, at the passive Wi-Fi device. Fig. 5.3(a) plots the results for all four 802.11b bit rates and shows that the interference from the side lobes of the mirror copy reduces as Δf increases. This is because, as Δf increases, the mirror copies are further separated in frequency, resulting in lower interference.

- An effect of this interference, however, is that it adds additional noise to the Wi-Fi signal, reducing the noise sensitivity at which each of the 802.11b bit rates can be decoded. Fig. 5.3(b) shows the loss in sensitivity for the four 802.11b bit rates, as a function of the frequency offset, Δf. The plots show that the sensitivity loss is slightly larger for higher 802.11b bit rates. This is because higher bit rates require a cleaner signal to successfully be decoded. Our system sets Δf to 12.375 MHz, where the sensitivity loss is less than 2 dB across all 802.11b bit rates. This also ensures that the passive Wi-Fi transmissions only occupy two adjacent Wi-Fi channels. Note that Wi-Fi applies filters to remove the interfering side lobes. Our implementation however does not do this.

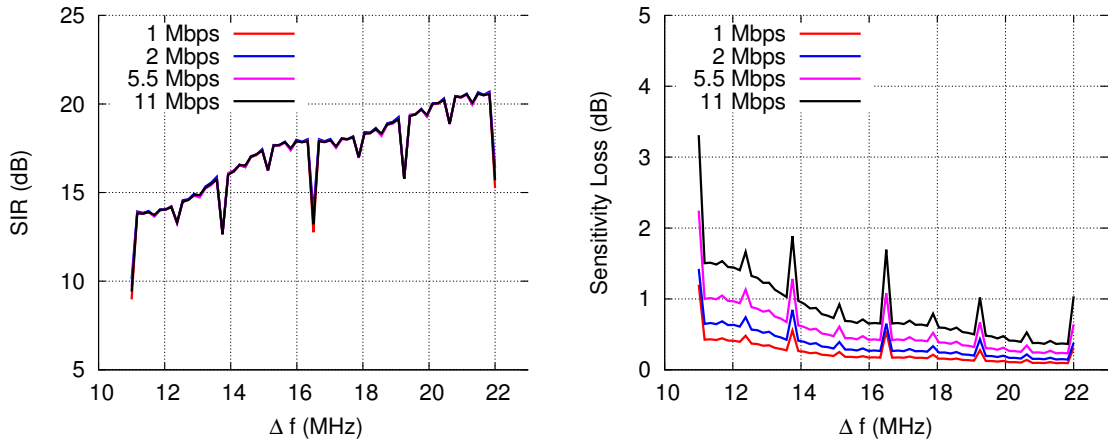

(a) Signal to Interference ratio *(b) Sensitivity Loss*

Figure 5.3: **SIR and loss in receiver sensitivity.** The plot shows the effect of different Δf's on the quality and the sensitivity of the synthesized Wi-Fi packets.

5.2.4 Generating 802.11b Wi-Fi packets using single side band backscatter

The technique described in §5.2.3 generates two copies of the Wi-Fi packet, one at f_{wifi} and another one of $f_{wifi} - 2\Delta f$. This is problematic because we are using twice the bandwidth which reduces the spectral efficiency of Wi-Fi by half and creates interference on neighboring channels in the crowded 2.4 GHz ISM band. Additionally, the mirror copy creates self-interference which degrades the sensitivity of the Wi-Fi packet as shown in Fig. 5.3. To improve the spectral efficiency of Wi-Fi, we present the first single sideband backscatter architecture that produces a single copy of Wi-Fi packet using complex sub-carrier backscatter modulation.

In traditional radios, 2.4 GHz oscillators are used to generate the orthogonal signals, $cos2\pi ft$ and $sin2\pi ft$. These are multiplied with digital in-phase, $I(t)$ and quadrature phase components, $Q(t)$ to create $I(t)cos2\pi ft + jQ(t)sin2\pi ft$. This generates the Wi-Fi packet at RF without any mirror copies. The challenge is that we cannot use oscillators running at 2.4 GHz since they consume significant power. Our insight is that mathematically we can imitate the above operations using complex impedances on the backscatter device without 2.4 GHz oscillators.

Let's assume that we could create the following complex signal, $e^{j2\pi\Delta ft}$ and use it as our complex sub-carrier. If we mix this signal with the baseband Wi-Fi packet, in the frequency domain

we can write

$$\mathscr{F}\left(S_{BB}(t) * e^{j2\pi\Delta ft}\right) = S_{BB}(f + \Delta f) \tag{5.2}$$

The above operation creates the desired shift without a mirror copy. So if we can create the complex signal $e^{j2\pi\Delta ft}$ using backscatter, we can achieve single-sideband backscatter modulation. We can write the complex sinusoid as,

$$e^{j2\pi\Delta ft} = cos(2\pi\Delta ft) + jsin(2\pi\Delta ft) \tag{5.3}$$

Step 1. We approximate the sin/cos terms in (5.3) using a square wave going between the two values, +1 to -1,[2] As described in §5.2.3, we approximate the square wave as $\frac{4}{\pi}sin(2\pi\Delta ft)$. For the cosine term, we time shift the square wave by a quarter of the time period.

Step 2. In §5.2.3, we showed that the baseband Wi-Fi packet can be generated in the digital domain. The next step is to implement DPSK and DQPSK modulation using the complex sub-carrier. We can write the Wi-Fi packet in terms of in phase and quadrature phase $S_{BB}(t) = (S_I(t) + jS_Q(t))$ and substitute this to (5.2) to get

$$\begin{aligned}(S_I(t) + jS_Q(t))\, e^{j2\pi\Delta ft} &= (S_I(t) + jS_Q(t)) * (cos2\pi\Delta ft + jsin2\pi\Delta ft) \\ &= S_I cos(2\pi\Delta ft) + jS_I sin(2\pi\Delta ft) + jS_Q cos(2\pi\Delta ft) - S_Q sin(2\pi\Delta ft) \\ &= \{S_I cos(2\pi\Delta ft) - S_Q sin(2\pi\Delta ft)\} + j\{S_I sin(2\pi\Delta ft) + S_Q cos(2\pi\Delta ft)\}\end{aligned} \tag{5.4}$$

In digital baseband we independently generate the the real and imaginary part shown in (5.4). We note that 802.11b Wi-Fi uses DBPSK and DQPSK modulation schemes. For DBPSK, S_I can take +1 and -1 values and S_Q is set to zero. The sine and cosine waveforms also take +1 and -1 values which translates to (5.4) taking values in the set $\{1+j,\ 1-j,\ -1+j,\ -1-j\}$. For DQPSK, $S_I(t) + jS_Q(t)$ are set to one of the values of the set $\{1,-1,j,-j\}$. Again, the complex sub-carrier

[2]Since digital operations are on 0 and 1 bits, in practice we perform step 1 using a square wave between 1 and 0 instead of +1 and -1. This is however a straightforward mapping with a DC offset. For simplicity however, we explain our design using +1 and -1.

modulated baseband Wi-Fi signal shown in (5.4) can take one of the following values: $\{1+j,\ 1-j,\ -1+j,\ -1-j\}$.

Step 3. The final step is the mixing operation at the backscatter switch. For a constant tone of $e^{j2\pi\left(f_{wifi}-\Delta f\right)t}$ from the plugged-in device, we can represent the complex mixing operation as

$$S_{BB}(t) * e^{j2\pi\Delta ft} * e^{j2\pi\left(f_{wifi}-\Delta f\right)t} = S_{BB}(t) * e^{j2\pi f_{wifi}t}$$

So, the baseband Wi-Fi packet is translated to the desired Wi-Fi channel using backscatter. However, to accomplish this we must map the $\{1+j,\ 1-j,\ -1+j,\ -1-j\}$ set of complex values to RF frequencies using a backscatter switch. We note that the set is a QPSK constellation map and we can create the complex values by changing the impedance of the backscatter switch [200]. Additionally, since the information is contained in the phase domain, first we normalize the constellation map to $\{\ (1+j)/\sqrt{2}\ ,\ (1-j)/\sqrt{2}\ ,\ (-1+j)/\sqrt{2}\ ,\ (-1-j)/\sqrt{2}\ \}$. Next, we note that the signal backscattered by the antenna is proportional to $(1+\Gamma^{*})$ where

$$\Gamma = \frac{Z_a^{*} - Z_{in}}{Z_a + Z_{in}}$$

Z_a represents the impedance of the antenna and Z_{in} is the impedance state connected to the antenna. We set the Z_{in} to four impedance states corresponding to $j(1-\sqrt{2})Z_a$, $j(-1+\sqrt{2})Z_a$, $j(-1-\sqrt{2})Z_a$ and $j(1+\sqrt{2})Z_a$ respectively to achieve the four desired QPSK complex values, $\{\ (1+j)/\sqrt{2}\ ,\ (1-j)/\sqrt{2}\ ,\ (-1+j)/\sqrt{2}\ ,\ (-1-j)/\sqrt{2}\ \}$. The antenna impedance, Z_a, is typically 50 ohms. By switching between these impedance states, we can generate the desired complex signal $S_{BB}(t)\, e^{j2\pi\Delta ft}$ and hence achieve single-sideband backscatter modulation.

Fig. 5.4 shows the frequency spectrum for the backscatter generated Wi-Fi signals at 1 Mbps using single-sideband backscatter with an arbitrary frequency shift of 22 MHz. For comparison, we also plot the spectrum using double side band backscatter approach [115]. The plots show that double sideband backscatter design creates a strong mirror copy on the other side of the single tone. However, single-sideband backscatter, introduced in this section, eliminates this mirror copy.

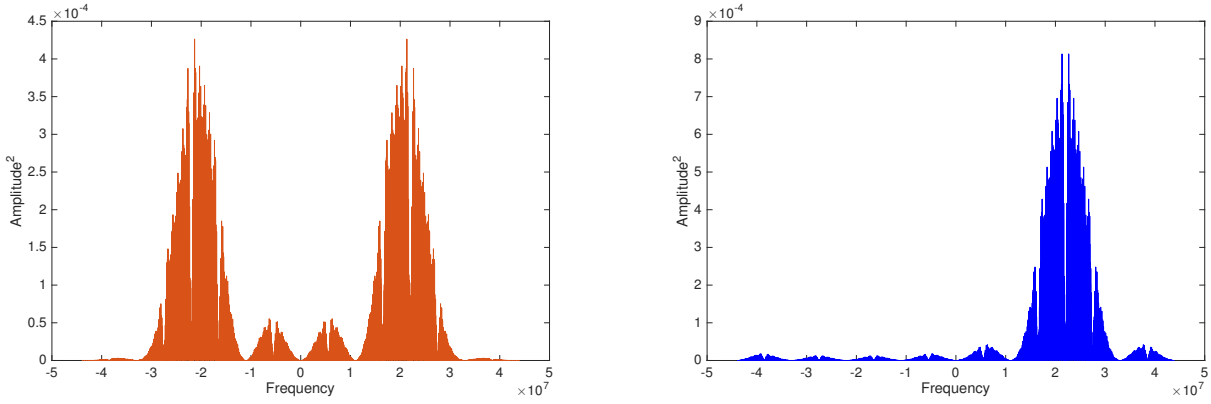

(a) Spectrum of baseband double side band Wi-Fi packet

(b) Spectrum of baseband single side band Wi-Fi packet

Figure 5.4: Comparing single sideband backscatter (blue signal) to prior double sideband backscatter approach (red signal).

5.2.5 Analyzing Passive Wi-Fi's Range

In passive Wi-Fi, the communication range depends on two parameters: the distance between the plugged-in device and the passive Wi-Fi transmitter and the distance between the passive Wi-Fi transmitter and the Wi-Fi receiver. Specifically, the signal strength at the receiver, P_r, can be modeled using Friis path loss [90, 153] as follows,

$$P_r = \left(\frac{P_t G_t}{4\pi d_1^2}\right)\left(\frac{\lambda^2 G_{passive}^2}{4\pi}\frac{|\Delta\Gamma|^2}{4}\alpha_{wifi}\right)\left(\frac{1}{4\pi d_2^2}\frac{\lambda^2 G_r}{4\pi}\right)$$

This equation has three key parts: the term in first parenthesis models signal propagation from the plugged-in device, with an output power P_t and an antenna gain G_t, to a passive Wi-Fi transmitter at a distance d_1 away. The third term, similarly, models the signal propagation from the passive Wi-Fi transmitter to a Wi-Fi receiver with an antenna gain G_r and at a distance d_2 away. Here, λ is the wavelength of the RF signal been transmitted. Finally, the middle parenthesis models the fraction of incident signal from the plugged-in device that is backscattered by a passive Wi-Fi transmitter with an antenna gain $G_{passive}$. $|\Delta\Gamma|^2$ is the backscatter coefficient which is a measure of the ef-

ficiency with which passive Wi-Fi can generate backscatter signals. As described in §5.2.2, this is 1.1 dB for double side band and 2.2 dB for single side band hardware prototypes. Finally, α_{wifi} models the loss in energy due to synthesis of Wi-Fi signals using backscatter. This is around 4.4 dB in double side band and includes half the power lost in the mirror copy generated by backscatter and the losses due to the side lobes as described in §5.2.3. For single side band, ideally this loss is 1.1 dB but in practice we see a loss of around 4 dB due to additional losses due to higher frequency switching in our COTS hardware implementation.

To gain a better intuition, consider the scenario in Fig. 5.5 where we place the plugged-in device and the Wi-Fi receiver separated by 45 feet. We move the passive Wi-Fi transmitter between these devices, along the line connecting them. We set P_t, G_t, G_r and $G_{passive}$ to 30 dBm, 6 dBi, 0 dBi, and 2 dBi respectively. Fig. 5.5 shows the received signal strength, P_r, as we move the passive Wi-Fi transmitter between the plugged-in device and the Wi-Fi receiver. The plots show two key points.

(1) The received signal strength increases as the passive Wi-Fi transmitter gets close to either the Wi-Fi receiver or the plugged-in device. This is because, maximizing the signal strength requires minimizing the product $d_1 d_2$, which is achieved either by reducing the distance d_1 or d_2.

(2) The mid-point between the plugged-in device and Wi-Fi receiver has the lowest strength. Fig. 5.5 shows this mid-point signal strength, as we change the distance between the plugged-in device and Wi-Fi receiver. The plot shows that this decreases with distance between the plugged-in device and the Wi-Fi receiver. As expected, it increases with plugged-in device's transmit power (P_t).

Understanding Deployment Scenarios

1. I want to deploy passive Wi-Fi devices in my home. Where do I place the plugged-in device so as to maximize their range? Fig. 5.6 shows the theoretical signal strength at the Wi-Fi receiver as a function of its distance from the passive Wi-Fi transmitter. We show the results for different distances between the passive Wi-Fi transmitter and the plugged-in device. We set G_t, G_r, $G_{passive}$, P_t to 6 dBi, 0 dBi, 2 dBi, and 30 dBm respectively. The plot shows that, in general, as the distance

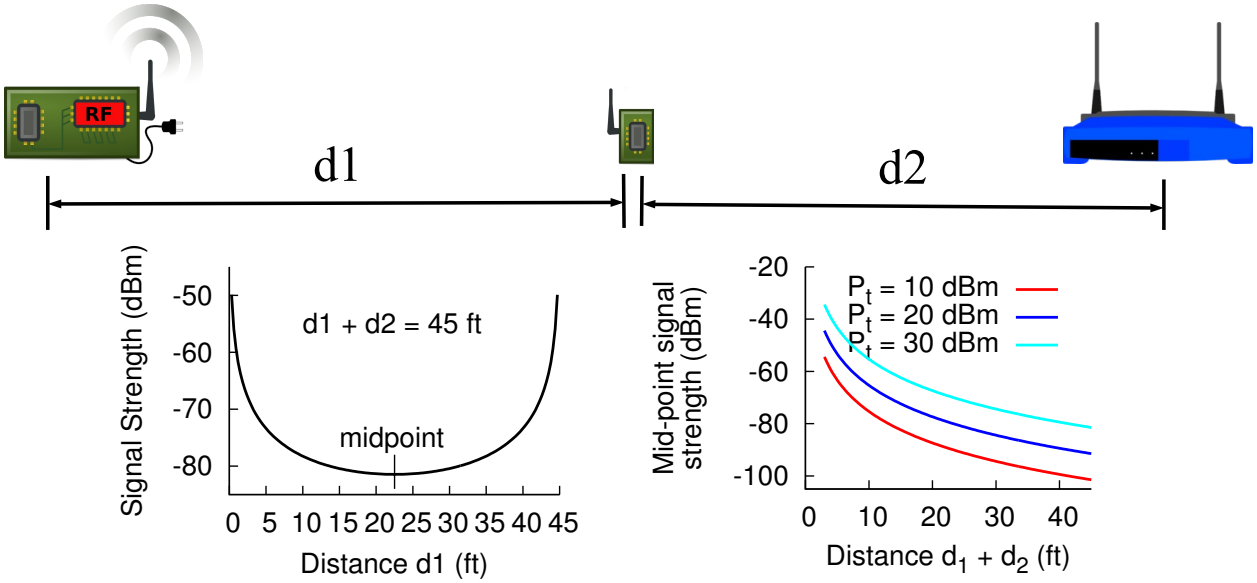

Figure 5.5: **Passive Wi-Fi's analytical received signal strength.** The passive Wi-Fi device moves along the line connecting the Wi-Fi router and plugged-in device.

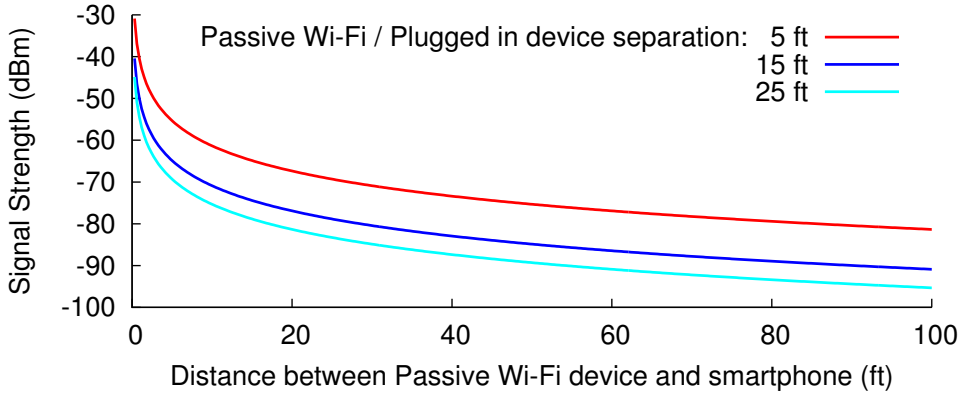

Figure 5.6: **Signal strength versus distance between passive Wi-Fi transmitter and Wi-Fi receiver.**

between the passive Wi-Fi transmitter and Wi-Fi receiver increases, the received signal strength reduces. More importantly, as the distance between the passive Wi-Fi transmitter and plugged-in device decreases, the coverage range increases. This is because, from our analysis, the signal strength can be increased either by reducing the distance between the passive Wi-Fi transmitter and the plugged-in device or that between the passive Wi-Fi transmitter and the Wi-Fi receiver. Since our goal is to maximize range, we should reduce the distance between the passive Wi-Fi transmitter and the plugged-in device. In the presence of multiple passive Wi-Fi devices, this would

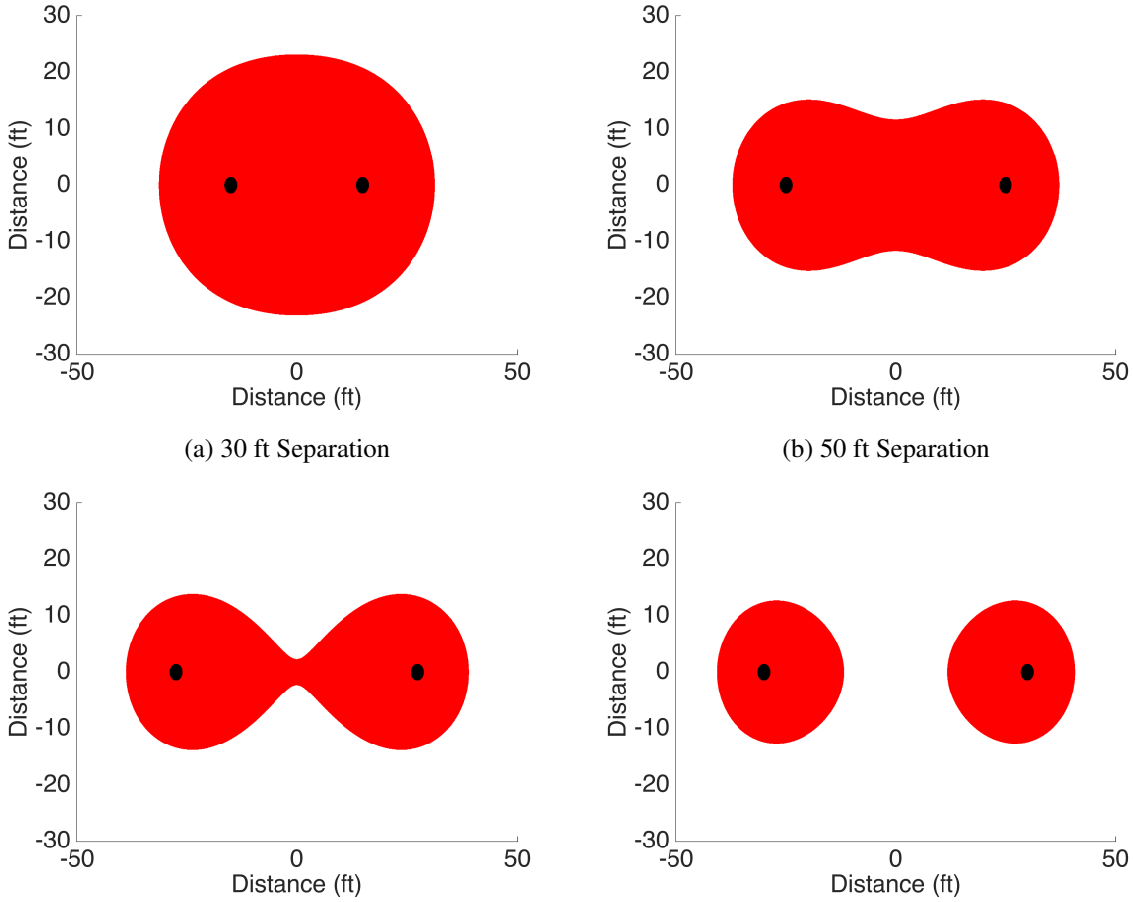

(a) 30 ft Separation

(b) 50 ft Separation

(c) 55 ft Separation

(d) 60 ft Separation

Figure 5.7: **Theoretical coverage maps for different distances between the plugged-in device and the Wi-Fi router.** The black dots denote the positions for these devices. The red region represents points in the 2D space where a passive Wi-Fi transmitter can be located, while ensuring that the signal from it to the Wi-Fi router is at least -85 dBm.

translate to minimizing the worst-case distance between the plugged-in device and all passive Wi-Fi transmitters.

2. Where do I place my Wi-Fi router and the plugged-in device, so that I can have passive Wi-Fi devices work from anywhere in my home? Fig. 5.7 shows the 2D coverage maps for different distances between the plugged-in device and the Wi-Fi router. The red region represents points in the 2D space where a passive Wi-Fi transmitter can be located, while ensuring that the signal from it to the Wi-Fi router is at least -85 dBm. These maps show that the coverage area is a union of two circles centered each at the Wi-Fi router and the plugged-in device. So, as a general rule of thumb,

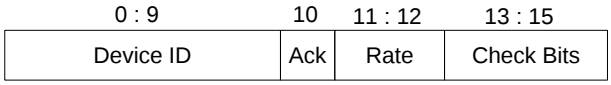

Figure 5.8: **Structure of the signaling packet.**

it is better to deploy the plugged-in device and the Wi-Fi router at either ends of the coverage area. Note however that at very large distances between the plugged-in device and Wi-Fi router (Figs. 5.7 (c) and (d)), we end up getting two islands of coverage. Such large distance deployments are suitable only when the passive Wi-Fi transmitters are going to be close to either the plugged-in device or the Wi-Fi router.

5.3 Passive Wi-Fi Network Stack Design

We first describe how passive Wi-Fi devices share the ISM band. We then address the issue of ACKs and retransmissions and finally, present our protocol to associate passive Wi-Fi devices with the network.

5.3.1 Sharing the ISM band

Wi-Fi uses carrier sense to share the ISM band. This however requires a Wi-Fi receiver that is ON before every transmission. Since Wi-Fi receivers require power-consuming RF components like LNA, frequency synthesizers, mixers and ADCs, this would eliminate the power savings from our design. Instead, we delegate the task of carrier sense to the plugged-in device, which also arbitrates access between multiple passive Wi-Fi devices.

We illustrate this with an example. Say a passive Wi-Fi transmitter wants to sent a packet on channel 6 and the plugged-in device transmits its tone between Wi-Fi channels 1 and 6. Before any of the above transmissions happen, the plugged-in device first uses carrier sense to ensure that there are no ongoing transmissions on any the frequencies including and in between channel 1 and 6.

Once the channels are found free, the plugged-in device sends a packet signaling a specific passive Wi-Fi device to transmit. This signal is sent and decoded using the ultra-low power receiver

described in §5.3.1. The packet starts with an ID unique to each passive Wi-Fi device (see Fig. 5.8). When the passive Wi-Fi device detects its ID, it transmits within a SIFS duration at the end of the signaling packet. The signaling packet is sent at the center of channel 1 and 6 as well as in between them. This prevents other devices in the ISM band from capturing the channel before the passive Wi-Fi device gets to transmit. The packet has 16 bits and adds a fixed overhead of 100 μs for every passive Wi-Fi transmission.

The above description assumes that the plugged-in device knows when to send the signaling packet to each of the passive Wi-Fi devices in the network. To see how this can be achieved let us focus on our target IoT applications. A device sending out beacons is configured to send them at a fixed rate. Temperature sensors, microphones and Wi-Fi cameras (e.g., Dropcam [20]) have a fixed rate at which they generate data. Similarly, motion sensors have an upper bound on the delay they can tolerate. The passive Wi-Fi devices convey this information to the plugged-in device during association (and can update it later using the protocol in §5.3.3). This information is used by the plugged-in device to signal each passive Wi-Fi device in accordance to its desired update rate.

Ultra-low power receiver design

The plugged-in device encodes the bits using ON-OFF keying. We use a passive energy detector with analog components and a comparator to distinguish between the presence and absence of energy. Our design is the same as that used in our prior work [113, 114] and we skip it for brevity. We implement the receiver using off-the-shelf components and it consumes 18 μW, while achieving a bit rate of 160 kbps.

5.3.2 ACKs and Rate Adaptation

ACKs and retransmissions. The plugged-in device listens to the ACKs and conveys this information back to the passive Wi-Fi sensor. Specifically, if the ACK is successfully decoded at the plugged-in device, it sets the ACK bit in the signaling packet shown in Fig. 5.8 to 1 and sends it to the passive Wi-Fi sensor, by piggybacking it during the next period when the sensor is scheduled to

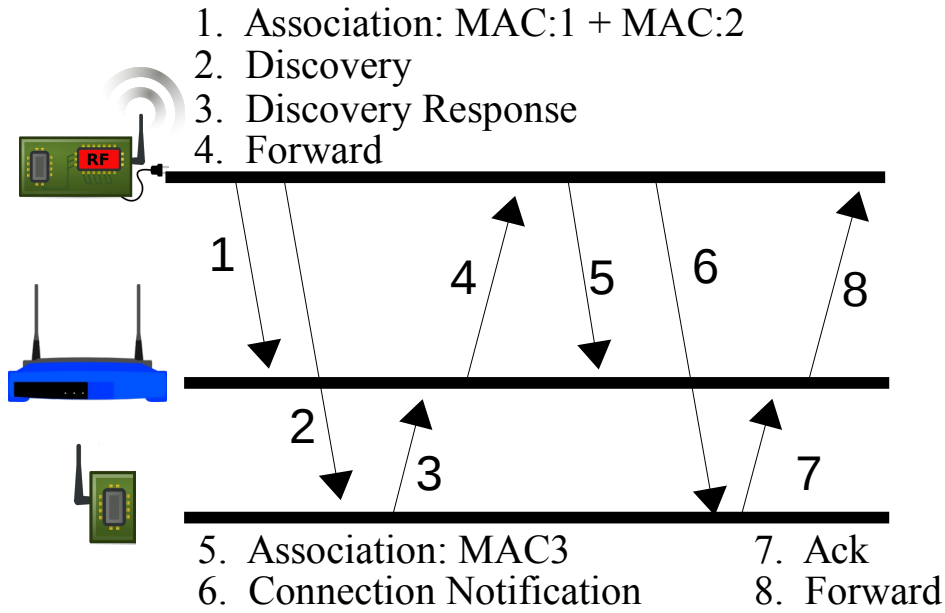

Figure 5.9: **Passive Wi-Fi association procedure.**

transmit. If the ACK is not received at the plugged-in device, it immediately performs carrier sense and sends a signaling packet with the ACK bit set to 0. When the passive Wi-Fi sensor receives this, it retransmits its sensor value. In our implementation, the plugged-in device detects an ACK by detecting energy for a ACK duration at the end of the passive Wi-Fi transmission.

Rate adaptation. Wi-Fi bit rate adaptation algorithms typically use packet loss as a proxy to adapt the transmitter bit rate. In our design, we delegate this function to the plugged-in device. Specifically, the plugged-in device estimates the packet loss rate for each of its associated passive Wi-Fi devices by computing the fraction of successfully acknowledged packets. It then estimates the best 802.11b bit rate and encodes this information in the bit rate field of the signaling packet. Since the plugged-in device knows the bit rate as well as the packet length (from association as described in §5.3.3), it knows how long the transmissions from each of its passive Wi-Fi devices would occupy on the wireless medium. Thus, it stops transmitting its tone at the end of the passive Wi-Fi transmission and listens for the corresponding ACKs.

5.3.3 Network Association

Finally, we describe how the passive Wi-Fi transmitters associate with the plugged-in device as well as with the Wi-Fi router in the network. The key challenge is that since the plugged-in device does not have a full-duplex radio (the lack of which is desirable to make it practical and low cost), there is no direct communication channel from the passive Wi-Fi device to the plugged-in device. Instead, as shown in Fig. 5.9, the plugged-in device associates with the Wi-Fi router with two MAC address (MAC:1 and MAC:2). The plugged-in device then broadcasts a discovery packet using ON-OFF keying modulation that contains these two MAC addresses and starts with a broadcast ID. The new passive Wi-Fi device then transmits a Wi-Fi packet with the source and destination addresses set to MAC:2 and MAC:1; this packet gets routed through the Wi-Fi router to the plugged-in device. The packet payload includes the sensor update rate, packet length, supported bit rates and its MAC address, MAC:3. The plugged-in device spoofs MAC:3 and associates it with the Wi-Fi router. It then picks a unique ID and sends it to the passive Wi-Fi device along with other Wi-Fi network credentials. Finally, the passive Wi-Fi device responds with a Wi-Fi packet with the source and destination addresses set to MAC:3 and MAC:1; this packet gets routed through the Wi-Fi router and confirms association at the plugged-in device.

After association, the passive Wi-Fi transmitter can send Wi-Fi packets to the plugged-in device through the router, and change its parameters including update rate and packet length. Note that the credentials for the spoofed MAC addresses could be sent securely using a manufacturer set secret key shared between the passive Wi-Fi devices and the plugged-in devices. Exploring this in detail however is not in the scope of this work.

5.4 Hardware Implementation

We first describe our implementation of passive Wi-Fi using off-the shelf components on an FPGA platform. We use this to characterize passive Wi-Fi in various deployment scenarios. We then present our IC design which we use to quantify our power consumption.

Off-the-shelf implementation. We implement a passive Wi-Fi prototype using off-the-shelf com-

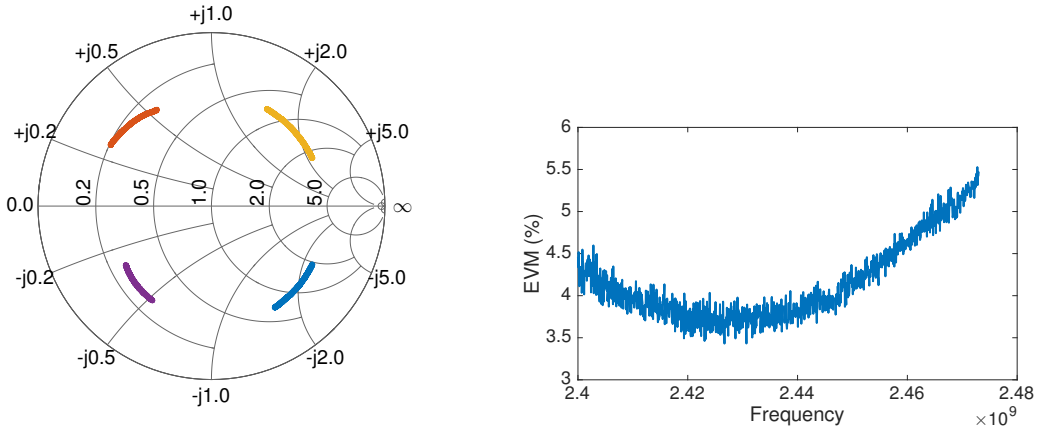

(a) Smith chart for single side band *(b) Error vector for single side band*

Figure 5.10: Smith chart showing the achieved constellation point for single side band backscatter design and the corresponding EVM as a function of frequency.

ponents for backscatter and an FPGA for digital processing. The backscatter modulator consisted of HMC190BMS8 SPDT RF switch network on a 2-layer Rodgers 4350 substrate [26]. The switch for double side band was designed to modulate between open and closed impedance states and had a 1.1 dB loss. For single side band, we used a cascaded SPDT switch network to connected to the four impedance states corresponding to QPSK constellation. The four impedance states corresponding to the QPSK constellation map and the corresponding Error Vector Magnitude (EVM) of the implementation for single side band backscatter implementation is shown in Fig. 5.10. All the required baseband processing including data scrambling, header generation, DSSS/CCK encoding, CRC computation and DBPSK/DQPSK modulation were written in Verilog. The Verilog code was synthesized and programmed on a DE1 Cyclone II FPGA development board by Altera [8]. We implement four different frequency shifts at the passive Wi-Fi device of 12.375, 16.5, 22 and 44 MHz. The digital output of the FPGA was connected to the backscatter switch to generate the Wi-Fi packets from the tone emitted by the plugged-in device. A 2 dBi omnidirectional antenna was used on the passive Wi-Fi device. The plugged-in device was set to transmit an equivalent isotropic radiated power (EIRP) of 30 dBm.

Integrated circuit implementation. CMOS technology scaling has enabled the exponential scaling in power and area for integrated circuits. Wi-Fi chipsets have tried to leverage scaling but with

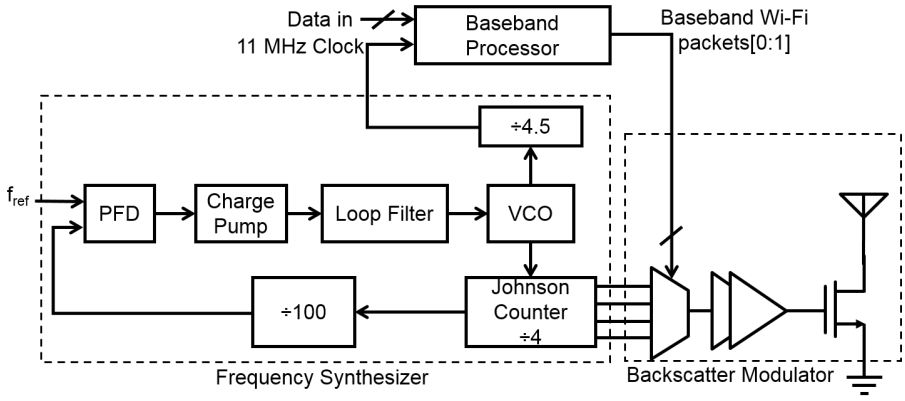

Figure 5.11: **Passive Wi-Fi's IC architecture for double side band backscatter.** The frequency synthesizer generates baseband clock. The baseband processor processes incoming data into Wi-Fi packets and the backscatter modulator performs phase modulation and backscatters using the RF switch.

limited success due to the need for power hungry analog components that do not scale in power and size with CMOS technology. However, baseband Wi-Fi operations are implemented in the digital domain and tend to scale very well with CMOS. For context, Atheros's AR6003 [10] and AR9462 [47] chipsets that were released in 2009 and 2012 use 65 nm CMOS and 55 nm CMOS node implementations respectively. For passive Wi-Fi devices integrated circuit implementation, we chose the 65 nm LP CMOS node by TSMC, which gives us power savings of baseband processing and ensures a fair comparison with current industry standards. The IC architecture of the passive Wi-Fi device for double side band backscatter is shown in Fig. 5.11 and has three main components:

Baseband frequency synthesizer. It generates the 11 MHz clock required for baseband processing as well as four phases at 12.375 MHz offsets required for DBPSK and DQPSK. We phase synchronize the 11 MHz and 12.375 MHz clocks to avoid glitches during phase modulation. We used an integer N charge pump and ring oscillator-based PLL to generate 49.5 MHz clock from a 12.375 kHz reference. The 49.5 MHz clock is fed to a quadrature Johnson counter to generate the four phases with the required timing offsets (corresponding to 0, $\frac{\pi}{2}$, π and $\frac{3\pi}{2}$ phases). The same 49.5 MHz carrier is divided by 4.5 to generate the 11 MHz baseband clock.

Baseband processor. It takes the payload bits as input and generates baseband 802.11b Wi-Fi

Table 5.1: **Passive Wi-Fi's IC Power Consumption**

	1 Mbps	**11 Mbps**
Baseband Frequency Synthesizer	5.6 μW	5.6 μW
Baseband Processor	5.0 μW	48 μW
Backscatter Modulator	3.9 μW	5.6 μW
Total Power	14.5 μW	59.2 μW

packet. We used the Verilog code that was verified on the FPGA and use the Design Compiler by Synopsis to generate the transistor level implementation of the baseband processor [54].

Backscatter modulator. It mixes the baseband data to generate DBPSK and DQPSK and drives the switch to backscatter the incident tone signal. The baseband data are the select inputs to a 2-bit multiplexer which switches between the four phases of the 12.375 MHz clock to generate the phase modulated data. The multiplexer output is buffered and used to drive the RF switch, which toggles the antenna between open and short impedance state.

Table 5.1 shows the power consumption of our design at 1 Mbps and 11 Mbps which was computed using the Cadence spectre and Synopsis Design Complier toolkits [14, 54]. Passive Wi-Fi's double side band IC implementation for 1 Mbps and 11 Mbps consumes a total of 14.5 and 59.2 μW of power respectively. The digital frequency synthesizer is clocked for DQPSK and consumes a fixed power for all data rates. The power consumption of the baseband processor that generates the 802.11b packets scales with the data rate and consumes 30% and 80% of total power for 1 and 11 Mbps respectively. The backscatter modulator consumes the rest of the power for performing phase modulation and running the switch.

5.5 Evaluation

We first evaluate passive Wi-Fi s physical layer performance and then evaluate its network performance.

5.5.1 Efficacy of single sideband backscatter

We compare the single sideband backscatter design with the double sideband backscatter hardware design [115] that create a mirror copy on the unintended side of the single tone. We use a USRP to transmit a single tone at 2.5085 GHz and set the backscatter devices to generate Wi-Fi packets at 2.462 GHz. In the case of prior double sideband backscatter designs, this creates a mirror copy on Wi-Fi channel 6. We configure a Wi-Fi transmitter-receiver pair on channel 6 and run an iperf connection using TCP with the default Wi-Fi rate adaptation algorithm. We use a Linksys WRT45g Wi-Fi AP as the Wi-Fi transmitter and a Nexus 4 smartphone as the Wi-Fi receiver and separate them by 10 feet. We set the backscatter device to generate 2 Mbps Wi-Fi packet with a 32-byte payload, at a distance of 2 feet from the Wi-Fi receiver.

Fig. 5.12 shows the iperf throughput in the presence of single and double-sideband backscatter hardware. The x-axis shows the rate at which Wi-Fi packets are generated using backscatter. For comparison, we show the baseline throughput in the absence of backscatter. The plot shows that,

- When the backscattering device generates a small number of packets (50 pkts/s), it has negligible impact on the concurrent iperf flow. This is true with both single-sideband and double sideband backscatter designs. This is expected because the backscattered packets are small and are transmitted at a very low rate that they do not affect the throughput of concurrent Wi-Fi connections.

- When we backscatter Wi-Fi packets at a higher rate, the iperf throughput reduces with prior double sideband backscatter hardware. This is because it creates a mirror copy on Wi-Fi channel 6 that reduces the throughput of any other Wi-Fi connection. However, we see negligible impact on the iperf throughput with our single-sideband hardware. This demonstrates that our single-sideband backscatter design can double spectral efficiency.

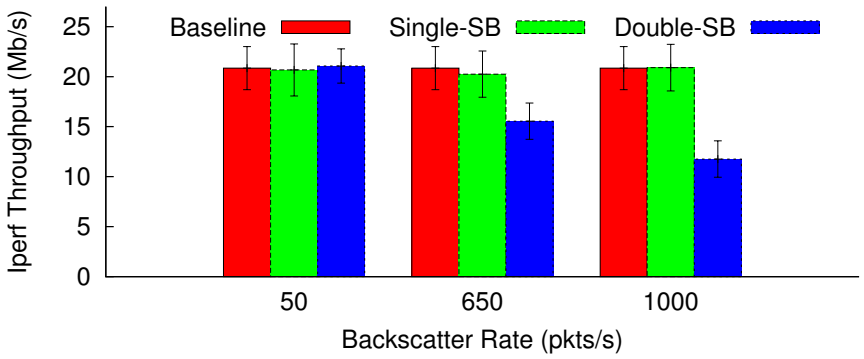

Figure 5.12: **Efficacy of single sideband backscatter.** We compare our design with double sideband backscatter designs on the throughput of an iperf flow on a concurrent Wi-Fi transmitter-receiver pair. Baseline is the throughput in the absence of any backscatter device.

5.5.2 Single side band evaluation

In the following sections we will evaluate the operating range of passive Wi-Fi in various deployment scenarios. The operating range of passive Wi-Fi is a function of the power of scattered Wi-Fi packet. We have introduced both single side band and double side band designs for generating Wi-Fi packets. In our hardware implementation, the single side band and double side band designs have similar output power. In theory, the single side band should have 3 dB higher power. However, 1.1 dB is lost in the additional switch required for QPSK constellation and we lose additional power due to switching attributed to complex sub-carrier modulation. To verify the spectrum and behavior of single side band backscatter, we connected the single side band backscatter switch, the transmitter and an RF scope using a circular and measured the spectrum. Fig. 5.13 shows the spectrum of single side band Wi-Fi packet and the measured in-band power in the Wi-Fi Channel 1 is within 0.5 dB of the power in double side band design. Since, the operating range of single side band and double side band systems would be similar, in subsequent sections, we will use double side band design to characterize the operating range of passive Wi-Fi.

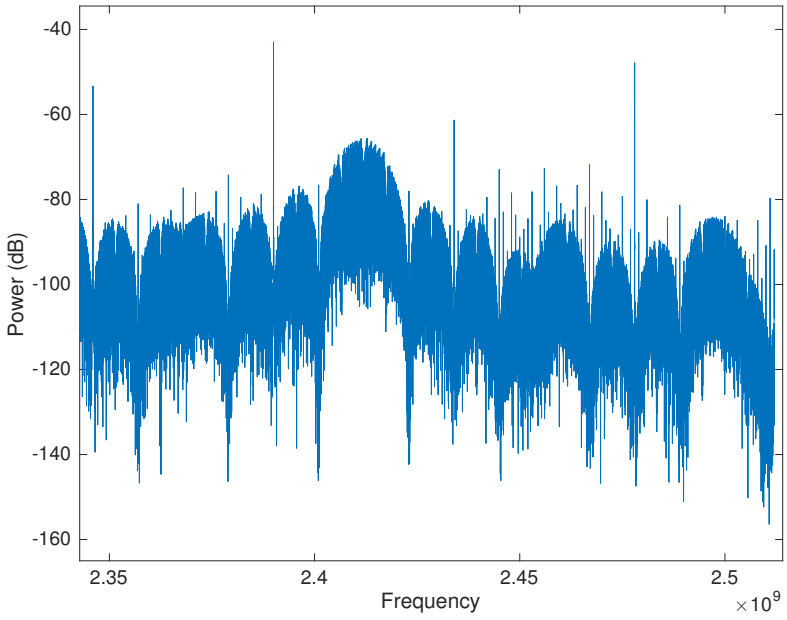

Figure 5.13: **Measured spectrum of single sideband backscatter.** We use a high frequency RF scope to measure the spectrum of the single side band backscattered Wi-Fi.

5.5.3 Physical Layer Performance

We first evaluate the range in various scenarios. We then evaluate the effect of the frequency shift used in our system on the packet loss rate. Finally, we present results for all four 802.11b bit rates.

RSSI in Line-of-sight scenarios

We run experiments in two line-of-sight scenarios.

Deployment scenario 1. We fix the distance between the passive Wi-Fi device and the plugged-in device. We then move the Wi-Fi receiver away from the passive Wi-Fi device and measure the RSSI of the passive Wi-Fi transmissions as seen by the receiver. We run the experiments in the CSE atrium where the maximum distance possible when the passive Wi-Fi device and Wi-Fi receiver were placed on either end was around 100 feet. In our experiments, we set the passive Wi-Fi device to generate 802.11b beacon packets at 1 Mbps. These packets have a payload of 68 bytes where

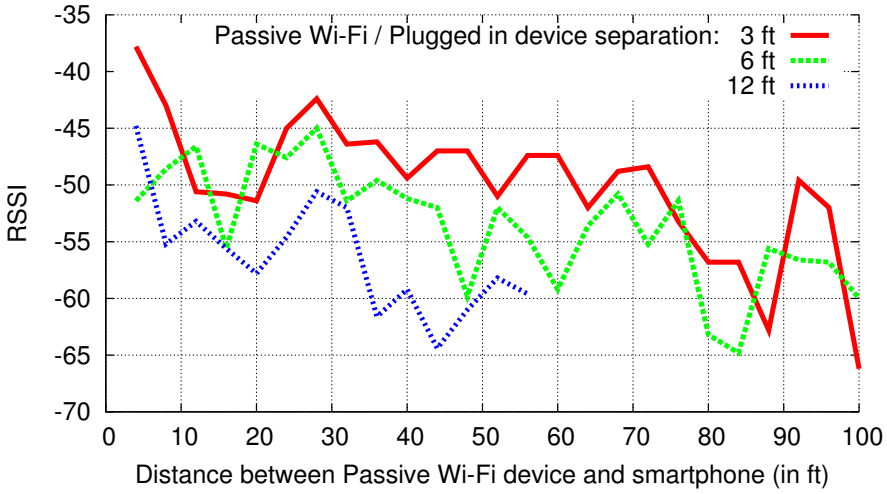

Figure 5.14: **RSSI in deployment scenario 1.** We move the phone away from the passive Wi-Fi device.

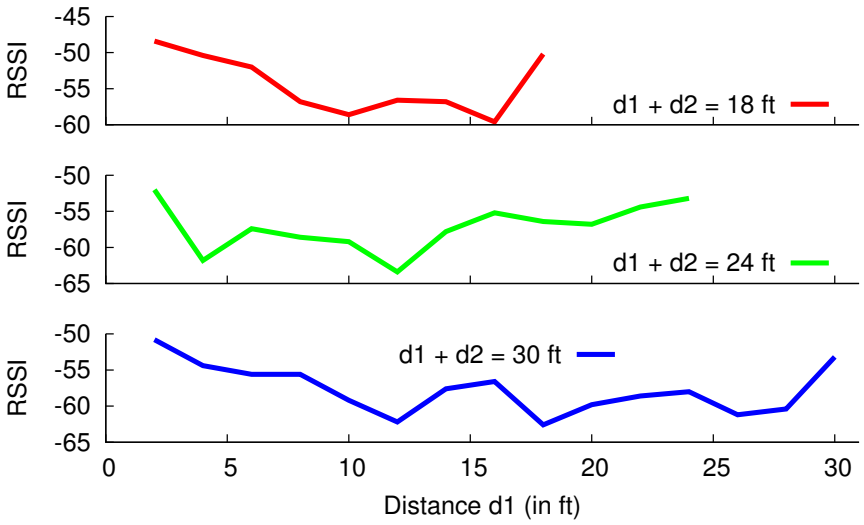

Figure 5.15: **RSSI in deployment scenario 2.** d_1 is the distance between the passive Wi-Fi transmitter and plugged-in device. d_2 is the distance between the passive Wi-Fi device and Wi-Fi receiver. The passive Wi-Fi device moves alone the line joining the other two devices.

the SSID is set to *WiLab_0000* and are transmitted every 15 ms. We set the plugged-in device to transmit its tone 12.375 MHz from the center of Wi-Fi channel 1 between channel 1 and 6. We use an HTC One (M7) phone as our Wi-Fi receiver. Since the passive Wi-Fi device is transmitting Wi-Fi beacons, it appears as a Wi-Fi AP at the smartphone. To measure the RSSI values of these

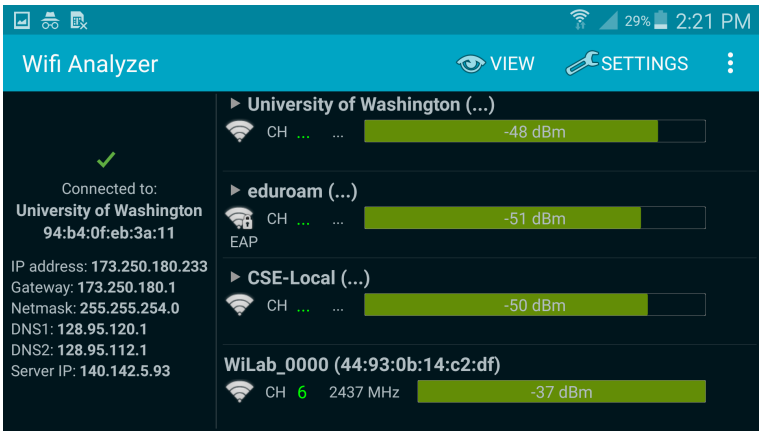

Figure 5.16: **Snapshot of the Wi-Fi analyzer app.** *WiLabAP_0000* corresponds to the beacons transmitted by a passive Wi-Fi device.

packets, we use a third party Android app called Wifi Analyzer [9]. The app provides the RSSI value computed by the Wi-Fi chipset at the smartphone using all the Wi-Fi beacon transmissions on the wireless medium as shown in Fig. 5.16.

In each experiment, we hold the smartphone in our hand and measure the reported RSSI values as we walk away from the passive Wi-Fi device. The measurements are taken at increments of 4 feet. Fig. 5.14 plots the results for three different values of the distance between the passive Wi-Fi transmitter and the plugged-in device. The x-axis plots the distance between the passive Wi-Fi transmitter and the Wi-Fi receiver while the y-axis plots the reported RSSI values. The plots show that as expected, the RSSI values reduce as the phone moves away from the passive Wi-Fi device. Further, as predicted by our analysis in §5.2.5, the range of our passive Wi-Fi transmissions reduce with the distance between the passive Wi-Fi transmitter and the plugged-in device. When the separation between the passive Wi-Fi transmitter and the plugged-in device is 3 or 6 feet, the range of the passive Wi-Fi transmissions spans the entire length of the CSE atrium. The range is around 55 feet when this separation is 12 feet. This reduced range is due to a combination of multipath and a weak backscatter signal.

Deployment scenario 2. Next we place the plugged-in device and the Wi-Fi receiver at a distance $d_1 + d_2$. We move the passive Wi-Fi transmitter along the line connecting these two devices. As above the passive Wi-Fi transmitter is set to generate 802.11b beacon packets at 1 Mbps and the

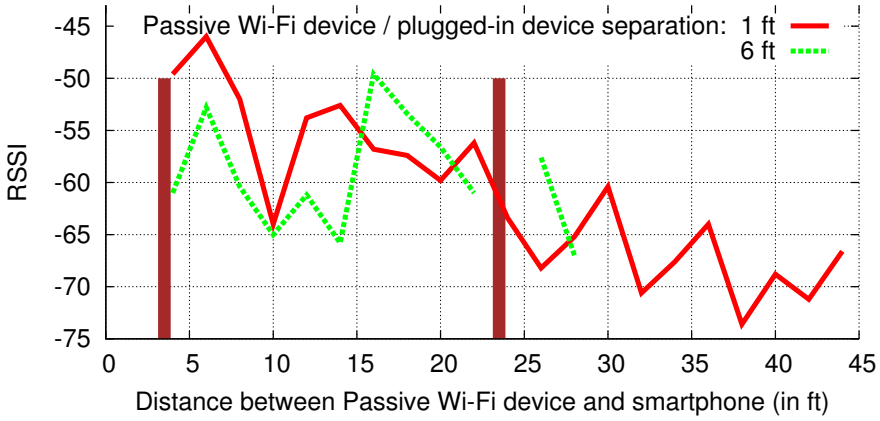

Figure 5.17: **RSSI in deployment scenario 1 in the presence of walls.** The brown blocks show the wall positions.

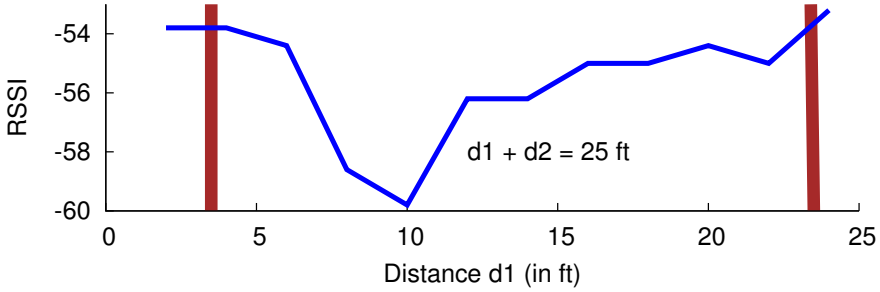

Figure 5.18: **RSSI in deployment scenario 2 in the presence of walls.** The brown blocks denote the position of the walls. d_1 is the distance between the passive Wi-Fi transmitter and plugged-in device. d_2 is the distance between the passive Wi-Fi device and Wi-Fi receiver.

plugged-in device transmits its tone at 12.375 MHz from the center of Wi-Fi channel 1. We collect the RSSI values from a HTC One (M7). Fig. 5.15 plots the results for three different values of the distance between the plugged-in device and the Wi-Fi receiver ($d_1 + d_2$). Each point on the x-axis denotes the distance between the passive Wi-Fi device and the plugged-in device (d_1). The plots show that the RSSI values are the highest when the passive Wi-Fi transmitter is either close to the Wi-Fi receiver or the plugged-in device. Further, the RSSI values are lower at the mid point between the two devices, confirming our theoretical analysis.

RSSI in Through-the-Wall Scenarios

We rerun experiments in the above deployment scenarios but now in the presence of walls. In the first deployment, we place the passive Wi-Fi device and the plugged-in devices at distances of 1

and 6 feet from each other. As the Wi-Fi receiver moves away from the passive Wi-Fi device, it is separated by multiple double sheet-rock (plus insulation) walls with a thickness of approximately 5.7 inches. As before, we use an HTC One (M7) phone as our Wi-Fi receiver and set the plugged-in device to transmit with a 12.375 MHz frequency offset from channel 1. The passive Wi-Fi device periodically transmits Wi-Fi beacons at 1 Mbps and we measure the RSSI values as reported by the Wi-Fi receiver. Fig. 5.17 shows that the range is now around 28 feet when the distance between the passive Wi-Fi device and the plugged-in device is 6 feet. This is expected because the signals get attenuated by two walls before arriving at the Wi-Fi receiver.

In the second deployment, we fix the location of the plugged-in device in the first room and place the Wi-Fi receiver in the third room at a distance of 25 feet. We then move the passive Wi-Fi device along the line connecting the above two devices and measure the RSSI reported by the Wi-Fi receiver. Fig. 5.18 plots the RSSI results and show that they follow a similar trend as before and work even in the presence of attenuation from walls.

Effect of different frequency shifts

We evaluate how different frequency shift values affect passive Wi-Fi performance. To do this, we place the passive Wi-Fi transmitter and plugged-in device 6 feet from each other. We move a Wi-Fi receiver away from the passive Wi-Fi device in a 50 foot long space. The passive Wi-Fi device transmits 1 Mbps Wi-Fi packets with a payload of 512 bytes on channel 1. We use the Intel 5350 chipset as a Wi-Fi receiver which runs tshark to log all the packets that are successfully decoded by it. The passive Wi-Fi transmitter consecutively transmits 200 unique sequence numbers in a loop using which we compute the packet error rate at the Wi-Fi receiver. We repeat these experiments for three different frequency shifts of 12.375, 16.5 and 44 MHz.

Fig. 5.19 plots the PER observed at the Wi-Fi receiver as a function of distance between the passive Wi-Fi transmitter and the Wi-Fi receiver. The figure shows that the PER is consistently around 20% when we use frequency shifts of 44 and 16.5 MHz. For comparison, we measured the PER for a conventional Wi-Fi transmitter placed 10 feet away and observed similar PER values. The interesting observation however is that when the frequency shift is 12.375 MHz, we see a large

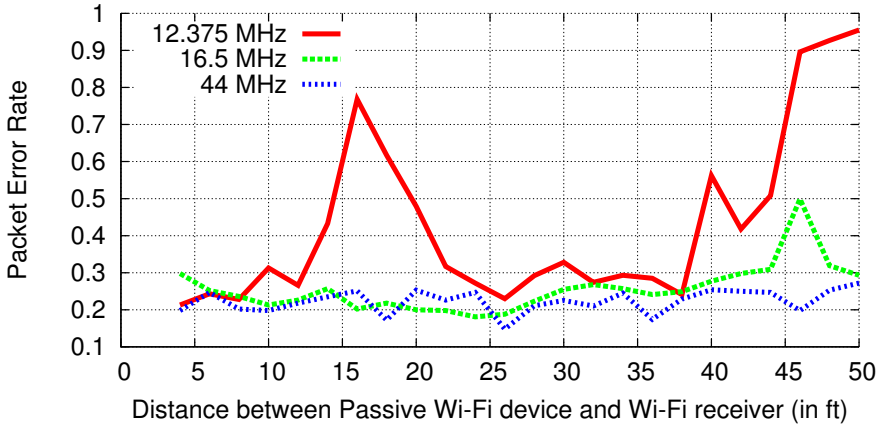

Figure 5.19: **Effect of different frequency shifts.** The PERs are very stable with 16.5 MHz and 44 MHz offsets.

variation in the PER as the location of the Wi-Fi receiver changes. This is because of two related reasons. First, when the frequency shift is small, the tone from the plugged-in device is very close to the desired Wi-Fi channel. Second, because of multipath, different locations see different signal strength differences between the passive Wi-Fi device and the out-of-band interference from the plugged-in device. When the frequency shift is small, this out-of-band interference can still be significant in certain locations to create losses. We note that while a 44 MHz shift is too high to be within the ISM band, a 16.5 MHz shift has PERs that are stable across locations and yet is small to be within the ISM band while generating Wi-Fi packets on channels 1, 6, and 11.

Higher 802.11b bit rates

Finally, we show that passive Wi-Fi can generate all four 802.11b bit rates. We place the passive Wi-Fi device and plugged-in device 6 feet from each other. We change the location of the Wi-Fi receiver to five different spots in a 15×24 ft office room. The plugged-in device is set to use a 12.375 MHz offset. For each Wi-Fi receiver location, the passive Wi-Fi device transmits 802.11b packets at 1, 2, 5.5 and 11 Mbps on channel 1. For each bit rate, the passive Wi-Fi transmitter is configured to send 200 packets with a 512-byte payload with different sequence numbers. The Wi-Fi receiver is configured to compute the effective PHY goodput achieved by multiplying the transmitted Wi-Fi bit rate with the fraction of packets that are decoded. We use an Intel 5350

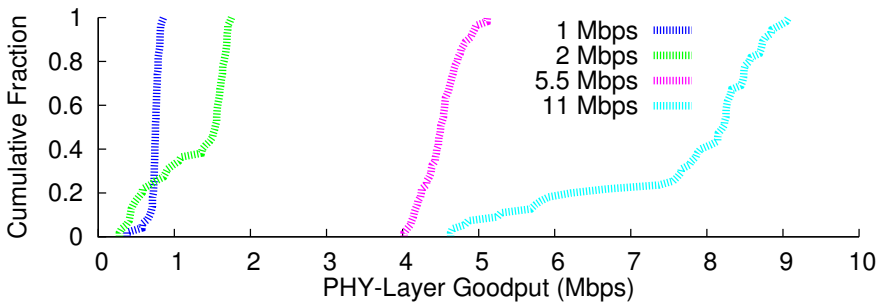

Figure 5.20: **All 802.11b bit rates.** Our design can generate 802.11b transmissions across all four bit rates.

chipset as our Wi-Fi receiver as before. Fig. 5.20 plots a CDF of the PHY-layer goodput across the five locations. The plots show that our design can indeed be used to generate 802.11b transmissions across all four bit rates.

5.5.4 Passive Wi-Fi Network Performance

As described in §5.3.1 to coexist in the ISM band, the plugged-in device first performs carrier sense and then signals the passive Wi-Fi device to transmit. In this section, we first evaluate how well the signaling mechanism works. We then describe how our overall carrier sense mechanism works in the presence of other Wi-Fi devices.

Evaluating the signaling mechanism

The plugged-in device transmits a packet with a 10-bit ID that is unique to each passive Wi-Fi device. We evaluate two aspects: (1) the probability with which the signal from the plugged-in device trigger transmissions from the correct passive Wi-Fi device and (2) the probability that it would trigger the wrong passive Wi-Fi device. To evaluate this, we consider the worst-case scenario: two devices that have IDs that differ by just one bit. We set the plugged-in device to transmit the signaling packet with the ID of the first device. We move the two passive Wi-Fi devices away from the plugged-in device. At each distance value, the plugged-in device is configured to transmit the signaling packet for a total of 1890 times. The passive Wi-Fi devices use an envelope detector to correlate for their specific ID. We compute the fraction of the 1890 signaling packets

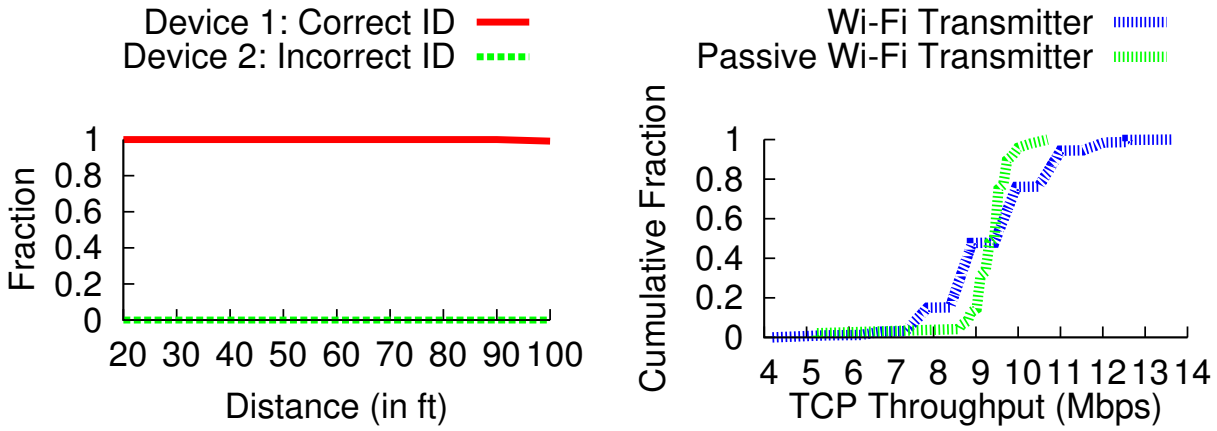

(a) ID Detection *(b) Wi-Fi Coexistence*

Figure 5.21: **Passive Wi-Fi network performance.**

that are decoded and match the ID of the passive Wi-Fi device. We run these experiments in the UW CSE atrium for increasing distances from the plugged-in device. Fig. 5.21(a) show the fraction of signaling packets that match the ID of the two passive Wi-Fi devices as a function of the distance from the plugged-in device. The plot shows that neither device incorrectly decodes the ID. This is because our receiver builds on our prior work [113, 114, 131, 187] and has gone through multiple iterations to improve its reliability.

Evaluating passive Wi-Fi s carrier sense

The plugged-in device performs carrier sense and signals a specific passive Wi-Fi device to transmit. The signaling mechanism adds a constant overhead of 100 μs to each passive Wi-Fi transmission. To compare how our mechanism compares to standard Wi-Fi, we compare the performance of a concurrent Wi-Fi transmitter-receiver pair in the presence of a passive Wi-Fi transmitter with that of a traditional Wi-Fi transmitter. We use two Intel 5350 Wi-Fi chipsets to transmit and receive Wi-Fi packets using iperf. The devices use the chipsets default bit rate adaptation. We run experiments in two scenarios: 1) we use a Ralink RT2070 Wi-Fi chipset to transmit Wi-Fi packets at 1 Mbps every 15 ms and 2) we set our passive Wi-Fi transmitter to transmit its packet every 15 ms at 1 Mbps using our carrier sense mechanism. We measure the throughput achieved by a concur-

rent Wi-Fi transmitter-receiver pair in the presence of these two devices. Fig. 5.21(b) plots the TCP throughput and shows that passive Wi-Fi has a similar impact on the ongoing flow as a traditional Wi-Fi transmitter. This is because, passive Wi-Fi adds only a small fixed 100 μs overhead. This small overhead is however overshadowed by transient changes in network conditions.

5.5.5 Applications

To understand how passive Wi-Fi can be used to reduce the power consumption of sensors with wireless connectivity, we analyze the power budgets of two kinds of applications. First we will consider low latency sensing platforms such as microphones and cameras that continuously transmits sensor information. We then analyze duty-cycled sensors that sporadically transmit data.

1) A low power microphone consumes 17 μW [6] and an ADC digitizing the microphone output would consume an additional 33 μW [6] resulting in a total power consumption of 50 μW for the sensing subsystem. If we use an IoT Wi-Fi chipset by Gainspan or TI to continuously transmit audio, the active Wi-Fi transmitter will consume 670 mW [24, 57]. This results in a total power budget of 670.05 mW which is dominated by the Wi-Fi chipset. However, if we substitute the active Wi-Fi chipset with passive Wi-Fi operating at 1 Mbps, the power budget of the system drops to 65 μW, a 1000x reduction in the total power budget of a wireless microphone.

2) A low power camera like OV7690 operating at VGA resolution and capturing one image per second consumes an average of 10 mW [43]. The camera outputs raw data at 2.45 Mbps which can be transferred wirelessly without the need for power hungry on-board compression. Using an IoT Wi-Fi chipset from Gainspan or TI, brings the total power consumption of the system to 680 mW. If we substitute an active Wi-Fi chipset which consumes 670 mW of power with 11 Mbps passive Wi-Fi, we can improve the battery life of Wi-Fi video camera by at least 50x [24, 57].

3) Finally, let us analyze duty cycled sensors such as iBeacons [27] and home proximity sensors [46] which periodically transmit data using Bluetooth Low Energy and ZigBee protocols respectively. iBeacons and proximity sensors typically transmit beacons/data packets a rate of 100 ms to 900 ms and last for 3 months to 3 years respectively on a coin cell battery [27]. If we replace the

active BLE/ZigBee transmitter which consumes 35 mW [55] in transmit mode with passive Wi-Fi consuming 15 μW, the battery life can be extended well beyond 10 years, which is the lifetime of a typical battery. Furthermore, the battery can be supplemented with harvested energy from RF signals such as Wi-Fi using systems such as power over Wi-Fi [196, 197], to develop beacons and sensors which would never require battery replacement.

5.6 Related Work

RFID systems. RFID tags backscatter the signal back to a dedicated 900 MHz RFID reader. The use of backscatter as a general communication mechanism, however, has been limited to RFID systems for two key reasons. First, to decode the weak backscattered signals, the reader eliminates the strong signal from the reader using full-duplex radios [116, 211]. This requires expensive components such as circulators and highly linear analog RF front end at the reader which contribute to the high cost of typical RFID readers (1000s of dollars). In contrast, chipsets for conventional radio technologies such as Wi-Fi do not require the specialized components, can be fully integrated in silicon and hence, are orders of magnitude less expensive. Second, enabling backscatter communication with existing devices such as smartphones, routers and laptop, requires a complete hardware change to their chipsets and incorporating a dedicated fully duplex radio; this is a high bar that has limited the adoption of backscatter below RFID systems.

Sub-carrier modulation in backscatter systems. Sub-carrier modulation is a widely used technique in backscatter systems. Commercial RFID readers use the EPC Gen2 protocol, and Miller based sub-carrier modulation with 2-ASK and 2-PSK is part of the EPC Gen2 RFID specification [92]. Additionally, [119, 120] and [90] have demonstrated the use of sub-carrier modulation for 2-FSK backscatter communication. In this work, we build on prior approaches of sub-carrier modulation to implement for the fist time, DBPSK and DQPSK modulation with DSSS/CCK encoding and synthesize 802.11b Wi-Fi packets. We also characterize the in-band and out of band interference as a function of the offset frequency. Finally, we introduce single side band backscatter which eliminates the mirror copy of backscatter data and increases the spectral efficiency of

backscatter.

Bi-static backscatter systems. Our work is also related to bi-static backscatter systems. In such systems, the transmitter of the carrier signal and the receiver of the backscattered signal are physically separated [90,118–120,183]. The use of this configuration, reduces out of band interference at the receiver, thereby eliminating the need for full-duplex radio and reducing the cost and complexity of the receiver. [119,120] uses a bi-static configuration to receive 2-FSK modulated backscatter signals on software defined radios. In [90], authors used a similar deployment but instead receive 2-FSK BLE packets synthesized using backscatter on Bluetooth radios. In this work, we use the bi-static configuration with interference resilient DSSS/CCK encoded DBPSK and DQPSK modulated signals which can be reliably received on Wi-Fi radios. We leverage the fact that Wi-Fi receivers have out of band interference rejection components and work even in the presence of interference in the adjacent band that is 35 dB stronger [28]. We analyze the operating range and the performance of passive Wi-Fi in various deployment scenarios. We also design a network-layer stack design that enable us to operate in the presence of multiple passive Wi-Fi transmissions as well as with existing devices in the ISM band.

Wi-Fi and ambient backscatter systems. In ambient and Wi-Fi backscatter [113, 131, 165], battery-free devices communicate with each other by backscattering ambient signals such as TV and Wi-Fi transmissions. The basic difference between these designs and passive Wi-Fi is that Wi-Fi backscatter systems create an additional narrowband data stream to ride on top of existing Wi-Fi signals. In contrast, passive Wi-Fi aims to use backscatter to generate 802.11b transmissions that can be decoded by billions of existing devices with Wi-Fi chipsets.

In particular, Wi-Fi backscatter [113] demonstrated that existing Wi-Fi chipsets can decode backscattered information from a tag using changes to the per-packet CSI/RSSI values at 1 kbps bitrates and a 2 m range. [168] improved the rate of this communication using a full-duplex radio to cancel the high-power Wi-Fi transmissions from the reader and decode the weak phase-modulated narrowband backscattered signal at the reader. This has allowed them to achieve data rates of up to 5 Mbps at a range of 1 m and 1 Mbps at a range of 5 m. A recent news release [37] claims to

achieve 330 Mbps Wi-Fi backscatter communication at 2.5 m using a custom IC that implements a full-duplex radio. The challenge with these full-duplex designs is that they have the same problem as conventional RFID designs— they require a custom full-duplex radio to be incorporated at the receiver and hence the backscattered signals cannot be decoded on any of the existing Wi-Fi devices.

Finally, [90] builds on Wi-Fi backscatter and uses subcarrier modulation to create 2 MHz narrowband 2-FSK signals. Instead, we create 22 MHz DSSS/CCK transmissions using backscatter and enable Wi-Fi transmissions at 10000x lower power than existing Wi-Fi systems. We also present a network-layer stack design that enable us to operate in the presence of multiple passive Wi-Fi transmissions as well as with existing devices in the ISM band.

Duty-cycled radios. Over the past decade, there has been significant work in the academic community to develop ultra-low power radios [176]. The key idea in these systems is to design a custom low power radio transmitter front end and use a wakeup receiver to duty cycle the transmitter and reduce the average transceiver power consumption. The power consumption of such transmitters at sub-milliwatt output power is in the order of 100 μW [162, 213] to few mWs [69, 161, 175]. Further, such radios use custom protocols supporting 10-100 kbps data rates that require deployment of special purpose receivers and hardware. In contrast, passive Wi-Fi demonstrates that we can generate Wi-Fi transmissions at tens of microwatts of power; given the ubiquity of Wi-Fi, this significantly lowers the bar for adoption. We also note that, the duty cycle operation is orthogonal to passive Wi-Fi transmissions and can be used to further reduce the overall power consumption of a system employing passive Wi-Fi.

Low power Wi-Fi transceivers. The Wi-Fi industry has designed Wi-Fi chipsets for IoT applications including QUALCOMM QCA4002 and QCA4004 [48]. These designs reduce the power consumption by decreasing the transmit power by up to a half when in proximity of another device. They also optimize the power consumption of their sleep mode to be less than 1 mW. Additionally, Gainspan and Texas Instruments also sell Wi-Fi chipsets for IoT applications which incorporate standby/hibernate modes with power consumption of less than 20 μ W and can switch to active

mode within 10's of milliseconds [24, 57]. However, the power consumption during active transmission is around of 600 mW [24, 57] which is orders of magnitude higher than passive Wi-Fi. Efforts including Intels Moores radio [17] design digital versions for RF components such as frequency synthesizers. This is targeted at reducing the cost and size of the RF chipset rather than its power consumption — a digital Wi-Fi frequency synthesizer consumes 10-50 mW [17, 99] which is similar in power consumption to its analog counterpart.

Finally, recent low power Wi-Fi receiver designs use techniques like dynamic voltage and frequency scaling [134] and compressive sensing [135]. In particular, SloMo [135] leverages the sparsity inherent to 802.11b DSSS signals using compressive sensing to operate the radio at a lower clock rate. Enfold [134] extends this to work with OFDM modulation. Our work on enabling ultra-low power Wi-Fi transmissions is complimentary to this work and can in principle be integrated together.

5.7 Conclusion

Wi-Fi has traditionally been considered a power-consuming communication system and has not been widely adopting in the sensor network and IoT space. We demonstrated for the first time that 802.11b transmissions can be generated using backscatter communication, while consuming 3–4 orders of magnitude lower power than existing Wi-Fi chipsets. Passive Wi-Fi transmissions can be decoded on any Wi-Fi device including routers, mobile phones and tablets. We built prototype hardwares and evaluated our design in various deployment scenarios.

Chapter 6

HYBRID ANALOG-DIGITAL BACKSCATTER: A NEW APPROACH FOR BATTERY-FREE SENSING

6.1 Introduction

Radio Frequency Identification (RFID) systems are widely used for identification and tracking in supply chain management, access management and transit systems. The RFID industry led this effort and developed long range and inexpensive tags, standardized protocols (EPC Gen2) and RFID reader infrastructure. Due to the progress made in this field, the sensing community has started to investigate RFID technology to develop wireless sensing platforms and solutions [185]. Traditionally, wireless sensors have batteries which are bulky, expensive, short lived and need frequent maintenance. However, with passive RFID, devices can now harvest power from the electromagnetic waves transmitted by the reader and use backscatter to send information to the reader thereby eliminating batteries. Due to their small form factor and battery free operation, passive RFID sensors have longer lifetime and can be widely distributed and embedded virtually everywhere.

6.1.1 Digital backscatter sensing

Digital backscatter is the conventional approach to sensing. Sensor output is sampled using an analog-to-digital converter (ADC) and the sensor values are wirelessly transmitted as digital packets to the reader. A digital temperature sensor using a custom integrated circuit operating in UHF frequency band has been reported in literature [164]. An alternative approach to custom IC, is the use of programmable platforms such as the Wireless Identification and Sensing Platform (WISP) which enable flexibility in the choice of sensing applications. Using WISP, a multitude of sensing applications such as accelerometer, light sensor, temperature sensor and strain gauge have been

reported [187]. Although digital backscatter provides noise immunity and enables the use of existing protocols (and features such as addressability and control) and reader infrastructure, the typical power requirements of active sensors and ADCs are higher than what is available from the reader at long distances. This limits the operating range and/or results in low duty cycles for the digital backscatter platform. Finally, existing protocols (such as EPC Gen2) are optimized for identification and are not suited for high data rate streaming of sensor data [185].

6.1.2 Analog backscatter sensing

Analog backscatter is an alternative sensing approach where the sensing quantity (e.g. temperature) directly modifies some characteristics (such as quality factor or resonant frequency) of a resonating structure or an antenna. For example, Theremin's cavity resonator (also known as the Great Seal Bug [15]) uses capacitive changes in a flexible metallic membrane to detect sound waves. In [95] electromagnetic and magneto mechanical resonance is to develop position and force sensors. Recent work in RFID sensing has focused on utilizing the antenna of commercial RFID tags as a sensor. RFID tag antennas coated with a sensing film or when brought in proximity of metals/dielectric result in changes in the read rate, operating range and turn-on power of the tag. These changes can be detected to develop humidity sensors [110], displacements sensors, and temperature sensors [73]. However, tag antenna based sensing has limited resolution and is only applicable to slowly varying quantities. Furthermore, a common drawback in tag based antenna sensing is the need for continuous calibration of the RF channel to mitigate path loss and multipath effects.

6.1.3 This work: hybrid analog-digital backscatter sensing

As described above, traditional analog backscatter requires specialized antennas/resonant structures which are incompatible with standard digital backscatter. In this work, we implement analog backscatter by directly connecting the antenna to a sensor whose impedance varies as a function of the sensed quantity. Using this approach of analog backscatter, first we undertake a comparative study of analog *vs.* digital backscatter. We conclude that the optimal solution for high data rate

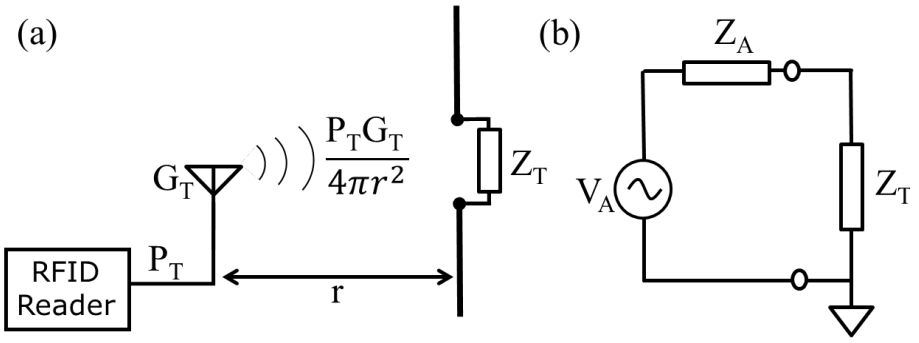

Figure 6.1: (a) A typical RFID system (b) Equivalent model of antenna

sensing is a hybrid platform that combines addressable digital backscatter with selectable sensors which use analog backscatter for high data rate sensing. We demonstrate our approach of hybrid backscatter by augmenting a digital RFID platform (WISP) with a microphone to develop the first digitally addressable battery free microphone.

The first prototype of the hybrid approach and preliminary results were presented in [195]. In this chapter, we undertake a comparative study of analog *vs.* digital backscatter to motivate the use of hybrid backscatter platforms. We also conduct an in-depth analysis of the working of the backscatter microphone. Based on this analysis, we optimize the backscatter microphone design and report substantial improvements in performance (compared to [195]). The outline of the paper is as follows. In §6.2 we discuss the basics of backscatter communication followed by a comparative study of analog and digital backscatter sensing in §6.3 and hybrid backscatter in §6.4. In §6.5 we discuss a backscatter microphone and in §6.6 we describe our implementation of hybrid backscatter platform. We evaluate our system in hybrid backscatter platform in §6.7 and conclude the paper in §6.8.

6.2 Background on Power Harvesting and Backscatter Communication

Figure 6.1(a) illustrates a typical RFID system which consists of an RFID reader transmitting electromagnetic waves with RFID tag(s) located in its vicinity. The power density of electromagnetic

waves incident on a tag placed at a distance r from the reader is given as

$$S = \frac{P_t G_t}{4\pi r^2} \tag{6.1}$$

where P_t is the power transmitted by the reader and G_t is the gain of the reader antenna.

The incident power is collected by the aperture of the tag antenna and delivered to the terminating impedance. At the same time, a fraction of the incident power is scattered by the antenna. We use the Thevenin equivalent circuit model shown in Figure 6.1(b) to analyze the received and backscattered power. Here, V_A is the open circuit voltage on the antenna terminals, $Z_A = R_A + jX_A$ is the complex antenna impedance and $Z_T = R_T + jX_T$ is the impedance across the antenna terminals [70]. Although it has been shown in [81] that Thevenin equivalent circuit model do not represent the true behavior of antennas, in cases such as thin wire antennas like dipole which is the topic of our study, the equivalent circuit model is sufficiently accurate.

6.2.1 Power Harvesting and Duty Cycle Operation

The RF power collected by the tag is determined by the effective aperture of its antenna (A_e) and is given by

$$P_{received} = SA_e = \frac{P_t G_t}{4\pi r^2}.\frac{\lambda^2 G_r}{4\pi}\left[1 - |\Gamma|^2\right] \tag{6.2}$$

where, G_r is the gain of the receive antenna, λ is the wavelength and Γ, the reflection coefficient is defined as

$$\Gamma = \frac{Z_T - Z_A^*}{Z_A + Z_T} \tag{6.3}$$

For maximum power transfer ($|\Gamma| = 0$), the antenna should be terminated by the complex conjugate of its impedance i.e. $Z_T = Z_A^*$. This represents the matched impedance state which is optimal for power harvesting.

As seen from (6.2), the power received by the tag drops rapidly as the distance between the tag and the reader increases. Typical sensing platforms have significant power consumption and they operate by duty cycling their operation [187]. During the sleep/off state, the circuit continuously harvests energy while consuming minimal power. Once sufficient energy is available, the system

wakes up and executes a task (sensing, computation and communication). By equating the energy harvested with the total energy consumed over each operating cycle and using (6.2), the duty cycle (D) of the system can be estimated as

$$D = \frac{T_{on}}{T_{on} + T_{sleep}} = \frac{P_{received} * \eta_{harvester} - P_{leakage}}{P_{load}} = \frac{\frac{P_t G_t}{4\pi r^2} . \frac{\lambda^2 G_r}{4\pi} \left[1 - |\Gamma|^2\right] * \eta_{harvester} - P_{leakage}}{P_{load}} \tag{6.4}$$

where $\eta_{harvester}$ is the efficiency of the power harvester, $P_{leakage}$ is leakage power consumption (during sleep/off mode) and P_{load} is the average power consumption of the active mode (excluding leakage power). It can be seen that the maximum operating range of the system is the distance at which the harvested power is equal to leakage power and the duty cycle drops to zero. Note that the maximum operating range is a function of leakage power and is independent of active power consumption.

6.2.2 Backscatter Communication

Passive RFID tags modulate the impedance across the antenna terminals which causes a change in the field backscattered by the tag. These changes in the field are detected by the reader and are used to decode the bits transmitted by the tags. For thin wire antennas such as dipoles which are used in RFID tags, the re-radiated can be computed using the equivalent circuit model [70] shown in Figure 6.1(b) and can be written as

$$P_{backscatter} = S\sigma = \frac{P_t G_t}{4\pi r^2} . \frac{\lambda^2 G_r^2}{4\pi} |1 - \Gamma|^2 \tag{6.5}$$

where, σ, the scalar radar cross section (RCS) of the antenna is defined as

$$\sigma = \frac{\lambda^2 G_r^2}{4\pi} |1 - \Gamma|^2 \tag{6.6}$$

The backscattered field from the tag antenna undergoes path loss in the reverse direction and the signal power received by the reader located at a distance r from the tag is given as

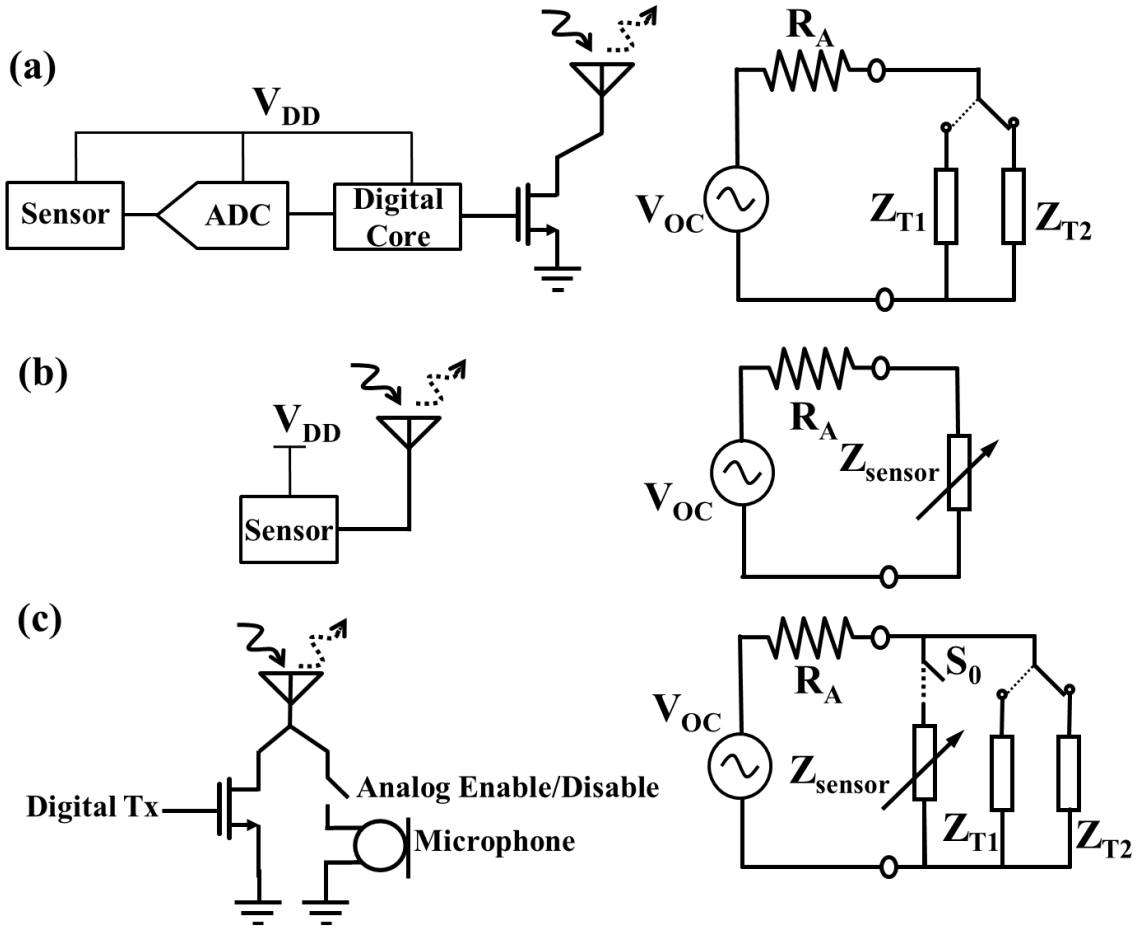

Figure 6.2: (a) Digital backscatter sensing platform (b) Analog backscatter sensing platform (c) Hybrid analog-digital backscatter sensing platform

$$P_{reader} = \frac{P_{backscatter} A_e}{4\pi r^2} = \frac{P_t G_t^2 \lambda^2 \sigma}{(4\pi)^3 r^4} \tag{6.7}$$

where σ is defined in (6.6). Subsequent analysis of backscatter communication will consider that all antennas are operating at their resonant frequency i.e. the frequency at which $Z_A = R_A$.

Digital Backscatter Sensing

A typical digital backscatter sensing platform (shown in Figure 6.2(a)) uses an ADC to sample the output of the sensor. The digitized sensor data is then processed and transmitted as binary data

packets by using a FET to switch the impedance across the antenna between matched ($Z_{T1} = R_A$) and short ($Z_{T2} = 0$) states. This switching operation implements binary amplitude/phase shift keying and maximizes the difference in the power backscattered by the tag (while being able to simultaneously harvest in one state). The difference in the backscattered digital power received by the RFID reader is typically analyzed using differential radar cross section [155], and in the ideal case of switching between matched and short states can be written as

$$P_{reader\_d} = \frac{P_t G_t^2 \lambda^2 \Delta\sigma}{(4\pi)^3 r^4} = P_t G_t^2 G_r^2 \left(\frac{\lambda}{4\pi r}\right)^4 \tag{6.8}$$

Analog Backscatter Sensing

Analog backscatter generalizes the concept of backscatter by continuously varying the terminating impedance of the antenna. Traditional approaches to analog backscatter use the antenna as the sensing element [73, 110]. However, our approach to analog backscatter sensing (as shown in Figure 6.2(b)) consists of a sensor (passive or active) directly connected to the terminals of the antennas. The impedance of the sensor (Z_{sensor}) varies as a function of the sensing quantity, thereby, modulating the backscatter power with sensor information. Using (6.3), (6.6) and (6.7), the analog backscattered power as a function of the impedance of the sensor can be written as

$$P_{reader\_a} = \frac{\lambda^4 P_t G_t^2 G_r^2}{(4\pi r)^4} \frac{4R_A^2}{|R_A + Z_{sensor}|^2} = f(Z_{sensor}) \tag{6.9}$$

The reader captures the backscattered power and uses (6.9) to decode, digitize and process the sensor information. Under certain conditions (as shown in Section 6.5) the decoding process can be as simple as a band pass filter.

Hybrid Analog-Digital Backscatter Sensing

Analog and digital backscatter can be combined to develop a hybrid backscattering platform. Figure 6.2(c) shows a schematic and equivalent circuit model for a hybrid backscattering platform which time multiplexes digital and analog backscatter modes. In the digital backscatter mode, the analog switch (S_0) is turned off, disconnecting the sensor from the antenna, and the platform

behaves as a normal digital backscatter platform as described in §6.2.2. In analog backscatter mode, the analog switch (S_0) is turned on, connecting the sensor (in parallel with the terminating impedance, Z_{T1}) to the antenna. The backscattered power as a function of the impedance of the microphone (Z_{sensor}) sensor can be written as

$$P_{reader\_h} = \frac{\lambda^4 P_t G_t^2 G_r^2}{(4\pi r)^4} \frac{4R_A^2}{|R_A + Z_{T1} \,||\, Z_{sensor}|^2} \tag{6.10}$$

6.3 Analog Backscatter versus Digital Backscatter

A comparative study of analog *vs*. digital computation has been reported in literature which shows that for a given bandwidth (set by the specifications of the task), analog performs better (in terms of both chip area and power) for low signal to noise ratios whereas digital has better performance for high signal to noise applications [188]. It was concluded that for any given task, the optimal solution was neither analog nor digital alone, but a combination/hybrid of both analog and digital computation. In this section, we undertake a similar **qualitative** study to understand the performance of analog backscatter sensing and digital backscatter sensing in terms of power consumption, duty cycle and signal to noise ratio. We have discussed analog backscatter and digital backscatter techniques for sensing platforms. However, it is not apparent which technique is best suited for a particular application. Each sensing application has certain specifications in terms of sampling frequency/data rate and operating distance from the reader. In this section we will **qualitatively** study the performance of RF powered analog and digital backscatter systems. The analysis is intended to be generic and serve as a guideline to understand the trends and inherent tradeoffs of analog and digital backscatter sensing.

6.3.1 Power Consumption

The power consumption ($P_{digital}$) of a digital backscatter platform (in Figure 6.2(a)) consisting of an active sensor (including amplifiers and anti-aliasing filters), ADC, digital core and the switching

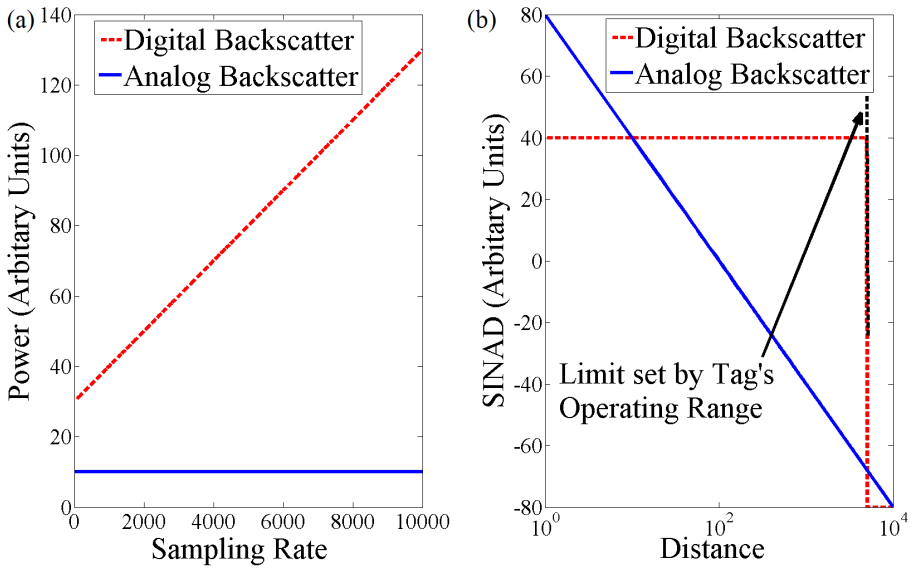

Figure 6.3: (a) Tradeoff between power consumption and sampling rate (b) Tradeoff between *SINAD* and distance

transistor can be written as

$$\begin{aligned} P_{digital} &= P_{sensor} + P_{core} + P_{ADC} \\ &= P_{DC} + CV_{DD}^2 f_s + FOM * 2^{ENOB} * f_s \end{aligned} \tag{6.11}$$

where, P_{DC} is the static power consumption of the sensor (and associated amplifiers and anti-aliasing filters). C represents the effective capacitance switching at sampling frequency (f_s) in the digital core and FET. *FOM* is the figure of merit of the ADC and *ENOB* represents the ADC's effective number of bits [206].

A typical analog backscatter sensing platform (Figure 6.2(b)) consists of a passive sensor (such as thermistors, light dependent resistor, microphone) directly connected to the antenna. However, for completeness let us express the power consumption as

$$P_{analog} = P_{sensor} = P_{analog\_DC} \tag{6.12}$$

In analog backscatter, sampling and subsequent filtering of sensor data is done by the reader. This enables the sensor data to be sampled and filtered/processed at high rates independent of the

power constraints of the sensing platform. Hence, the power consumption of an analog backscatter platform is independent of sampling frequency.

Figure 6.3(a) shows a comparison of the power consumption in terms of sampling rate. The power consumption of the digital system increases linearly with increasing sampling frequency whereas the power consumption of the analog system is constant. Hence, analog backscatter is a more power efficient technique for sensing. Please note the preceding analysis does not take into account the loss in the harvested power during backscatter communication.

6.3.2 Signal to Noise and Distortion Ratio (SINAD)

In typical digital backscatter systems, as long as the tag is within the operating range of the reader, the digital communication from the tag to the reader can be considered error free. In such a scenario, the noise is limited by the ADC front end and the signal to noise and distortion ratio (*SINAD*) of digital sensor data is a function of *ENOB* and can be written as [206]

$$SINAD_{digital} = ENOB * 6.02 + 1.76dB \tag{6.13}$$

For analog backscatter systems, the sensor information is modulated with the backscattered power. Using (6.7), *SINAD* of sensor data as a function of distance between the reader and the tag can be written as

$$SINAD_{analog} = 10 * log_{10}\left(\frac{P_0}{r^4\left(P_d + N\right)}\right) \tag{6.14}$$

where, P_0 is the backscatter power received at unit separation, P_d is the distortion and N is the noise floor of the RF front end of the reader.

Figure 6.3(b) shows a comparison of *SINAD* between analog and digital backscatter systems across a range of operating distances. It can be seen that *SINAD* of analog backscatter drops at the rate of 40 dB/decade with increase in distance (equivalently 10 dB/decade drop with path loss) whereas digital backscatter has a near constant value till the maximum operating range. So, for applications which require high signal to noise ratio, digital backscatter is better suited (assuming there is no power constraint).

6.3.3 Duty Cycle

As shown in Section 6.2.1, the high power consumption of digital backscatter systems requires that RF powered digital backscatter systems be duty cycled. Using (6.4), the duty cycle of digital backscatter systems can be expressed as

$$D_{digital} = \frac{P_{DC\_received}}{r^2 P_{digital}} = \frac{P_{DC\_received}}{r^2 \left(P_{DC} + CV_{DD}^2 f_s + FOM * 2^{ENOB} * f_s\right)}$$

where, $P_{DC\_received}$ is available DC power from the harvester at unit distance. As sampling frequency increases the power consumption of digital systems increases, which forces a reduction in the duty cycle of digital systems. Similarly, the power available from the reader reduces with increase in distance, which results in reduction of duty cycle for the digital backscatter system.

On the other hand, typical analog backscatter systems are passive, consuming zero/minuscule power. For the distances under consideration, analog backscatter systems are always active or in other words they operate at 100% duty cycle. Hence, for applications which require high duty cycles (response rate), analog backscatter is better suited (neglecting *SINAD* of sensor data).

6.3.4 SINAD for reconfigurable ADCs

Typical digital backscatter systems have ADCs with fixed resolution (*ENOB*) which determines the *SINAD* of digital sensor data. This has been the premise in previous sections. Here we will consider the scenario where the resolution of the ADC can be modified. This could be accomplished in the design phase of the ADC, or in the case of reconfigurable ADC's, the digital core can modify the resolution (and hence the power consumption) on the fly. This enables us to study the maximum achievable *SINAD* for digital backscatter as a function of sampling frequency and distance. (6.15) can be rearranged and the *ENOB* of the ADC can be written as

$$ENOB = log_2\left(\frac{P_{DC\_received}}{r^2 D} - P_{DC} - CV_{DD}^2 f_s\right) - log_2\left(f_s FOM\right)$$

Using (6.13) and (6.15), *SINAD* of digital backscatter system can be ascertained. Note that *ENOB* is a function of three parameters (r, f_s and D) which complicates the analysis. Since duty

cycle was analyzed above, we will assume a fixed duty cycle digital backscatter system for this analysis.

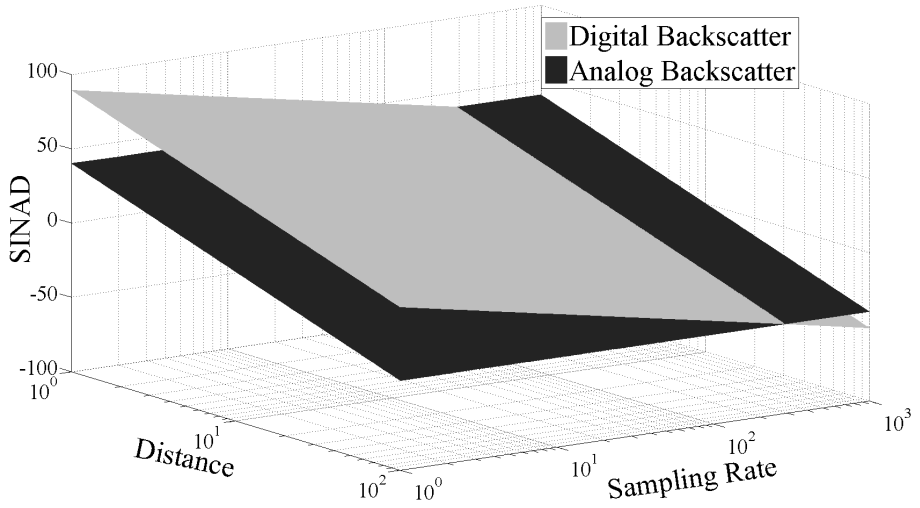

Figure 6.4: Study of tradeoff between *SINAD*, distance and sampling rate for analog and digital backscatter sensing

For analog backscatter system, (6.14) can be used to determine *SINAD* as the function of distance and sampling rate. Figure 6.4 plots the variation of *SINAD* of analog and fixed duty cycle digital backscatter systems with distance and sampling frequency. As discussed above, *SINAD* of analog backscatter is independent of sampling rate but drops with increasing distance. For digital backscatter systems, since power received from the reader decreases with increases in distance, *SINAD* has to be decreased to reduce power consumption. Similarly, as sampling frequency increases, the power consumption of digital backscatter increases which has to be compensated with a reduction of *SINAD*. Hence, in terms of *SINAD*, digital backscatter performs better at low sampling rates whereas analog backscatter is preferred for high sampling rates.

In conclusion it can be seen that digital backscatter is better suited for low sampling rate and/or low response rate sensing applications, whereas, analog backscatter is better suitable for high sampling rate applications. Note that the above analysis is generic and device independent and has been performed to understand the inherent tradeoffs and trends in analog and digital backscatter sensing. The actual regions of tradeoff depend on the characteristics of the sensor employed and

the specific requirements of the sensing application.

However, it should be acknowledged that as technology scales, digital computation becomes more power efficient [123]. This implies that digital backscatter sensing platforms will be able to operate at higher sampling rates, at higher duty cycles and at farther distances from the reader. Unfortunately, analog sensors and analog backscatter do not follow the same trend. Additionally, new digital signal processing techniques such as compressing sensing could potentially enable digital platforms to accomplish the sensing goal at lower sampling rates. Hence, trends in semiconductor technology and signal processing favor digital backscatter for sensing applications.

6.4 Hybrid Sensing

In addition to the quality (*SINAD*) of sensor data, typical applications require multiple sensors (with a wide range of sensor data quality and sampling rate specifications), addressability, selectivity and control. Unfortunately, all the requirements cannot be satisfied by either analog or digital backscatter independently. The optimal solution is to use digital backscatter for addressability, control and low sampling/update rate sensing and analog backscatter (addressed and controlled digitally) for high data rate sensing. This can be accomplished by combining analog and digital backscatter into a hybrid backscatter platform. Another advantage of hybrid sensing is that hybrid backscatter utilizes the same quasi-static RF channel (for typical switching periods) for digital and analog backscatter communication. Since digital backscatter switches between two pre-determined states, it can be used to estimate the characteristics of the channel and in turn calibrate analog backscattered signal/sensor data.

6.5 Real Time Battery Free Microphone

Audio or speech is a very popular sensed quantity which finds numerous applications in human smart spaces. Battery free wireless detection of heart sounds (250 Hz bandwidth) has been reported in [137]. However, wireless transmission of speech on RF powered systems has been a challenge due to the sampling rate ($\geq$ 6.8 kHz) requirements. These are relatively high rates for backscatter sensing. Typical applications (such as speech recognition) require a minimum time

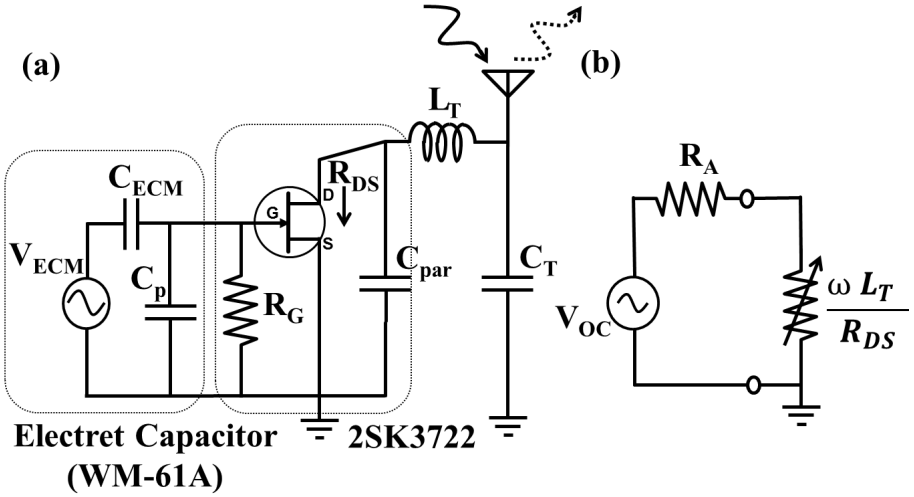

Figure 6.5: (a) Equivalent model of the WA61A microphone connected to antenna to transmit audio data using analog backscatter (b) Equivalent model for analog backscatter analysis

duration (typically 75 ms or more) of continuously sampled data for processing. This implies that each operating cycle must capture, process, packetize and transmit 75 ms long samples. Such high cumulative power consumption severely cripples the duty cycle of the digital system. As an example, let us consider integrating an off the shelf low power digital microphone (*MP45DT02*) with the WISP. The active power consumption of WISP is 1.12 mW and *MP45DT02* has an additional power consumption of 1.17 mW (650 μA at 1.8V). At a nominal distance of 2 m from the reader, ignoring protocol limitations, the maximum achievable duty cycle for the system is 7.8% i.e. the WISP can transmit one cluster of samples (75 ms long) every second. Moreover, this rate rapidly drops with each increase in distance. Such low and varying data rates are impractical for current implementations of typical applications (such as speech recognition and event detection). Although the use of WISP and *MP45DT02* is not the optimal approach, it illustrates the limitations of digital backscatter platforms in transmission of speech.

Analog backscatter systems can transmit continuously at high data rates in real time. Furthermore, human ears and speech recognition software are capable of processing noisy speech data. Taking these factors into account, analog backscatter which can accommodate high data rates (at near zero power), albeit at distance-dependent SNR, is a good candidate for transmission of speech

on RF powered platforms.

To demonstrate analog backscatter of speech, we chose an electret condenser microphone (*WM61A* by Panasonic). Electret condenser microphones (ECM) use electret material (which has a quasi-permanent electrostatic charge) as the microphone's diaphragm. The electret diaphragm is separated from a fixed metal back plate by an air gap, forming a capacitor. As sound waves move the diaphragm, the distance between the electret diaphragm and the metal back place changes, resulting in a change in the capacitance. Since the charge stored on the electret diaphragm is fixed, this results in small signal voltage change. We can model the electret capacitor as a capacitive voltage source (V_{ECM}) in series with a $2 \sim 5pF$ capacitor (C_{ECM}) as shown in Figure 6.5(a). Because the output impedance of the electret capacitor is very high, a high input-impedance (modeled as R_B) device such as a JFET ($2SK3722$) is connected to its output.

Let P_{audio} be the pressure of the audio waves striking the diaphragm. The voltage induced at the drain source terminal can be written as

$$V_{GS} = k_{ECM} * P_{audio} \tag{6.15}$$

where k_{ECM} accounts for the transducer gain of the electret diaphragm and the voltage transfer from V_{ECE} to V_{GS}. In traditional applications, the JFET is biased in saturation region using an external DC bias and a load resistor to operate as a common source amplifier. However, to implement analog backscatter, we directly connect the source terminal of the JFET to the antenna (using an L-C tuning network) as shown in Figure 6.5. Since there is no DC voltage on the drain terminal, the voltage across the drain and source terminals of the JFET, $V_{DS} \approx 0$, biasing the device in the triode region. In triode region, the impedance looking into the drain terminal of the JFET can be written as

$$R_{DS} = \frac{R_{DS\_ON}}{\left(1 - \frac{V_{GS}}{V_p}\right)} \tag{6.16}$$

where, V_p is the pinch-off voltage and R_{DS_ON} is the impedance looking into the drain terminal for $V_{GS} = 0$. In order to tune the input impedance of the JFET to the impedance of the antenna, an

L-C tuning network is employed and for typical values of quality factor ($\geq$ 4), the impedance of the microphone as seen by the antenna can be written as

$$R_{sensor} = \frac{\omega^2 L_T^2}{R_{DS}} = R_s \left(1 - \frac{V_{GS}}{V_p}\right) \tag{6.17}$$

where, R_s is the impedance seen by the antenna for zero audio/speech input. The signal received by an RFID reader is typically analyzed in the voltage domain. Using the equivalent circuit model shown in Figure 6.5(b) and combining (6.17) and (6.9), the backscatter signal (in terms of scalar voltage) received at the reader can be written as

$$\begin{aligned} V_{reader\_a} &= \frac{K}{R_A + R_s\left(1 - \frac{V_{GS}}{V_p}\right)} \\ &\approx \frac{K}{R_A + R_s}\left[1 + \frac{V_{GS}}{V_p}\left(\frac{R_s}{R_A + R_s}\right)\right] \\ &= \frac{K}{R_A + R_s} + \frac{k_{ECM}}{V_p}\frac{KR_s}{(R_A + R_s)^2}P_{audio} \end{aligned} \tag{6.18}$$

where K incorporates the gain of the RF front end, path loss, transmit power and antenna parameters. For typical values of V_{GS} (10's of mV), the binomial approximation can be used to write the voltage as a linear function of the sound input (as shown in (6.18)). For maximum gain/sensitivity, we should set $R_s = R_A$ i.e. the impedance of microphone for zero input should be matched to antenna impedance. Compared to the un-optimized system first reported in [195], this matching increases the sensitivity of the microphone by at least 15 dB.

Similarly, for a hybrid microphone using (6.17) and (6.10) wherein $Z_{T1} = R_A$ (for optimal power harvesting), the backscatter signal (in terms of scalar voltage) received at the reader can be written as

$$V_{reader\_h} = \frac{K(R_A + R_s)}{R_A(R_A + 2R_s)} + \frac{k_{ECM}}{V_p}\frac{KR_s}{(R_A + 2R_s)^2}P_{audio} \tag{6.19}$$

In case of a hybrid microphone, maximum sensitivity is achieved for $R_s = R_A/2$. Note that the hybrid microphone is less sensitive than pure analog backscatter microphone by a factor of 3 dB.

However, given the additional benefits of hybrid backscatter, the loss in sensitivity is acceptable for typical sensing applications.

6.6 Hybrid Analog-Digital Backscatter Sensing Platform

To demonstrate the concept of hybrid backscatter, we developed a hybrid analog-digital backscatter platform by leveraging the micro-controller based Wireless Sensing and Identification platform (WISP) [187]. WISP is an RF-powered platform featuring a fully programmable 16-bit micro-controller (*MSP*430) and an array of sensors, that communicates with commercial RFID readers at 915 MHz using the EPC global Class-1 Generation-2 (Gen 2) protocol. We integrated the backscatter microphone discussed in section 6.5 with the WISP to design a hybrid WISP. The hybrid WISP by default operates in digital mode and switches into analog mode to backscatter analog audio sensor data.

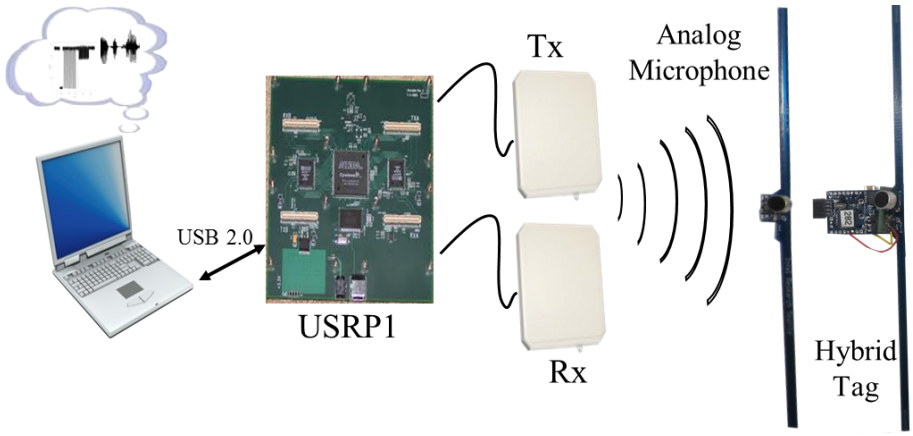

Figure 6.6: Experimental Setup consists of a USRP based RFID reader and the backscatter microphones

6.6.1 Hybrid WISP Design

As shown in Figure 6.2(c), the RF front end of the hybrid WISP consists of a dipole antenna connected to the digital backscatter MOSFET (*BF1212WR*) in parallel to the microphone. The microphone is switched in and out of the backscatter network using an RF switch (*ADG902*). The

RF switch and the digital backscatter MOSFET are both controlled by the output ports of the WISP micro-controller. WISP implements the Gen 2 protocol completely in firmware, and we modified its state-machine so that it transitions into analog backscatter mode when it receives a Gen 2 *READ* command that is addressed to it. After a configurable amount of time, the firmware switches the microphone out of the antenna network and returns to digital mode. During the analog mode, although the microphone consumes zero power, the MSP430 micro-controller operates in low power mode (LPM3) consuming 4.5 μW. This power consumption is miniscule and only fractionally more than the leakage power (2.9 μW) of WISP which enables the hybrid WISP to remain in analog mode for substantial period of time (e.g. few seconds for 100 μF storage capacitor).

6.6.2 Custom RFID Reader Design

To experimentally demonstrate our approach to hybrid backscatter, we extended a software-radio based Gen 2 reader previously developed using the Universal Software Radio Peripheral and GNU Radio [76]. The RFID reader and the hybrid WISPs by default operate in digital mode, where the reader continuously executes inventory rounds to query for WISPs that are present in the vicinity. During a query round, the reader can selectively transition a given WISP to analog mode by transmitting Gen 2 *READ* command addressed to that WISP. Once in analog mode, the hybrid WISP backscatters sensor (audio) data for a predetermined period of time and the RFID reader logs the digitized sensor data. At end of this time, the WISP transitions back to digital mode and the reader continues to query for other WISPs. The sensor (audio) data is recovered by passing the logged data through a band pass filter (300 Hz-3.4 kHz in case of speech) using GNU Radio signal processing blocks on the reader.

6.7 Results and Discussion

Figure 6.7 shows a communication cycle between the hybrid WISP and the software defined RFID reader. The reader initiates a query round and receives the WISP ID as a part of the EPC code. The reader then issues a READ command upon which the hybrid WISP transitions into the analog

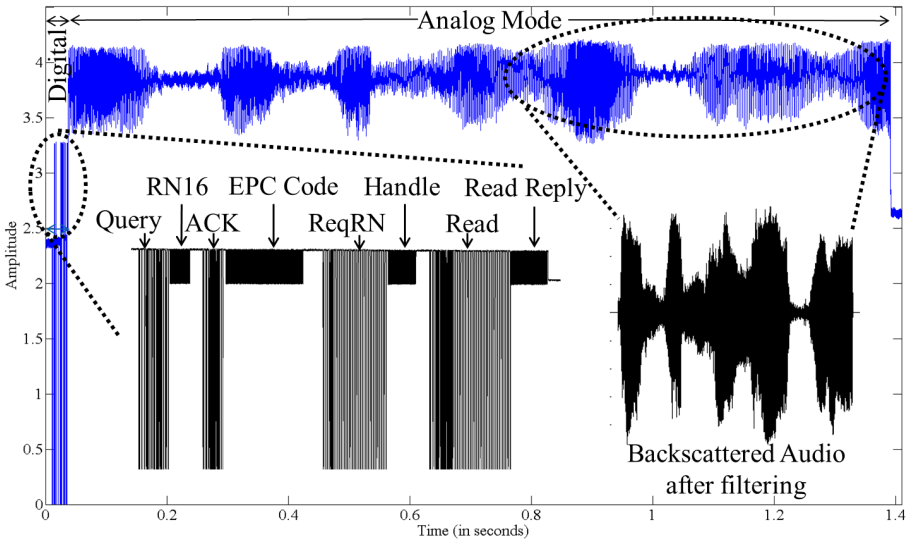

Figure 6.7: The signal trace of a communication cycle between the hybrid WISP and the software defined RFID reader

mode, backscattering analog sensor (audio) data. The inset on the left shows the EPC Gen2 commands transmitted between the reader and the WISP. The right inset shows a magnified image of the recovered speech by the reader.

The experimental setup shown in Figure 6.6 was used to evaluate the performance of the backscatter microphone. The analog/hybrid microphone and the reader antennas were placed 0.5 m apart at a height of 1 m from the ground and the reader was configured to transmit at maximum power (26.7 dBm). To minimize multipath effects, path loss was introduced using variable attenuators in the forward and return path. A constant tone at mid band frequency (1.75 Khz at 90 decibels) and a reference audio clip (at 90 decibels) were played as input to the microphone to evaluate *SINAD* and Perceptual Evaluation of Sounds Quality (*PESQ*) scores [64] respectively.

Figure 6.8 shows the quality of speech received from the analog and hybrid microphone as a function of RF signal strength and Friis equivalent distance. The *SINAD* of received speech decreases linearly (at rate of 10 dB/decade) with signal attenuation which agrees with our hypothesis. Analog backscatter microphone performs better than the hybrid backscatter microphone by a factor

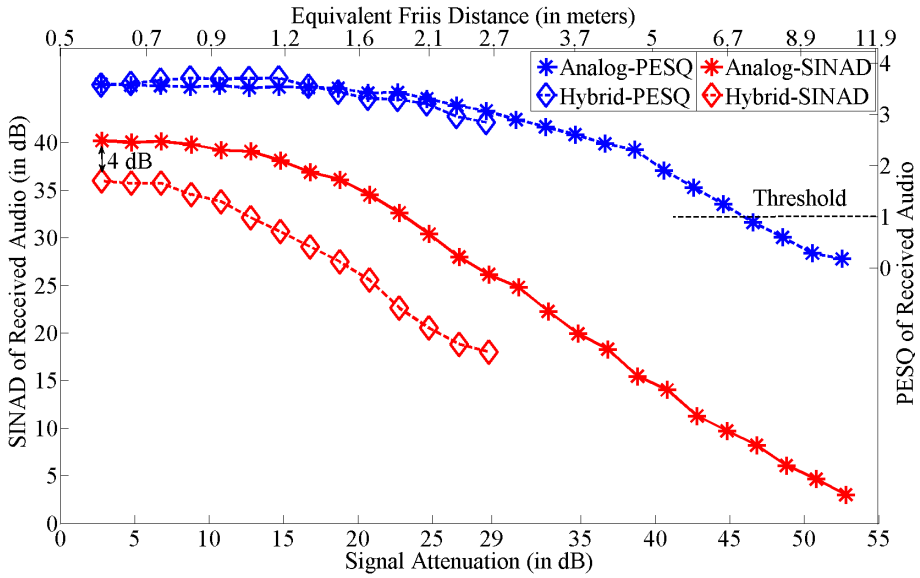

Figure 6.8: Performance of analog backscatter and hybrid backscatter microphone

of 4 dB which is along expected lines. 3 dB difference was predicted in §6.5 and the additional 1 dB can be attributed to loss in the switch and impedance mismatch. However, for short distances (up to 1 m), the power received by the tag is very high, which introduces non-linearity in the JFET and results in saturation. Lastly, as path loss increases, hybrid backscatter starts to degrade at a rate higher than the predicted 10 dB/decade. This is most likely due to mismatches introduced by the non-linear variations in the input impedance of the harvester.

Since the quality of the received audio is a function of distance, the maximum operating range of the system is determined by the minimum acceptable *SINAD* and/or *PESQ* by the application. As an example, for human hearing, $PESQ \geq 1.0$ is decipherable and the expected operating range of analog microphone in this case is 7.4 m. The hybrid microphone works up to 2.7 m which is the maximum operating range of the WISP for 26.7 dBm reader output power. However, if the microphone is integrated with an RFID platform with larger operating range, extrapolating the graph, the expected operating range of the hybrid microphone would be 4.7 m.

Although the use of USRP as RFID reader provides flexibility, it has some limitations. The maximum transmit power of USRP is 26.7 dBm which is 3.3 dBm lower than the limit imposed

by FCC (for 36 dBm EIRP). Increase in transmit power of the reader to 30 dBm would result in a 20% increase in the operating range of the analog backscatter microphone and 40% increase in the operating range of hybrid backscatter microphone. The RF front end of the USRP is not optimal and has higher noise floor than commercial RFID readers, which results in significant signal degradation. Use of a lower noise factor RF front end can substantially increase the *SINAD* of received audio and extend the operating range of the microphone.

6.8 Conclusion

We have undertaken a comprehensive study of analog and digital backscatter sensing. After careful analysis, we concluded that the optimal strategy for high data rate battery free sensing is a hybrid of analog and digital backscatter sensing. This approach combines the addressability and control of digital backscatter with high data rate (and low/zero power) analog backscatter sensing. We have demonstrated a practical implementation of hybrid backscatter and developed an addressable real-time battery free microphone. The battery free hybrid microphone operates at a distance upto 2.7 m with a 26.7 dBm transmit power reader. Use of a 30 dBm output power reader should extend the operating range to 4.5 m (operating range of WISP). Such battery free microphones can be used to develop battery free voice communicators, audio event detectors and audio based localizers for pervasive computing applications.

Chapter 7

THESIS CONCLUSION AND FUTURE WORK

The computer revolution began with the main frame computer ENIAC in 1946. In the 1950's and 60's, mainframe computers used to occupy rooms and were available to a very tiny subset of population. The invention of the transistor followed by the CMOS integrated circuits ushered in a technology revolution. The PC revolution began in the 1970's which gave us the personal computer. Development of wireless communication, CMOS scaling and batteries led to laptops, cellphones, smart watches, wearable devices and other portable devices which have become commonplace.

Although we have made tremendous progress, our current devices are still either tethered to power cords or use batteries which require constant supervision and maintenance. We started with the goal of designing systems which defy this rule. We introduced ambient backscatter, a new communication primitive where devices leverage ambient RF signal to harvest power and communicate with each other. Next, we designed full-duplex backscatter, which enables instantaneous feedback between two battery free devices. We used full-duplex backscatter to design a network stack which minimizes the energy wastes that occur due to collisions and packet errors making deployment of battery-free systems robust and practical. In addition to this, we also introduce hybrid analog digital backscatter where we demonstrated the first battery-free microphone. Although these projects were intellectually challenging and rewarding, the solutions were tailored to only niche applications.

We turned our attention to Wi-Fi, one of the most popular standard in the world, which is found ubiquitously in indoor environments such as homes, offices and public spaces. If we can develop solutions using the Wi-Fi standard, we can impact a wide range of applications. This realization led to *power over Wi-Fi*, the first system that delivers power to low-power sensors and devices using existing Wi-Fi chipsets. We use the Wi-Fi router, a ubiquitous part of wireless

communication infrastructure to provide far field wireless power to a range of sensors and batteries without significantly impacting the Wi-Fi network performance. Next, we used Wi-Fi to tackle the issue of providing connectivity at low power. We introduced Passive Wi-Fi, the first system to demonstrate that 802.11b Wi-Fi packets can be synthesized using backscatter communication while consuming 3–4 orders of magnitude lower power than traditional Wi-Fi chipsets. Passive Wi-Fi transmission can be decoded on any standard complaint Wi-Fi device and coexist with other devices in the ISM band, without incurring the power consumption of carrier sense and medium access control operations. We showed that Passive Wi-Fi brings low power connectivity to all devices and could be a game changer in bringing Wi-Fi connectivity to the Internet of Things.

Moving forward, we believe that the approaches that have been proposed in this thesis could be basis for future research. Some of the potential avenues are listed below.

7.1 Wireless Power Delivery

(b) High sensitivity RF rectifiers. We have demonstrated the use of TV, RFID and Wi-Fi signals for delivering far field wireless power to devices. However, all the prototype implementations in this thesis are based on commercial off the shelf components which limit the performance of the RF harvester. Specifically, we note that we were able to achieve sensitivity in the range of -15 to -20 dBm which greatly limit the scope of wireless power delivery. For wireless power to be truly ubiquitous, there is a need for an integrated circuit solution which can push the sensitivity lower and increase the coverage range of wireless power transfer [101]. Specifically, the holy grail would be to achieve high sensitivities (-30 dBm or lower) for cold start/boot up operation (when every node starts at 0V) where battery free systems can harvest power. In addition to the obvious need for lower threshold and lower leakage Schottkey diodes for the rectifier, the key to solving this challenge will be an integrated circuit DC-DC or boost converter optimized to efficiently convert the output DC output of the RF rectifier to usable voltages in the range of 1-1.8 V.

(b) Wireless power delivery with MIMO. The PoWi-Fi system uses multiple antennas to transmit concurrently on different Wi-Fi channels. In principle, we can reuse the PoWi-Fi architecture and

use MIMO techniques such as beam forming to deliver higher power to devices. Specifically, using beam forming techniques, we can counter the multi-path fading effects and provide a more uniform coverage for wireless power delivery. However, the challenge with using beam forming for wireless power delivery is estimation of the channel without incurring a power penalty. Traditional beam forming estimates the channel by transmitting sounding packets which are extremely power intensive. In [68] authors showed that backscatter communication in combination with an extensive characterization of the non-linear behavior of the rectifier can be used to estimate the wireless channel. Techniques on similar lines adapted for the Wi-Fi protocol could potentially open 802.11 MIMO power delivery to low power devices.

(e) Future clean-slate designs. The PoWi-Fi is a general design for power delivery in the ISM bands. As Wi-Fi access and densities continue to grow in the ISM band, solutions that deteriorate Wi-Fi performance by jamming any specific frequency are not desirable. Our power delivery solution is integrated with the Wi-Fi protocol and hence can deliver power while having minimal impact on Wi-Fi traffic. Future designs would generalize our multi-channel approach to operate across multiple ISM bands (e.g., 900 MHz, 2.4 GHz and 5 GHz). We believe that this paper takes a significant step towards that goal. Furthermore, our design was limited by kernel levels access to a Wi-Fi chipset. If we can get access to the firmware on the Wi-Fi chipset, the PoWi-Fi system can be further optimized for power delivery.

(d) Security implications of wireless power delivery. As networks capable of delivering both power and data become prevalent, one can imagine a "power denial-of-service" (PDoS) attack in which a rogue device causes power starvation for other members of the network by generating signals designed to cause carrier sense events at the router. This opens up interesting research opportunities for understanding the tradeoffs for security mechanisms that protect against such attacks in an efficient manner.

7.2 Low Power Connectivity

(a) Rate adaptation for communication between battery-free devices. Ambient backscatter and full-duplex backscatter systems support data rates ranging form 100 bps to 10 kbps for communication between battery-free devices. However, during the prototype evaluation we manually set the filter components at the receiver to optimize the system for a single data rate. But in practice one could imagine that as the operating range changes, the tags would adapt their data rate. Algorithms to determine the optimal data rate and hardware design techniques which can switch between different filter components values are interesting areas for future research to develop robust tag to tag communication systems.

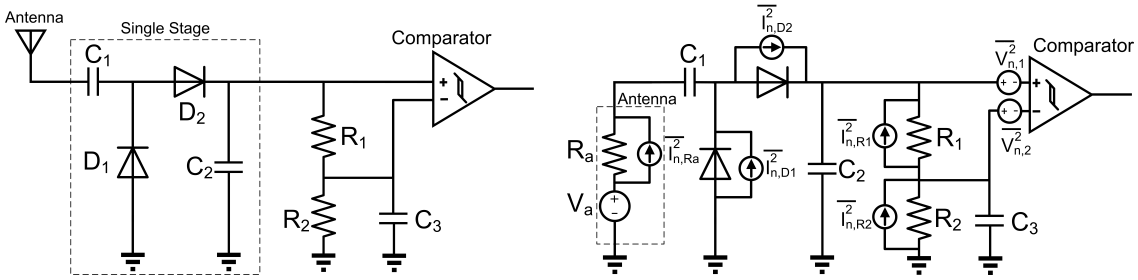

(a) The RF front end of a typical passive receiver in backscatter systems

(b) Noise model for the RF front end of a typical passive receiver in backscatter systems

Figure 7.1: **Architecture of a passive receiver.** The circuit implementation and the corresponding noise model of a passive receiver is shown.

(b) Theoretical analysis of operating range of tag to tag communication. In this thesis, we have experimentally demonstrated the operating range of tag to tag backscatter communication systems. In literature, the limits and tradeoffs in the design of traditional RFID backscatter has been extensively studied [80]. However, such an analysis on the limits and inherent tradeoffs in other relatively new backscatter configurations of tag to tag communication is missing. As an example, let us consider the passive receiver used in ambient backscatter (and tag to tag RFID) and its equivalent noise model as shown in Fig. 7.1. Similar to the analysis of a radio receive chain, the

noise model of a passive receiver can be analyzed analytically or in a circuit simulator to determine the signal to noise ratio (SNR) and noise figure (as a function of received signal strength) of the front end. Furthermore, this model can be used to minimize the noise figure by the redesign of the voltage multiplier and/or the passive elements of the filters. The SNR in combination with modulation schemes and coding techniques can be used to maximize and ascertain the maximum operating range of ambient backscatter and tag to tag RFID communication. Similarly, the downlink communication in Passive Wi-Fi can be analyzed and optimized to enable greater operating ranges. The noise model of the passive receiver (described above) can be used to minimize the noise figure and maximize the downlink communication.

(c) Antenna design. In our work, we limited our prototypes to dipole antennas for ambient backscatter and full duplex backscatter systems. The goal of the work was to demonstrate the novel tag to tag communication systems and use of standard dipole antennas eliminated the risk associated with new antenna topologies. In future design, one can imagine using PCB antennas such as PIFA and investigating the effect of antenna orientation between the tags and with respect to the signal source on the performance of the backscatter communication.

(d) Security for Passive Wi-Fi. The passive Wi-Fi devices could use WPA/WPA2 and ensure that their Wi-Fi transmissions comply with the Wi-Fi security specifications. Since these are digital operations, these could be easily implementable on the IC using baseband processing.

7.3 Low Power Sensing

(a) Microphones optimized for analog backscatter. In our work we used commercially available electret microphones wherein, the JFET is designed and optimized for active signal amplification. Instead, for analog backscatter, we reused the JFET for impedance modulation. If one custom manufactures microphones such as capacitive membranes, these can be directly interfaced with the antenna to greatly improve the performance of analog backscatter microphones.

(b) Increasing the range of analog backscatter. Our design was limited to the use of inherently noisy FLEX 900 daughter boards with 27 dBm output power on USRP1. If we upgrade our design

to the latest lower noise figure daughterboard and increase the transmit power to 30 dBm we can improve the quality of the audio and extend the operating range of analog backscatter based sensing systems.

(b) Exploring additional sensors for analog backscatter. The concept of analog backscatter can be extended to other sensing quantities as well. We can interface sensors such as thermistors and light dependent resistors to antenna and decode the sensor information using the backscattered signal. This technique however, requires using digital backscatter to calibrate the wireless channel prior to each sensor measurement.

BIBLIOGRAPHY

[1] 0402HP Series Inductors by Coilcraft. http://www.coilcraft.com/pdfs/0402hp.pdf.

[2] 41 dBu service contours around ASRN 1226015, FCC TV query database. http://transition.fcc.gov/fcc-bin/tvq?list=0&facid=69571.

[3] 8VSB vs. COFDM. http://www.axcera.com/downloads/technotes-whitepapers/technote_4.pdf.

[4] ADG902 RF switch datasheet. http://www.analog.com/static/imported-files/data_sheets/ADG901_902.pdf.

[5] ADG919 RF switch datasheet, Analog Devices. http://www.analog.com/static/imported-files/data_sheets/ADG918_919.pdf.

[6] ADMP801. http://www.cdiweb.com/datasheets/invensense/ADMP801_2_Page.pdf.

[7] Alexa – Top Sites in United States. http://www.alexa.com/topsites/countries/US. Loaded January 13, 2015.

[8] Altera DE1 FPGA development board. http://www.terasic.com.tw/cgi-bin/page/archive.pl?No=83.

[9] Android Wi-Fi analyzer. https://play.google.com/store/apps/details?id=com.farproc.wifi.analyzer&hl=en.

[10] Atheros targets cellphone with Wi-Fi chip. http://www.eetimes.com/document.asp?doc_id=1172134.

[11] Average U.S. home now receives a record 118.6 TV channels, according to Nielsen. http://www.nielsen.com/us/en/insights/press-room/2008/average_u_s__home.html.

[12] AVX BestCap. http://www.voti.nl/docs/OV7670.pdf.

[13] bq25570 by Texas Instruments. http://www.ti.com/lit/ds/symlink/bq25570.pdf.

[14] Cadence RFSpectre. http://www.cadence.com/products/rf/spectre_rf_simulation/pages/default.aspx.

[15] Cavity resonator microphone (also known as the great seal bug). http://www.spybusters.com/Great_Seal_Bug.html.

[16] Co-existence of Wi-Fi and Bluetooth radios by Marvell. http://www.marvell.com/wireless/assets/Marvell-WiFi-Bluetooth-Coexistence.pdf.

[17] Connecting the future: The latest research from Intel Labs. http://download.intel.com/newsroom/kits/idf/2012_fall/pdfs/IDF2012_Justin_Rattner.pdf.

[18] Cota by Ossia. http://www.ossiainc.com/.

[19] DiBEG — the launching country. http://www.dibeg.org/world/world.html.

[20] Dropcam. https://nest.com/camera/meet-nest-cam/?dropcam=true.

[21] The encounternet project. http://encounternet.net/.

[22] Energous Wattup wireless charging demo. http://www.engadget.com/2015/01/05/energous-wattup-wireless-charging-demo/.

[23] Fitbit. http://www.fitbit.com/.

[24] Gainspan GS1500M. http://www.alphamicro.net/media/412417/gs1500m_datasheet_rev_1_4.pdf.

[25] Google glass. http://www.google.com/glass/start/.

[26] HMS190BMS8 by Hittite Microwave Devices. https://www.hittite.com/content/documents/data_sheet/hmc190bms8.pdf.

[27] iBeacons. http://beekn.net/2014/04/will-apple-pull-plug-ibeacon-devices/.

[28] IEEE 802.11 standard, 2012. http://standards.ieee.org/getieee802/download/802.11-2012.pdf.

[29] ifixit teardown of fitbit flex. https://www.ifixit.com/Teardown/Fitbit+Flex+Teardown/16050.

[30] Invensense nine-axis mems motion tracking device. http://www.invensense.com/mems/gyro/documents/PS-MPU-9255.pdf.

[31] Jawebone. https://jawbone.com/fitness.

[32] LMT84 Temp Sensor by Texus Instruments. http://www.ti.com/lit/ds/symlink/lmt84.pdf.

[33] MAX2830 by Maxim. https://datasheets.maximintegrated.com/en/ds/MAX2830.pdf.

[34] MEMS digital microphone (MP45DT02) datasheet. http://www.st.com/internet/com/TECHNICAL_RESOURCES/TECHNICAL_LITERATURE/DATASHEET/DM00025467.pdf.

[35] MS412FE-FL26E micro-battery by SEIKO. http://www.sii.co.jp/compo/catalog/battery_en.pdf.

[36] MSP430FR5969 micro-controller by Texas Instruments. http://www.ti.com/lit/ds/symlink/msp430fr5969.pdf.

[37] NASA news release: A Wi-Fi reflector chip to speed up wearables. http://www.jpl.nasa.gov/news/news.php?feature=4663.

[38] Nest thermostat. https://nest.com/.

[39] Neurovista. http://www.neurovista.com.

[40] New policies for part 15 devices, FCC, TCBC workshop, 2005. https://transition.fcc.gov/oet/ea/presentations/files/may05/New_Policies_Pt._15_SD.pdf.

[41] Number of IoT devices estimate by Cisco. http://www.zdnet.com/article/the-internet-of-things-and-big-data-unlocking-the-power/.

[42] Ov7670 by OmniVision. http://www.voti.nl/docs/OV7670.pdf.

[43] Ovm 7690 camera module. http://www.ovt.com/uploads/parts/OVM7690_PB(1.0)_web.pdf.

[44] Panasonic BK-4MCCA12SA AAA Nicekl Metal Hydride Batteries. http://www.amazon.com/Panasonic-BK-4MCCA12SA-Pre-Charged-Rechargeable-Batteries/dp/B00JHKSMG8.

[45] PhantomJS. http://phantomjs.org/. Loaded January 14, 2015.

[46] Proximity sensors. https://www.ia.omron.com/products/category/sensors/proximity-sensors/.

[47] QUALCOMM Atheros 9462. http://www.qca.qualcomm.com/wp-content/uploads/2013/11/AR9462.pdf.

[48] QUALCOMM QCA4002 and QCA4004. http://www.eeworld.com.cn/zt/wireless/downloads/QCA4002-4004FIN.pdf.

[49] Rodgers RO4000 Series Specifications. http://www.rogerscorp.com/documents/726/acm/RO4000-Laminates---Data-sheet.pdf.

[50] S-882Z Series by SEIKO. http://www.eet-china.com/ARTICLES/2006MAY/PDF/S882Z_E.pdf.

[51] Samsung galaxy gear. http://www.samsung.com/us/mobile/wearable-tech/SM-V7000ZKAXAR.

[52] SMS7630-061 by Skyworks. http://www.skyworksinc.com/uploads/documents/SMS7630_061_201295G.pdf.

[53] Software defined radio hardware survey, scott johnston. http://people.bu.edu/mrahaim/NEWSDR/Presentations/NEWSDR_Johnston.pdf.

[54] Synopsis design complier. http://www.synopsys.com/Tools/Implementation/RTLSynthesis/DesignCompiler/Pages/default.aspx.

[55] TI CC2541. http://www.ti.com/lit/ds/symlink/cc2541.pdf.

[56] TI CC2650. http://www.digikey.com/product-detail/en/CC2650F128RHBR/CC2650F128RHBR-ND/5189550.

[57] TI CC3100MOD. http://www.ti.com/lit/ds/symlink/cc3100mod.pdf.

[58] Toq. http://toq.qualcomm.com/.

[59] TS 881 datasheet, STMicroelectronics, july 2012. http://www.st.com/internet/com/TECHNICAL_RESOURCES/TECHNICAL_LITERATURE/DATASHEET/DM00057901.pdf.

[60] Wattup by Energous. http://www.energous.com/overview/.

[61] Wi-Fi device shipments to surpass 15 billion by end of 2016. http://www.wi-fi.org/news-events/newsroom/wi-fi-device-shipments-to-surpass-15-billion-by-end-of-2016.

[62] WM-61A Panasonic electret microphone datasheet. http://industrial.panasonic.com/www-data/pdf/ABA5000/ABA5000CE22.pdf.

[63] ATSC digital television standard. ATSC Standard A/53, 1995.

[64] Tool for pesq analysis of speech. 2008. http://www.utdallas.edu/~loizou/speech/software.htm.

[65] Ian F Akyildiz, Tommaso Melodia, and Kaushik R Chowdhury. A survey on wireless multimedia sensor networks. Computer networks, 51(4):921–960, 2007.

[66] G Andia Vera, Apostolos Georgiadis, Ana Collado, and Selva Via. Design of a 2.45 ghz rectenna for electromagnetic (em) energy scavenging. In IEEE RWS 2010.

[67] Steven R Anton and Henry A Sodano. A review of power harvesting using piezoelectric materials (2003–2006). Smart materials and Structures, 16(3):R1, 2007.

[68] Daniel Arnitz and Matthew S Reynolds. Multitransmitter wireless power transfer optimization for backscatter rfid transponders. Antennas and Wireless Propagation Letters, IEEE, 12:849–852, 2013.

[69] James Ayers, Napong Panitantum, Kartikeya Mayaram, and Terri S Fiez. A 2.4 ghz wireless transceiver with 0.95 nj/b link energy for multi-hop battery-freewireless sensor networks. In VLSI Circuits (VLSIC), 2010 IEEE Symposium on, pages 29–30. IEEE, 2010.

[70] S.R. Best and B.C. Kaanta. A tutorial on the receiving and scattering properties of antennas. Antennas and Propagation Magazine, IEEE, 51(5):26 –37, oct. 2009.

[71] Dinesh Bharadia, Kiran Raj Joshi, and Sachin Katti. Full duplex backscatter. In Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks, 2013.

[72] Dinesh Bharadia, Emily McMilin, and Sachin Katti. Full duplex radios. In Proceedings of the ACM SIGCOMM, pages 375–386. ACM, 2013.

[73] R. Bhattacharyya, C. Floerkemeier, and S. Sarma. Low-cost, ubiquitous rfid-tag-antenna-based sensing. Proceedings of the IEEE, 98(9):1593 –1600, sept. 2010.

[74] J.L. Bohorquez, A.P. Chandrakasan, and J.L. Dawson. A 350μW CMOS MSK transmitter and 400μW OOK super-regenerative receiver for medical implant communications. Solid-State Circuits, IEEE Journal of, 44(4):1248 –1259, april 2009.

[75] C. Boyer and S. Roy. Coded qam backscatter modulation for rfid. Communications, IEEE Transactions on, 60(7):1925 –1934, july 2012.

[76] M. Buettner and D. Wetherall. A software radio-based uhf rfid reader for phy/mac experimentation. In RFID (RFID), 2011 IEEE International Conference on, pages 134–141. IEEE, 2011.

[77] Michael Buettner. Backscatter Protocols and Energy-Efficient Computing for RF-Powered Devices. PhD thesis, University of Washington, Seattle, 2012.

[78] Michael Buettner, Ben Greenstein, Alanson Sample, Joshua R. Smith, and David Wetherall. Revisiting smart dust with RFID sensor networks. In Proc. 7th ACM Workshop on Hot Topics in Networks (Hotnets-VII), October 2008.

[79] Michael Buettner, Ben Greenstein, and David Wetherall. Dewdrop: An energy-aware runtime for computational rfid. In Proceedings of NSDI'11, 2011.

[80] Ritochit Chakraborty, Sumit Roy, and Vikram Jandhyala. Revisiting rfid link budgets for technology scaling: range maximization of rfid tags. Microwave Theory and Techniques, IEEE Transactions on, 59(2):496–503, 2011.

[81] R.E. Collin. Limitations of the thevenin and norton equivalent circuits for a receiving antenna. Antennas and Propagation Magazine, IEEE, 45(2):119 – 124, april 2003.

[82] Wireless Power Consortium. Qi wireless power specification.

[83] G.A. Covic and J.T. Boys. Inductive power transfer. Proceedings of the IEEE, 2013.

[84] J-P Curty, Norbert Joehl, C Dehollaini, and Michel J Declercq. Remotely powered addressable uhf rfid integrated system. IEEE Journal of Solid-State Circuits, 2005.

[85] Daniel M Dobkin. The RF in RFID: UHF RFID in Practice. Newnes, 2012.

[86] M. Duarte and A. Sabharwal. Full-duplex wireless communications using off-the-shelf radios: Feasibility and first results. In 2010 ASILOMAR, pages 1558–1562, Nov 2010.

[87] Melissa Duarte. Full-duplex Wireless: Design, Implementation and Characterization. PhD thesis, Rice University, May 2012.

[88] D.N. Duc, H. Lee, and K. Kim. Enhancing security of EPCglobal Gen-2 RFID against traceability and cloning. Auto-ID Labs Information and Communication University, White Paper, 2006.

[89] Jan Ebert and Marian Kazimierczuk. Class e high-efficiency tuned power oscillator. Solid-State Circuits, IEEE Journal of, 16(2):62–66, 1981.

[90] J.F. Ensworth and M.S. Reynolds. Every smart phone is a backscatter reader: Modulated backscatter compatibility with bluetooth 4.0 low energy (ble) devices. In RFID, 2015 IEEE International Conference on.

[91] J.F. Ensworth, S.J. Thomas, Seung Yul Shin, and M.S. Reynolds. Waveform-aware ambient rf energy harvesting. In IEEE RFID 2014.

[92] Epc class 1 gen 2 uhf rfid standard. 2008. `http://www.gs1.org/gsmp/kc/epcglobal/uhfc1g2/uhfc1g2\_1\_2\_0-standard-20080511.pdf`.

[93] Sinem Coleri Ergen. Zigbee/ieee 802.15.4 summary, 2004.

[94] K. Finkenzeller. RFID Handbook. J. Wiley & Sons, New York, 2003.

[95] Richard Fletcher, Jeremy A. Levitan, and Joel Rosenberg. Application of smart materials to wireless id tags and remote sensors. In In Materials for Smart Systems II, Materials Research Society, pages 557–562, 1997.

[96] Simone Gambini, Nathan Pletcher, and Jan M Rabaey. Sensitivity analysis for am detectors. EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2008-31, 2008.

[97] Shyamnath Gollakota, Nabeel Ahmed, Nickolai Zeldovich, and Dina Katabi. Secure in-band wireless pairing. In USENIX Security, 2011.

[98] M. Gorlatova, P. Kinget, I. Kymissis, D. Rubenstein, X. Wang, and G. Zussman. Energy-harvesting active networked tags (EnHANTs) for ubiquitous object networking. IEEE Wireless Commun., 2010.

[99] Kate Greene. Intel's tiny wi-fi chip could have a big impact. MIT Technology review, 2012.

[100] J.D. Griffin and G.D. Durgin. Complete link budgets for backscatter-radio and rfid systems. Antennas and Propagation Magazine, IEEE, 51(2):11 –25, april 2009.

[101] Kenneth Gudan, Shuai Shao, Jonathan J Hull, Joshua Ensworth, and Matthew S Reynolds. Ultra-low power 2.4 ghz rf energy harvesting and storage system with- 25dbm sensitivity. In RFID (RFID), 2015 IEEE International Conference on, pages 40–46. IEEE, 2015.

[102] Dongning Guo and Lei Zhang. Virtual full-duplex wireless communication via rapid on-off-division duplex. CoRR, abs/1010.2667, 2010.

[103] Joseph A Hagerty, Florian B Helmbrecht, William H McCalpin, Regan Zane, and Zoya B Popovic. Recycling ambient microwave energy with broad-band rectenna arrays. IEEE Transactions on Microwave Theory and Techniques, 2004.

[104] Joseph A Hagerty, Tian Zhao, Regan Zane, and Zoya Popovic. Efficient broadband RF energy harvesting for wireless sensors. Department of Electrical and Computer Engineering, University of Colorado at Boulder, Boulder, CO, pages 80309–0425, 2005.

[105] R.C. Hansen. Relationships between antennas as scatterers and as radiators. Proceedings of the IEEE, 77(5):659 –662, may 1989.

[106] Allen M. Hawkes, Alexander R. Katko, and Steven A. Cummer. A microwave metamaterial with integrated power harvesting functionality. Applied Physics Letters, 2013.

[107] Ettus Inc. Universal software radio peripheral. http://ettus.com.

[108] Jouya Jadidian and Dina Katabi. Magnetic mimo: How to charge your phone in your pocket. MOBICOM, 2014.

[109] Mayank Jain, Jung Il Choi, Taemin Kim, Dinesh Bharadia, Siddharth Seth, Kannan Srinivasan, Philip Levis, Sachin Katti, and Prasun Sinha. Practical, real-time, full duplex wireless. In Proceedings of the 17th annual international conference on Mobile computing and networking, pages 301–312. ACM, 2011.

[110] S. Johan, Xuezhi Zeng, T. Unander, A. Koptyug, and H.-E. Nilsson. Remote moisture sensing utilizing ordinary rfid tags. In Sensors, 2007 IEEE, pages 308 –311, oct. 2007.

[111] Yoshihiro Kawahara, Xiaoying Bian, Ryo Shigeta, Rushi Vyas, Manos M. Tentzeris, and Tohru Asami. Power harvesting from microwave oven electromagnetic leakage. In UbiComp 2013.

[112] Yoshihiro Kawahara, Hoseon Lee, and Manos M. Tentzeris. Sensprout: Inkjet-printed soil moisture and leaf wetness sensor. In Ubicomp 2012.

[113] Bryce Kellogg, Aaron Parks, Shyamnath Gollakota, Joshua R. Smith, and David Wetherall. Wi-Fi backscatter: Internet connectivity for rf-powered devices. In Proceedings of the 2014 ACM Conference on SIGCOMM, 2014.

[114] Bryce Kellogg, Vamsi Talla, and Shyamnath Gollakota. Bringing gesture recognition to all devices. In NSDI, 2014.

[115] Bryce Kellogg, Vamsi Talla, Shyamnath Gollakota, and Joshua Smith. Passive wi-fi: Bringing low power to wi-fi transmissions. In Usenix NSDI, 2016.

[116] Pradeep Basappa Khannur, Xuesong Chen, Dan Lei Yan, Dan Shen, Bin Zhao, M Kumarasamy Raja, Ye Wu, Rendra Sindunata, Wooi Gan Yeoh, and Rajinder Singh. A universal uhf rfid reader ic in 0.18-μm cmos technology. Solid-State Circuits, IEEE Journal of, 43(5):1146–1155, 2008.

[117] Sukun Kim, Rodrigo Fonseca, Prabal Dutta, Arsalan Tavakoli, David Culler, Philip Levis, Scott Shenker, and Ion Stoica. Flush: a reliable bulk transport protocol for multihop wireless networks. In Proceedings of the 5th international conference on Embedded networked sensor systems, pages 351–365. ACM, 2007.

[118] John Kimionis, Aggelos Bletsas, and John N Sahalos. Bistatic backscatter radio for tag read-range extension. In RFID-Technologies and Applications (RFID-TA), 2012 IEEE International Conference on, pages 356–361. IEEE, 2012.

[119] John Kimionis, Aggelos Bletsas, and John N Sahalos. Bistatic backscatter radio for power-limited sensor networks. In Global Communications Conference (GLOBECOM), 2013 IEEE, pages 353–358. IEEE, 2013.

[120] John Kimionis, Aggelos Bletsas, and John N Sahalos. Increased range bistatic scatter radio. Communications, IEEE Transactions on, 62(3):1091–1104, 2014.

[121] L. Kleinrock and F.A. Tobagi. Packet switching in radio channels: Part I–carrier sense multiple-access modes and their throughput-delay characteristics. Communications, IEEE Trans. on, 23(12):1400–1416, 1975.

[122] Murali Kodialam and Thyaga Nandagopal. Fast and reliable estimation schemes in rfid systems. In Proceedings of the 12th annual international conference on Mobile computing and networking, pages 322–333. ACM, 2006.

[123] J.G. Koomey, S. Berard, M. Sanchez, and H. Wong. Implications of historical trends in the electrical efficiency of computing. Annals of the History of Computing, IEEE, 33(3):46 –54, march 2011.

[124] Andre Kurs, Aristeidis Karalis, Robert Moffatt, J. D. Joannopoulos, Peter Fisher, and Marin Soljacic. Wireless power transfer via strongly coupled magnetic resonances. Science, 2006.

[125] John Kymissis, Clyde Kendall, Joseph Paradiso, and Neil Gershenfeld. Parasitic power harvesting in shoes. In Wearable Computers, 1998. Digest of Papers. Second International Symposium on, pages 132–139. IEEE, 1998.

[126] J. Landt. History of rfid. IEEE Potentials, 24(4):8–11, October 2005.

[127] A. Lazarus. Remote, wireless, ambulatory monitoring of implantable pacemakers, cardioverter defibrillators, and cardiac resynchronization therapy systems: analysis of a worldwide database. Pacing and clinical electrophysiology, 30:S2–S12, 2007.

[128] T.H. Lee. The Design of CMOS Radio-Frequency Integrated Circuits. Cambridge University Press, 1998.

[129] Chieh-Jan Mike Liang, Nissanka Bodhi Priyantha, Jie Liu, and Andreas Terzis. Surviving wi-fi interference in low power zigbee networks. In Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems, pages 309–322. ACM, 2010.

[130] Mingliang Michael Liu. Demystifying switched capacitor circuits. Newnes, 2006.

[131] Vincent Liu, Aaron Parks, Vamsi Talla, Shyamnath Gollakota, David Wetherall, and Joshua R Smith. Ambient backscatter: wireless communication out of thin air. ACM SIGCOMM Computer Communication Review, 43(4):39–50, 2013.

[132] Y. Liu, C. Huang, H. Min, G. Li, and Y. Han. Digital correlation demodulator design for RFID reader receiver. In Wireless Communications and Networking Conference 2007, pages 1664–1668. IEEE.

[133] Zhen Ning Low, R.A. Chinga, Ryan Tseng, and Jenshan Lin. Design and test of a high-power high-efficiency loosely coupled planar wireless power transfer system. Industrial Electronics, IEEE Transactions on, 2009.

[134] Feng Lu, Patrick Ling, Geoffrey M Voelker, and Alex C Snoeren. Enfold: downclocking ofdm in wifi. In Proceedings of the 20th annual international conference on Mobile computing and networking, pages 129–140. ACM, 2014.

[135] Feng Lu, Geoffrey M Voelker, and Alex C Snoeren. Slomo: Downclocking wifi communication. In NSDI, pages 255–258, 2013.

[136] J.A. MacLellan, R.A. Shober, G. Vannucci, and G.A. Wright. Tag for use in a radio communication system, July 15 1997. US Patent 5,649,296.

[137] S. Mandal, L. Turicchia, and R. Sarpeshkar. A battery-free tag for wireless monitoring of heart sounds. In Wearable and Implantable Body Sensor Networks, 2009. BSN 2009. Sixth International Workshop on, pages 201 –206, june 2009.

[138] Justin Manweiler and Romit Roy Choudhury. Avoiding the rush hours: Wifi energy management via traffic isolation. In MobiSys, 2011.

[139] S. Manzari, C. Occhiuzzi, S. Nawale, A. Catini, C. Di Natale, and G. Marrocco. Polymer-doped uhf rfid tag for wireless-sensing of humidity. In RFID (RFID), 2012 IEEE International Conference on, pages 124 –129, april 2012.

[140] G. Marrocco and F. Amato. Self-sensing passive rfid: From theory to tag design and experimentation. In Microwave Conference, 2009. EuMC 2009. European, pages 001 –004, 2009.

[141] G. Marrocco, L. Mattioni, and C. Calabrese. Multiport sensor rfids for wireless passive sensing of objects :basic theory and early results. Antennas and Propagation, IEEE Transactions on, 56(8):2691 –2702, aug. 2008.

[142] J.J. Mastrototaro. The MiniMed continuous glucose monitoring system. Diabetes technology & therapeutics, 2(1, Supplement 1):13–18, 2000.

[143] Gaurav Mathur, Peter Desnoyers, Paul Chukiu, Deepak Ganesan, and Prashant Shenoy. Ultra-low power data storage for sensor networks. ACM Transactions on Sensor Networks (TOSN), 5(4):33, 2009.

[144] Robert M. Metcalfe and David R. Boggs. Ethernet: Distributed packet switching for local computer networks. Commun. ACM, 19(7):395–404, July 1976.

[145] Robert G Meyer. Low-power monolithic rf peak detector analysis. Solid-State Circuits, IEEE Journal of, 30(1):65–67, 1995.

[146] Arunesh Mishra, Shravan Rayanchu, Dheeraj Agrawal, and Suman Banerjee. Supporting continuous mobility through multi-rate wireless packetization. In HotMobile, 2008.

[147] Radhika Mittal, Aman Kansal, and Ranveer Chandra. Empowering developers to estimate app energy consumption. In MobiCom, 2012.

[148] Carlo Mutti and Christian Floerkemeier. Cdma-based rfid systems in dense scenarios: Concepts and challenges. In RFID, 2008.

[149] Saman Naderiparizi, Aaron Parks, Zerina Kapetanovic, Benajamin Ransford, and Joshua R Smith. Wispcam: A battery-free rfid camera. In IEEE RFID 2015.

[150] L.M. Ni, Dian Zhang, and M.R. Souryal. Rfid-based localization and tracking technologies. Wireless Communications, IEEE, 18(2):45 –51, april 2011.

[151] Pavel Nikitin, A.N. Parks, and Joshua R. Smith. RFID-Vox: A Tribute to Leon Theremin. In Joshua R. Smith, editor, Wirelessly powered sensor networks and computational RFID. Springer SBM, 2012.

[152] P.V. Nikitin, S. Ramamurthy, R. Martinez, and K.V.S. Rao. Passive tag-to-tag communication. In RFID, 2012.

[153] P.V. Nikitin and K.V.S. Rao. Theory and measurement of backscattering from RFID tags. Antennas and Propagation Magazine, IEEE, 2006.

[154] P.V. Nikitin and K.V.S. Rao. Antennas and propagation in uhf rfid systems. In RFID, 2008 IEEE International Conference on, pages 277 –288, april 2008.

[155] P.V. Nikitin, K.V.S. Rao, and R.D. Martinez. Differential rcs of rfid tag. Electronics Letters, 43(8):431 –432, 12 2007.

[156] I. Obeid and P.D. Wolf. Evaluation of spike-detection algorithms for a brain-machine interface application. Biomedical Engineering, IEEE Transactions on, 51(6):905 –911, june 2004.

[157] Cecilia Occhiuzzi, Stefano Cippitelli, and Gaetano Marrocco. Modeling, design and experimentation of wearable rfid sensor tag. Antennas and Propagation, IEEE Transactions on, 58(8):2490–2498, 2010.

[158] U. Olgun, C.-C. Chen, and J.L. Volakis. Design of an efficient ambient wifi energy harvesting system. IET Microwaves, Antennas Propagation, 2012.

[159] U. Olgun, Chi-Chih Chen, and J.L. Volakis. Efficient ambient wifi energy harvesting technology and its applications. In IEEE APSURSI 2012.

[160] U. Olgun, Chi-Chih Chen, and J.L. Volakis. Wireless power harvesting with planar rectennas for 2.45 ghz rfids. In URSI 2010.

[161] BP Otis, YH Chee, R Lu, NM Pletcher, and JM Rabaey. An ultra-low power mems-based two-channel transceiver for wireless sensor networks. In VLSI Circuits, 2004. Digest of Technical Papers. 2004 Symposium on, pages 20–23. IEEE, 2004.

[162] J. Pandey and B.P. Otis. A sub-100μW MICS/ISM band transmitter based on injection-locking and frequency multiplication. Solid-State Circuits, IEEE Journal of, 46(5):1049 –1058, may 2011.

[163] Joseph A Paradiso and Thad Starner. Energy scavenging for mobile and wireless electronics. Pervasive Computing, IEEE, 4(1):18–27, 2005.

[164] Sunghyun Park, Changwook Min, and Seong-Hwan Cho. A 95nw ring oscillator-based temperature sensor for rfid tags in 0.13 um cmos. In Circuits and Systems, 2009. ISCAS 2009. IEEE International Symposium on, pages 1153 –1156, may 2009.

[165] Aaron N. Parks, Angli Liu, Shyamnath Gollakota, and Joshua R. Smith. Turbocharging ambient backscatter communication. In Proceedings of the 2014 ACM Conference on SIGCOMM, 2014.

[166] Aaron N Parks, Alanson P Sample, Yi Zhao, and Joshua R Smith. A wireless sensing platform utilizing ambient rf energy. In Biomedical Wireless Technologies, Networks, and Sensing Systems (BioWireleSS), 2013 IEEE Topical Conference on, pages 154–156. IEEE, 2013.

[167] K. Penttila, M. Keskilammi, L. Sydanheimo, and M. Kivikoski. Radar cross-section analysis for passive rfid systems. Microwaves, Antennas and Propagation, IEE Proceedings, 153(1):103 – 109, feb. 2006.

[168] Dinesh Pharadia, Kiran Raj Joshi, Manikanta Kotaru, and Sachin Katti. Backfi: High throughput wifi backscatter. In Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication, 2015.

[169] V. Pillai, H. Heinrich, D. Dieska, P.V. Nikitin, R. Martinez, and K.V.S. 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. IEEE Circuits and Systems Trans. on, 54(7):1500–1512, 2007.

[170] J. Polastre, R. Szewczyk, and D. Culler. Telos: enabling ultra-low power wireless research. In Information Processing in Sensor Networks, 2005. IPSN 2005. Fourth International Symposium on, pages 364 – 369, april 2005.

[171] David M Pozar. Microwave engineering. John Wiley & Sons, 2009.

[172] J.G. Proakis. Digital Communications. Communications and signal processing. McGraw-Hill, 1995.

[173] John G Proakis and Masoud Salehi. Digital communications. 2005. McGraw-Hill, New York.

[174] X. Qing and N. Yang. A folded dipole antenna for RFID. In Antennas and Propagation Society International Symposium, 2004. IEEE, volume 1, pages 97–100. IEEE, 2004.

[175] J Rabaey, J Ammer, B Otis, F Burghardt, YH Chee, N Pletcher, M Sheets, and H Qin. Ultra-low-power design. Circuits and Devices Magazine, IEEE, 22(4):23–29, 2006.

[176] Jan M Rabaey, M Josie Ammer, Julio L da Silva, Danny Patel, and Shad Roundy. Picoradio supports ad hoc ultra-low power wireless networking. Computer, 33(7):42–48, 2000.

[177] J.M. Rabaey, J. Ammer, T. Karalar, Suetfei Li, B. Otis, M. Sheets, and T. Tuan. PicoRadios for wireless sensor networks: the next challenge in ultra-low power design. In Solid-State Circuits Conference, 2002. Digest of Technical Papers. ISSCC. 2002 IEEE International, volume 1, pages 200 –201 vol.1, 2002.

[178] Vijay Raghunathan, Aman Kansal, Jason Hsu, Jonathan Friedman, and Mani Srivastava. Design considerations for solar energy harvesting wireless embedded systems. In Proceedings of the 4th international symposium on Information processing in sensor networks, page 64. IEEE Press, 2005.

[179] Yogesh K Ramadass and Anantha P Chandrakasan. A batteryless thermoelectric energy-harvesting interface circuit with 35mv startup voltage. Institute of Electrical and Electronics Engineers, 2010.

[180] Benjamin Ransford, Jacob Sorber, and Kevin Fu. Mementos: system support for long-running computation on RFID-scale devices. SIGPLAN Not., 46(3):159–170, March 2011.

[181] K.V.S. Rao, P.V. Nikitin, and S.F. Lam. Antenna design for UHF RFID tags: A review and a practical application. Antennas and Propagation, IEEE Transactions on, 53(12):3870–3876, 2005.

[182] Behzad Razavi. RF microelectronics. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1998.

[183] Heejun Roh and Wonjun Lee. Feasibility of a low-cost tx-rx separated reader system for rf-powered computers. In Computer Communications Workshops (INFOCOM WKSHPS), 2015 IEEE Conference on, pages 51–52. IEEE, 2015.

[184] George Roussos and Vassilis Kostakos. Rfid in pervasive computing:state-of-the-art and outlook. Pervasive Mob. Comput., 5(1):110–131, February 2009.

[185] S. Roy, V. Jandhyala, J.R. Smith, D.J. Wetherall, B.P. Otis, R. Chakraborty, M. Buettner, D.J. Yeager, You-Chang Ko, and A.P. Sample. RFID: From supply chains to sensor nets. Proceedings of the IEEE, 2010.

[186] Alanson Sample and Joshua R Smith. Experimental results with two wireless power transfer systems. In Radio and Wireless Symposium, 2009. RWS'09. IEEE, pages 16–18. IEEE, 2009.

[187] 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, November 2008.

[188] R. Sarpeshkar. Analog versus digital: Extrapolating from electronics to neurobiology. Neural computation, 10(7):1601–1638, October 1988.

[189] G. Seigneuret, E. Bergeret, and P. Pannier. Auto-tuning in passive UHF RFID tags. In NEWCAS Conference (NEWCAS), 2010 8th IEEE International, pages 181–184, 2010.

[190] Souvik Sen, Romit Roy Choudhury, and Srihari Nelakuditi. CSMA/CN: Carrier sense multiple access with collision notification. In MobiCom, 2010.

[191] Joshua R Smith. Range scaling of wirelessly powered sensor systems. In Wirelessly powered sensor networks and computational RFID, pages 3–12. Springer New York, 2013.

[192] D. Smithies and F. Fietkau. minstrel: MadWiFi and Linux kernel rate selection algorithm. 2005.

[193] Jungmin So and Nitin H Vaidya. Multi-channel mac for ad hoc networks: handling multi-channel hidden terminals using a single transceiver. In Proceedings of the 5th ACM international symposium on Mobile ad hoc networking and computing, pages 222–233. ACM, 2004.

[194] Kannan Srinivasan, Prabal Dutta, Arsalan Tavakoli, and Philip Levis. An empirical study of low-power wireless. ACM Transactions on Sensor Networks (TOSN), 6(2):16, 2010.

[195] V. Talla, M. Buettner, D. Wetherall, and J. Smith. Hybrid analog-digital backscatter platform for high data rate, battery-free sensing. In Wireless Sensors and Sensor Networks (WiSNet), 2013 IEEE Topical Conference on, 2013.

[196] Vamsi Talla, Bryce Kellogg, Benjamin Ransford, Saman Naderiparizi, Shyamnath Gollakota, and Joshua R Smith. Powering the next billion devices with wi-fi. arXiv preprint arXiv:1505.06815, 2015.

[197] Vamsi Talla, Bryce Kellogg, Benjamin Ransford, Saman Naderiparizi, Shyamnath Gollakota, and Joshua R Smith. Powering the next billion devices with wi-fi. CoNEXT, 2015.

[198] Nikola Tesla. My Inventions: The Autobiography of Nikola Tesla. Hart Bros., 1982.

[199] S.J. Thomas and M.S. Reynolds. A 96 Mbit/sec, 15.5 pJ/bit 16-QAM modulator for UHF backscatter communication. In RFID (RFID), 2012 IEEE International Conference on, pages 185 –190, april 2012.

[200] S.J. Thomas, E. Wheeler, J. Teizer, and M.S. Reynolds. Quadrature amplitude modulated backscatter in passive and semipassive uhf rfid systems. Microwave Theory and Techniques, IEEE Transactions on, 60(4):1175 –1182, april 2012.

[201] Matthew S Trotter and Gregory D Durgin. Survey of range improvement of commercial rfid tags with power optimized waveforms. In IEEE RFID 2010.

[202] Matthew S Trotter, Joshua D Griffin, and Gregory D Durgin. Power-optimized waveforms for improving the range and reliability of rfid systems. In IEEE RFID 2009.

[203] M. Tubaishat and S. Madria. Sensor networks: an overview. Potentials, IEEE, 22(2):20–23, 2003.

[204] C Valenta and G Durgin. Harvesting wireless power: Survey of energy-harvester conversion efficiency in far-field, wireless power transfer systems. IEEE Microwave Magazine, 2014.

[205] H.J. Visser, A.C.F. Reniers, and J.A.C. Theeuwes. Ambient rf energy scavenging: Gsm and wlan power density measurements. In EuMC 2008.

[206] R.H. Walden. Analog-to-digital converter survey and analysis. Selected Areas in Communications, IEEE Journal on, 17(4):539 –550, apr 1999.

[207] BH Waters, AP Sample, P Bonde, and JR Smith. Powering a ventricular assist device (vad) with the free-range resonant electrical energy delivery (free-d) system. Proceedings of the IEEE 2012.

[208] E. Welbourne, L. Battle, G. Cole, K. Gould, K. Rector, S. Raymer, M. Balazinska, and G. Borriello. Building the internet of things using RFID: The RFID ecosystem experience. Internet Computing, IEEE, 13(3):48–55, may-june 2009.

[209] Lih-Chyau Wuu, Yen-Ju Chen, Chi-Hsiang Hung, and Wen-Chung Kuo. Zero-collision rfid tags identification based on cdma. In International Conference on Information Assurance and Security, 2009.

[210] Jun Yi, Wing-Hung Ki, and Chi-Ying 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, 54(1):153–166, 2007.

[211] Chen Ying and Zhang Fu-Hong. A system design for uhf rfid reader. In Communication Technology, 2008. ICCT 2008. 11th IEEE International Conference on, pages 301–304. IEEE, 2008.

[212] J.L. Zalesky and A. Wakefield. Integrating segmented electronic paper displays into consumer electronic devices. In Consumer Electronics (ICCE), 2011 IEEE International Conference on, pages 531 –532, jan. 2011.

[213] Fan Zhang, Yanqing Zhang, Jason Silver, Yousef Shakhsheer, Manohar Nagaraju, Alicia Klinefelter, Jagdish Pandey, James Boley, Eric Carlson, Aatmesh Shrivastava, et al. A batteryless 19μw mics/ism-band energy harvesting body area sensor node soc. In Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2012 IEEE International, pages 298–300. IEEE, 2012.

[214] Jiansong Zhang, Haichen Shen, Kun Tan, Ranveer Chandra, Yongguang Zhang, and Qian Zhang. Frame retransmissions considered harmful: improving spectrum efficiency using micro-acks. In Proceedings of MOBICOM 2012.

VITA

Vamsi Talla earned his Ph.D. in Electrical Engineering from the University of Washington, Seattle in March 2016. He received his M.S. in Electrical Engineering from University of Washington in August 2012 and B.Tech in Electronics and Communication Engineering from the Indian Institute of Technology, Guwahati in May 2009. Vamsi also earned a Technology Entrepreneurship Certificate (TEC) from the University of Washington Foster Business School. During his graduate study, Vamsi was a member of the Sensor Systems Lab led by Prof. Joshua R. Smith and worked closely with Prof. Shyamnath Gollakota. For a period of 6 months, Vamsi worked as an RFIC intern at Intel labs. Currently, Vamsi is a co-founder of Jeeva Wireless Inc. and will join University of Washington as a research associate.

Vamsi received best paper awards at IEEE WiSNet 2013, ACM SIGCOMM 2013 and USENIX NSDI 2016. He has also been the recipient of the Intel PhD fellowship and QUALCOMM Innovation Fellowship in 2014. Vamsi has authored several articles on near field wireless power transfer, far field wireless power, backscatter communication and backscatter based sensing. His research interests lie in circuit and system level design of wireless power transfer, energy harvesting and low power sensing and communication.