

©Copyright 2021

Mehrdad Hesar

Platforms and Technologies for Low-Power Wireless Systems

Mehrdad Hesar

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

Doctor of Philosophy

University of Washington

2021

Reading Committee:

Shyamnath Gollakota, Chair

Joshua R. Smith

Tom Anderson

Program Authorized to Offer Degree:
Computer Science and Engineering

University of Washington

Abstract

Platforms and Technologies for Low-Power Wireless Systems

Mehrdad Hesar

Chair of the Supervisory Committee:
Associate Professor Shyamnath Gollakota
Paul G. Allen School of Computer Science and Engineering

Achieving the vision of ubiquitous connectivity requires addressing multiple challenges across low-power wireless communication, network protocols and sensing. This thesis presents four technologies that address some of the challenges in achieving ubiquitous connectivity. First, the thesis describes an ultra-low-power high-definition video streaming system that enables video streaming from a low-power, wearable camera to a nearby mobile device. Second, we present LoRa backscatter that enables the first wide-area backscatter communication system to achieve long-range low-power cheap wireless communication at tens of microwatts of power. Third, we build NetScatter, the first wireless protocol that can scale to hundreds and thousands of concurrent transmissions from backscatter devices. Finally, we describe TinySDR platform, a standalone, fully programmable low power software-defined radio solution that can be duty cycled for battery operation like a real low-power sensor endpoint, and more importantly, can be programmed over the air to allow for large scale deployment.

TABLE OF CONTENTS

	Page
List of Figures	iii
List of Tables	viii
Chapter 1: Introduction	1
1.1 Low-power high data-rate IoT applications	2
1.2 Low-power long-range connectivity for low-rate applications	3
1.3 Supporting applications with large-scale backscatter devices	4
1.4 Supporting applications with programmable radio networks	5
1.5 Thesis Organization	6
Chapter 2: Low-Power High Data-Rate IoT Applications	7
2.1 Introduction	7
2.2 A Case for Our Architecture	12
2.3 System Design	13
2.4 Implementation	20
2.5 Evaluation	24
2.6 Low-Resolution Security Camera	28
2.7 Related Work	32
2.8 Limitations and Conclusion	34
Chapter 3: Low-Power Long-Range Connectivity for Low-Rate Applications	36
3.1 Introduction	36
3.2 The case for LoRa backscatter and Related Work	40
3.3 System Design	44
3.4 Hardware Implementation	55
3.5 Evaluation	58

3.6	Application Deployments	65
3.7	Discussion and Conclusion	70
Chapter 4:	Supporting Applications with Large-Scale Backscatter Devices	72
4.1	Introduction	72
4.2	CSS Primer & Existing Approaches	76
4.3	NetScatter Design	79
4.4	Evaluation	92
4.5	Related Work	101
4.6	Conclusion	102
Chapter 5:	Supporting Applications with Programmable Radio Networks	104
5.1	Introduction	104
5.2	SDR Requirements for IoT Nodes	109
5.3	TinySDR Platform	110
5.4	Case Studies: LoRa and BLE Beacons	122
5.5	Evaluation	126
5.6	Research Study: Concurrent Reception	133
5.7	Conclusion and Research Opportunities	137
Chapter 6:	Conclusion	139
Bibliography	141

LIST OF FIGURES

Figure Number	Page
2.1 Target Application. Our ultra-low-power HD streaming architecture targets wearable cameras. We achieve this by performing analog video backscatter from the wearable camera to a nearby mobile device (e.g., smartphone).	8
2.2 The amplifier, AGC, ADC and compression module consume orders of magnitude more power than is available on a low-power device. In our design, these power hungry modules have been delegated to the reader, eliminating their power consumption overhead from the wireless camera.	9
2.3 Sample HD video frame streamed with our analog video backscatter design. Our prototype was placed 4 feet from the reader.	10
2.4 Architecture of the PWM Converter.	14
2.5 CDF of effective bandwidth for zig-zag pixel scanning (a), and improvement it provides over raster scanning (b).	18
2.6 Distributed Compression. The sequence of frames and pixels transmitted by the camera to the reader.	20
2.7 HD Wireless Camera Prototype. A laptop emulating an analog output camera feeds raw pixel data into our backscatter prototype. Then the PWM encoded pixels are backscattered to the reader.	21
2.8 Signal Processing Steps at the Reader. Recovering video from a PWM backscatter signal.	23
2.9 ENOB of the Received Video. ENOB of the received video versus distance of our wireless camera prototype from the reader.	25
2.10 Effective Number of Bits (ENOB) versus video quality. The frame corresponding to an ENOB of 3 (d) shows video quality degradation.	26
2.11 ENOB of Received Video Under Different Head Motions.	27
2.12 Power Harvesting. We plot the average power harvested by the battery-free hardware over different distances from the reader.	28
2.13 Evaluation of Distributed Inter-frame Compression Algorithm for HD Video. The plots show PSNR at different compression ratio for two HD videos.	29
2.14 Prototype of Our Low-resolution, Video Streaming Home Security Camera. Image of the analog camera, FPGA digital core and pulse width modulated (PWM) backscatter, all implemented using COTS components. The overall board measures 3.5 cm by 3.5 cm by 3.5 cm.	30

2.15	Operational Range of Low-resolution Security Camera.	31
2.16	Distributed Interactive Compression with Our Low-resolution Security Camera. PSNR of the compressed video as a function of compression parameters in a typical lab setting.	32
2.17	Face Detection with Our Low-resolution Security Camera. We show the accuracy of face detection on the video streamed from our wireless camera.	33
3.1	LoRa backscatter deployment. The LoRa Backscatter device consumes $9.25 \mu\text{W}$, operates at 100s of meters and can be powered by flexible printed batteries and button cells (10 cents), a capability that cannot be achieved with radios. The RF source transmits a single tone that the backscatter device uses to synthesize CSS signals. The challenge is that at the receiver, the backscatter signal is not only drowned by noise but also suffers interference from the RF source.	37
3.2	We compare the operating distance of existing backscatter systems.	40
3.3	Communication Technologies. We show the sensitivity and supported data rates for different communication technologies and the feasibility of powering them from different sources.	40
3.4	We compare power consumption and cost of different communication technologies including backscatter and radio techniques.	41
3.5	CSS Modulation & Receiver FFT Processing. The two figures on the left shows two chirp symbols, where the second symbol is a cyclic shifted version of the first, offset by half the chirp duration. BW and SF are the bandwidth and spreading factor respectively. The time delays in the chirp translate to a peak in the FFT domain. The two plots on the right show the FFT output corresponding to the two symbols.	45
3.6	Hybrid Analog-Digital Backscatter. The digital baseband processor generates the frequency plan that is converted in the analog domain to control the output frequency of a VCO. The time shifted versions of the VCO output are mapped according to the dataset of the approximated exponential signal to their respective backscatter impedance values using a SP8T RF switch.	46
3.7	Four CSS symbols when spreading factor is 2.	47
3.8	Approximation of a cosine wave with multi-level signal. We approximate the cosine wave as a sum of three digital signals $S_0(t)$, $S_1(t)$ and $S_2(t)$ resulting in a multi-level signal.	50
3.9	LoRa Packet Structure.	53
3.10	We characterize our receiver by measuring PER as a function of reported RSSI in presence and absence of out of band interference.	58
3.11	RSSI in Deployment Scenario 1. d is the distance between the RF source and receiver. We move the backscatter device along the line joining them. This figure also shows the corresponding LoRa bit rate at which we successfully receive all our ten packets from the backscatter device without any loss.	61

3.12	RSSI in Deployment Scenario 2. d_1 (d_2) is the distance between the backscatter device and RF source (receiver). We fix the location of the backscatter device and RF source and move the receiver away from the backscatter device.	62
3.13	The plot on the left demonstrates the efficacy of our harmonic cancellation technique. The right plots CDF of RSSI of a LoRa backscatter device in the presence and absence of another LoRa backscatter device concurrently transmitting in adjacent band.	63
3.14	Office Deployment. We receive backscattered packets across one floor of an office building with $13,024 \text{ ft}^2$ (1210 m^2) area.	65
3.15	Home Deployment. We receive packets across the $4,800 \text{ ft}^2$ (446 m^2) house spread across three floors.	66
3.16	Precision Agriculture. We receive backscattered packets across a one-acre (4046 m^2) farm.	67
3.17	Smart Contact Lens and Flexible epidermal patch sensor. We show the RSSI and data rate distribution across an entire $3,328 \text{ ft}^2$ (309 m^2) atrium.	69
4.1	Large-Scale Network Deployment of Backscatter Devices. We deploy 256 backscatter devices across a floor of an office building covering multiple rooms.	73
4.2	NetScatter Overview. In traditional CSS systems, a single device uses different cyclic shifts to convey bits. In distributed CSS coding, each cyclic shift is assigned to a different backscatter device. Each device then uses the presence and absence of cyclic shift to send ‘1’ and ‘0’ bits.	74
4.3	CSS Primer. We show upchirp and downchirp symbols and FFT results of their multiplication. (a) Baseline upchirp symbol, (b) frequency shifted upchirp symbol and (c) cyclically shifted upchirp symbol.	76
4.4	Choir Approach. We evaluate FFT variation of chirp symbols when $BW = 500 \text{ kHz}$ and $SF = 9$ for both active LoRa radios and backscatter devices.	78
4.5	Bandwidth Aggregation. Here we use an aggregate bandwidth of $2BW$ but each device transmits only using BW . Upchirps with different cyclic shifts shown in different colors. Each upchirp is assigned to a device.	79
4.6	Timing Mismatch, in detecting beginning of a chirp symbol and its translation to FFT bin variation.	81
4.7	Power Adjustment for Backscatter. (a) gain normalized to maximum power as a function of Z_0 impedance and (b) switch network to support multiple power levels.	85
4.8	Normalized Power Spectrum. We show power spectrum of an upchirp multiplied by a baseline downchirp in FFT domain. This plot shows the main lobe and side lobes of a single chirp transmission. We assign devices to high and low power regions based on their power level.	86
4.9	Backscatter Devices SNR Variance. CDF of SNR variance of backscatter devices in an office environment, when people were walking around, over 30 mins.	87

4.10	NetScatter Network Association Process. We show the association process of an incoming NetScatter device (#2) to the network, while there are existing devices associated with the network (i.e., device #1).	89
4.11	Structure of AP's Query Message.	91
4.12	Near-Far BER Results. We show the effect of the second device's power on the first device's BER vs. SNR for different ratios of the second device's to first device's power with power aware cyclic shift assignments.	92
4.13	Our Backscatter Devices. They are arranged closely for this picture. They are spread out across more than ten rooms in our deployment.	93
4.14	Frequency Offset FFT Bin Variation. (a) frequency offset of backscatter devices, and (b) effect of residual time and frequency offset for different configurations.	94
4.15	Doppler Effect and Power Dynamic Range Evaluation. We evaluate (a) Doppler effect, and (b) we show power difference between two concurrent transmissions at different locations of FFT domain. One transmission is fixed and the other is sweeping across different chirp symbols.	95
4.16	Spectrogram of Backscattered Signal at the Different Power Levels.	96
4.17	Network Physical Rate. We evaluate NetScatter network physical rate and compare it with other schemes.	98
4.18	Link-layer Data Rate. We evaluate link-layer data rate for NetScatter and compare it with other schemes.	100
4.19	Network Latency. We evaluate the latency of NetScatter and compare it with other schemes. We define latency as total time for transmitting all the devices' data.	101
5.1	TinySDR Hardware Platform. It has two antenna ports for running IoT PHY and MAC protocols at 2.4 GHz and 900 MHz. This image is the actual size of the platform on printed paper.	105
5.2	Radio Module Power Consumption for Each Platform. The TX output power of each radio module is shown on top of it.	110
5.3	TinySDR System Block Diagram. A complete system diagram showing all of the components of tinySDR. This includes the software radio consisting of the radio, amplifiers, and FPGA, OTA programmer which uses a LoRa radio and flash memory to store programs, and a power management system with the flexibility to turn off power consuming components. Each of these subsystems are controlled in software running on the MCU.	112
5.4	I/Q Word Structure Used by I/Q Radio.	115
5.5	LoRa Packet Structure.	119
5.6	LoRa Implementation Block Diagrams.	120
5.7	Evaluation Testbed Map.	124
5.8	TinySDR Single-Tone Frequency Spectrum.	124

5.9	Single-Tone Transmitter Power Consumption. We show the total power consumption of tinySDR including I/Q radio, FPGA, MCU and regulators at different transmitter output power. This is 15-16 times lower power consumption than the USRP E310 embedded SDR.	128
5.10	LoRa Modulator Evaluation. We evaluate our LoRa modulator in comparison with Semtech LoRa chip.	129
5.11	LoRa Demodulator Evaluation. We evaluate our LoRa demodulator by demodulating chirp symbols at different RSSI.	130
5.12	BLE evaluation. BLE beacons at different power levels.	133
5.13	BLE Beacons Signal. We show BLE beacon transmissions on three advertising channels from tinySDR using an envelope detector.	133
5.14	OTA Programming Time. We show CDF of OTA programming time for programming LoRa and BLE implementations on tinySDR.	134
5.15	Orthogonal LoRa Demodulation Evaluation	136

LIST OF TABLES

Table Number	Page
2.1 Raw digital video sampling and bitrate requirement	11
2.2 Intra-frame compression for average/worst-case scenarios across 100 videos.	13
2.3 Power Consumption	22
3.1 Comparison of LoRa Backscatter with Existing Backscatter Systems. d_1 is the distance between the backscatter device and the RF source. d_2 is the distance between the backscatter device and the receiver.	43
4.1 NetScatter Different Modulation Configurations, with maximum time/freq. mismatch that can be tolerated.	84
5.1 Comparison Between Different SDR Platforms. Costs are based on sale prices for commercial products without a public bill of materials (BOM) and published BOM prices for research prototypes. OTA refers to over-the-air programming capabilities.	106
5.2 Existing Off-the-Shelf I/Q Radio Modules.	113
5.3 Power Domains in TinySDR.	118
5.4 Different Operation Timing for TinySDR.	125
5.5 TinySDR Cost Breakdown for 1000 Units.	127
5.6 FPGA Utilization for LoRa Protocol.	131

ACKNOWLEDGMENTS

This dissertation is the result of support and encouragements of many people who supported me during my Ph.D. and I am very grateful to have them in my life.

My advisor, Shyam Gollakota, was my primary support and help from my first day. Shyam always has been a great advisor to help me to find the right research problem and steps to solve the problem. Since the first project that we worked on together until the very last one he was there with me patiently as an advisor to guide me to be an independent thinker and also as a peer to teach me collaborations on solving large research projects. I would like to specially thank him for being flexible on the projects that I selected to work on to develop my skills both as a researcher and also as an engineer.

I would like to thank my committee members for their interest in my work and their advises. Specially, I would like to thank Prof. Josh Smith who I collaborated with on two exciting projects.

Without the support of my labmates, this dissertation would not have been possible. Vamsi and Bryce, it was fun to work with you on LoRa Backscatter project and doing experiments in unusual places. Vikram was a great labmate and an awesome friend. He was always patient to teach me working with new tools and solve my hardware questions. Also, I want to thank Justin and Anran for the great time that we have in the lab, coffee breaks and amazing courses that we took together. Finally, I want to specially thank Ali who I had honor to collaborate with on multiple projects. Working with Ali was always fun and changed my Ph.D. years to be amazing! In addition to my labmates, I want to thank members of *Sensor Systems Lab*. Specially, I want to thank Saman Naderiparizi and Ali Saffari from Sensor Systems Lab who I worked with on battery-free camera project and I want to thank Mohammad Katanbaf for answering my hardware questions.

My Ph.D. experience would have not been this perfect without the experience of working with

my mentors during my internships. First, I want to thank Dr. Krishna Chintalapudi for his mentorship during my internship at Microsoft Research Lab. Working with him gave me a different perspective in my Ph.D. studies. I had an opportunity to work as an intern with Dr. Thierry Moreau at OctoML. This was a great experience for me to become exposed to a complete different research area and it opened up new opportunities. Finally, I would like to thank Dr. Aakanksha Chowdhery and Dr. Michael B. Crouse for their mentorship during my internship at Google. It was a great experience to work on a Google product and bring new capabilities to it under their supervision.

I wish to express my deepest gratitude to my parents and family members who always encouraged me to pursue my dreams. I want to thank my mom who always scarified herself to bring so much to my life and my dad who showed me to aim higher. I want to specially thank my brother who strongly supported my in every decision that I made.

And finally I like to thank all of my friends in Seattle who were like a family to me in the past five years. In the last five years I did not have a chance to visit my family in my home country. It was hard to keep up with all the struggles and my amazing friends turned Seattle to be my new home.

DEDICATION

To my parents for their love and support,
and to my sister, Farkhondeh

Chapter 1

INTRODUCTION

Wireless connectivity is increasingly underpinning much of our surrounding services and infrastructure. Internet of Things (IoT) target to connect everyday objects to the internet and enable these devices to share and process essential information. This vision could be transformative for smart agriculture, medical technologies and industrial sensing.

The vision of connecting billions of devices around the world is becoming more realistic thanks to the ubiquity of wireless networks. However, IoT wireless networks have four key requirements. First, the power consumption of IoT devices is a major challenge. Each IoT device should be low-power and be able to run on a small battery for a long time. This is important since the cost of maintaining an IoT network running at large scale is high. The IoT network requires to support long range communication. In addition, we need to be able to support a dense network of devices. Finally, the cost of an ideal IoT network should be low so that it can be adapted by various industries.

Existing communication technologies cannot satisfy all these requirements for IoT networks. To tackle this challenge for wireless communication in IoT applications, we take a systematic approach to identify the power-hungry parts of the system in different IoT applications and innovate to optimize power consumption. In addition, IoT protocols require testing and deployment at large scale to demonstrate their effectiveness. However, the research community is constrained by the lack of an inexpensive IoT platform that provides enough flexibility to develop any low-power IoT protocol and deployment of large scale networks. In this thesis, we divide IoT applications into multiple categories and explain our high-level motivations and technical solutions.

1.1 Low-power high data-rate IoT applications

There has been recent interest in wearable cameras for applications ranging from lifecasting, video blogging and live streaming concerts, political events and even surgeries [84]. Unlike smartphones, these wearable cameras must be both ultra-lightweight and cause no overheating during continuous operation. This has resulted in a trade-off between the usability of the device and its streaming abilities since higher resolution video streaming requires a bigger and heavier battery as well as power-consuming communication and processing units. For example, Snap Spectacle, while lightweight and usable, cannot stream live video [23] and can record only up to one hundred 10-second videos (effectively less than 20 minutes) [23, 19] on a single charge.

The question we ask is: *can we design a low-power camera that can perform HD video streaming to a nearby mobile device such as a smartphone?* A positive answer would enable a wearable camera that is lightweight, streams high quality video, and is safe and comfortable to wear. Specifically, reducing power consumption would reduce battery size, which in turn addresses the key challenges of weight, battery life and overheating.

We present a novel architecture that enables HD video streaming from a low-power, wearable camera to a nearby mobile device. Instead of independently optimizing the imaging and communication modules, we jointly design these components to reduce the power consumption. We present an analog video backscatter technique that feeds analog pixels from the photo-diodes directly to the backscatter hardware, thereby eliminating power-consuming hardware components, such as ADCs and codecs from the wearable camera and shift these power-hungry operations to the reader.

To evaluate our design, we simulate an ASIC, which achieves 60 fps 720p and 1080p HD video streaming for 321 μW and 806 μW , respectively. Our empirical results also show that we can harvest sufficient energy to enable battery-free 30 fps 1080p video streaming at up to 8 feet. Because no HD camera currently provides access to its raw pixel voltages, we connect the output of a DAC converter to our backscatter hardware to emulate the analog camera and stream raw analog video voltages to our backscatter prototype. Our system successfully backscatters 720p HD video at 10 fps up to 16 feet. Finally, we implement a proof-of-concept prototype of our analog

backscatter design on an ultra-low-power FPGA platform and a custom implementation of the backscatter module that backscatters video in 112×112 resolution in real-time.

1.2 Low-power long-range connectivity for low-rate applications

Active radio technologies including Wi-Fi, ZigBee, SigFox [64], LoRa [46] and LTE-M [61] provide reliable coverage and long ranges but are power consuming and cost at least 4–6 dollars [55, 67]; making them too expensive for embedding into objects at scale. Further, given their high peak current and power requirements, active radios significantly deteriorate battery life and are incompatible with emerging small, flexible and ultra-thin printed batteries [57, 53, 58] that promise innovative applications across healthcare, wearable devices and cosmetics [71, 139, 135].

Backscatter promises to be an extremely low power, smaller and cheaper alternative to active radios. Given the absence of expensive radio analog components including RF oscillators, decoupling capacitors and crystals, backscatter designs including passive RFID and Wi-Fi backscatter [120, 114] cost only a few cents to manufacture at scale [45]. We question the conventional wisdom that backscatter is a *short-range* system. Specifically, we ask if *one can achieve wide-area backscatter communication with a range of hundreds of meters, if not kilometers*. A positive answer would give us the best of both worlds: long-range reliable communication capabilities of radios at the low-power and cost of backscatter hardware. This enables, wide area relatively lower data-rate connectivity for everyday objects and opens applications in domains like smart cities [29], precision agriculture [35], industrial, medical and whole-home sensing [114], where backscatter is currently infeasible.

We presents LoRa Backscatter which adapts Chirp Spread Spectrum (CSS) coding used by LoRa protocol. LoRa uses CSS modulation where a ‘0’ bit can be represented as a continuous chirp that increases linearly with frequency, while a ‘1’ bit is a chirp that is cyclically shifted in time. Thus, CSS requires continuously changing the frequency as a function of time, which has not been demonstrated on backscatter hardware. we design the first backscatter design that can synthesize continuous frequency modulated chirps, while consuming as low as $9.25 \mu W$.

Our design can successfully backscatter from any location between an RF source and receiver,

separated by 475 m, while being compatible with commodity LoRa hardware. Further, when our backscatter device is co-located with the RF source, the receiver can be as far as 2.8 km away. We deploy our system in a 4,800 ft^2 house spread across three floors, a 13,024 ft^2 office area covering 41 rooms, as well as a one-acre vegetable farm and show that we can achieve reliable coverage, using only a single RF source and receiver. We also build a contact lens prototype as well as a flexible epidermal patch device attached to the human skin. We show that these devices can reliably backscatter data across a 3,328 ft^2 room.

1.3 Supporting applications with large-scale backscatter devices

We have shown above that we can support long-range low-power communication for IoT devices using LoRa Backscatter technology. However, LoRa Backscatter and prior designs are evaluated in a network of 1–2 devices [159, 187]. Our question here is *Can we design a network protocol that enables these low-power backscatter networks to support hundreds to thousands of concurrent transmissions?* This is challenging because the resulting design must operate reliably with weak backscatter signals that can be close to or below the noise floor.

We present NetScatter, the first wireless protocol that scales to hundreds of concurrent transmissions from backscatter devices. Our key innovation is a distributed coding mechanism, Distributed Chirp Spread Spectrum (DCSS), that works below the noise floor, operates on backscatter devices and can decode all the concurrent transmissions at the receiver using a single FFT operation. DCSS is a combination of CSS modulation and ON-OFF keying. In existing CSS systems the AP transmits a continuous wave signal which each device backscatters and encodes bits using different cyclic shifts of a chirp signal. In contrast, in DCSS, we assign a different cyclic shift of the chirp to each of the concurrent devices. Each device uses ON-OFF keying over these cyclic shifted chirps to convey bits, i.e., the presence and absence of the corresponding cyclic shifted chirp correspond to a ‘1’ and ‘0’ bit respectively.

In comparison to CSS coding where each device transmits $\log_2 N$ bits using N cyclic shifts, our distributed design enables N concurrent devices, each of which transmits a single bit, using ON-OFF keying. Thus, our design transmits a total of N bits within a chirp duration, which results

in a theoretical gain of $\frac{N}{\log_2 N}$. Our design addresses practical issues such as timing/frequency synchronization as well as the near-far problem. We deploy our design using a testbed of backscatter hardware and show that our protocol scales to concurrent transmissions from 256 devices using a bandwidth of only 500 kHz. Our results show throughput and latency improvements of 14–62x and 15–67x over existing approaches and 1–2 orders of magnitude higher transmission concurrency.

1.4 Supporting applications with programmable radio networks

While, NetScatter was an exciting work that led us to a major challenge in IoT network deployment, we learned that building an IoT network with a custom protocol is challenging. This is important because in the last few years we have seen development of numerous wireless protocols for IoT devices. In addition to longtime standards such as Bluetooth and Zigbee, a number of new protocols including LoRa, Sigfox and NB-IoT have been developed that achieve long range communication and new MIMO techniques has been introduced for IoT endpoints [99]. Due to the lack of a de-facto standard, this space remains an active area of research for both industry and academia.

Ideally, we would like a large scale IoT network testbed with the flexibility to run *any* IoT protocol at the PHY and MAC layers. Further, since many of these IoT testbeds can span hundreds of endpoints across a large campus or even a city, we need the ability to push changes to the PHY and MAC layers, using simple over-the-air software updates. This would allow for performance comparisons on a single testbed to investigate the trade-offs between existing standards as well as showcase the advantages of an entirely new custom protocol. Moreover, to make such a system representative of real-world deployments, individual network nodes should model the constraints of IoT endpoints. Specifically, these devices should have appropriate power controls and options to duty cycle transmissions, have an ultra-low power sleep mode and also support interfaces to connect sensors.

We introduce TinySDR [107], the first software-defined radio platform tailored to the needs of power-constrained IoT endpoints. TinySDR provides a standalone, fully programmable low power software-defined radio solution that can be duty cycled for battery operation like a real IoT endpoint, and more importantly, can be programmed over the air to allow for large scale

deployment. We present extensive evaluation of our platform showing it consumes as little as 30 μ W of power in sleep mode. We present two case studies by implementing LoRa and BLE beacons on the platform and achieve sensitivities of -126 dBm and -94 dBm respectively while consuming 11% and 3% of the FPGA resources.

1.5 Thesis Organization

The rest of this dissertation is organized as follows. In Chapter 2, we explain our battery-free video streaming system by explaining our analog backscatter approach in details. We show how we redesign video streaming systems by moving the power-hungry component of the system from video streaming device to the base station to increase operation time of the video streaming device. Then we explain our deployments in different scenarios and report power consumption evaluation of our system.

In Chapter 3, we describe LoRa Backscatter technology in details. Specifically, we show how we synthesize CSS on a backscatter device to achieve long-range low-power communication as well as the several deployments and evaluation of LoRa Backscatter. In Chapter 4, we explain our Distributed CSS mechanism and compare it we CSS coding to show how we can achieve concurrent transmissions from backscatter devices. Then we explain two main challenges with NetScatter, timing/frequency synchronization and near-far problem and provide our solutions in details.

Finally, in Chapter 5 we explain TinySDR. In this chapter, we explain our systematic approach to design TinySDR hardware, over-the-air programming mechanism and our power management system to reduce power consumption in different SDR operations.

Chapter 2

LOW-POWER HIGH DATA-RATE IOT APPLICATIONS

2.1 Introduction

There has been recent interest in wearable cameras like Snap Spectacles [23] for applications ranging from lifecasting, video blogging and live streaming concerts, political events and even surgeries [84]. Unlike smartphones, these wearable cameras have a spectacle form-factor and hence must be both ultra-lightweight and cause no overheating during continuous operation. This has resulted in a trade-off between the usability of the device and its streaming abilities since higher resolution video streaming requires a bigger (and heavier) battery as well as power-consuming communication and processing units. For example, Snap Spectacle, while lightweight and usable, cannot stream live video [23] and can record only up to one hundred 10-second videos (effectively less than 20 minutes) [23, 19] on a single charge.

In this paper, we ask the following question: Can we design a low-power camera that can perform HD video streaming to a nearby mobile device such as a smartphone? A positive answer would enable a wearable camera that is lightweight, streams high quality video, and is safe and comfortable to wear. Specifically, reducing power consumption would reduce battery size, which in turn addresses the key challenges of weight, battery life and overheating. Finally, since users typically carry mobile devices like smartphones that are comparatively not as weight and power constrained, they can relay the video from camera to the cloud infrastructure.

To understand this challenge, let us look at the different components of a video-streaming device: image sensor, video compression and communication. Image sensors have an optical lens, an array of photo-diodes connected to amplifiers, and an ADC to translate analog pixels into digital values. A video codec then performs compression in the digital domain to produce compressed video, which is then transmitted on the wireless medium. Existing approaches optimize the camera






Figure 2.1: **Target Application.** Our ultra-low-power HD streaming architecture targets wearable cameras. We achieve this by performing analog video backscatter from the wearable camera to a nearby mobile device (e.g., smartphone).

and communication modules individually to minimize power consumption. However, designing an efficient video streaming device requires power-consuming hardware components and video codec algorithms that interface to the camera and communication modules. Specifically, optical lens and photo-diode arrays can be designed to consume as little as $1.2 \mu\text{W}$ [104]. Similarly, recent work on backscatter can significantly lower the power consumption of communication to a few microwatts [120, 114] using custom ICs. Interfacing camera hardware with backscatter, however, requires ADCs and video codecs that significantly add to power consumption.

Table 2.1 shows the sampling rate and data rate requirements for the ADC and video codec, respectively. HD video streaming requires an ADC operating at a high sampling rate of more than at least 10 MHz. While the analog community has reduced ADC power consumption at much lower sampling rates [169, 91], state-of-the-art ADCs in the research community consume at least a few milliwatts at the high sampling rates [140]. Additionally, the high data rate requires the oscillator and video codec to run at high clock frequencies, which proportionally increases power consumption. Specifically, video codecs at these data rates consume 100s of milliwatts to a few watts of power [3].



((a)) Conventional camera design



((b)) Our camera design approach

Figure 2.2: The amplifier, AGC, ADC and compression module consume orders of magnitude more power than is available on a low-power device. In our design, these power hungry modules have been delegated to the reader, eliminating their power consumption overhead from the wireless camera.

We present a novel architecture that enables video streaming on low-power devices. Instead of independently optimizing the imaging and the communication modules, we jointly design these components to significantly reduce system power consumption. Our architecture, shown in Fig. 2.2(b), takes its inspiration from the Great Seal Bug [7], which uses changes in a flexible metallic membrane to reflect sound waves. Building on this idea, we create the first “analog” video backscatter system. At a high level, we feed analog pixels from the photo-diodes directly to the backscatter antenna; we thus eliminate power-hungry ADCs, amplifiers, AGCs and codecs.

Our intuition for an analog video backscatter approach is to shift most of the power-hungry, analog-to-digital conversion operations to the reader. Because analog signals are more susceptible to noise than digital ones, we split the ADC conversion process into two phases, one performed at the video camera, and one accomplished at the reader. At the video camera, we convert analog pixel voltage into a pulse that is discrete in amplitude but continuous in time. This signal is sent via backscatter from the camera to the reader. Avoiding the amplitude representation in the wireless link provides better noise immunity. The reader measures the continuous length pulse it receives to produce a binary digital value. Philosophically, the continuous-time, discrete amplitude pulse



Figure 2.3: **Sample HD video frame streamed with our analog video backscatter design.** Our prototype was placed 4 feet from the reader.

representation used in the backscatter link resembles that used in an extremely power-efficient biological nervous system, which encodes information in spikes that do not vary in amplitude but are continuous in time [168].

Specifically, our design synthesizes three key techniques. First, we show how to interface pixels directly to the backscatter hardware without using ADCs. To do, this we transform analog pixel values into different pulse widths using a passive ramp circuit and map these pulses into pixels at the reader. Second, we achieve intra-frame compression by leveraging the redundancy inherent in typical images. Intuitively, the signal bandwidth is proportional to the rate of change across adjacent pixels; since videos tend to be redundant, the bandwidth of the analog signal is inversely proportional to the redundancy in the frame. Thus, by transmitting pixels consecutively, we can implicitly perform compression and significantly reduce wireless bandwidth. Finally, to achieve inter-frame compression, we design a distributed algorithm that reduces the data the camera transmits while delegating most inter-frame compression functionality to the reader. At a high level, the camera performs averaging over blocks of nearby pixels *in the analog domain* and transmits these averaged values using our backscatter hardware. The reader compares these averages with those from the previous frame and requests only the blocks that have seen a significant change in the average pixel value, thus reducing transmission between subsequent video frames.

We implement a proof-of-concept prototype of our analog backscatter design on an ultra-low-

Table 2.1: Raw digital video sampling and bitrate requirement

Video Quality	Frame Rate: 60 fps		Frame Rate: 30 fps		Frame Rate: 10 fps	
	Sampling	Data Rate	Sampling	Data Rate	Sampling	Data Rate
	Rate (MHz)	(Mbps)	Rate (MHz)	(Mbps)	Rate (MHz)	(Mbps)
1080p (1920x1080)	124.4	995.3	62.2	497.7	20.7	165.9
720p (1280x720)	55.3	442.4	18.4	221.2	9.2	73.7
480p (640x480)	18.4	147.4	9.2	73.7	3.1	24.58
360p (480x360)	10.4	82.9	5.2	41.5	1.7	13.8

power FPGA platform and a custom implementation of the backscatter module. Because no HD camera currently provides access to its raw pixel voltages, we connect the output of a DAC converter to our backscatter hardware to emulate the analog camera and stream raw analog video voltages to our backscatter prototype. Fig. 2.3 shows a sample frame from an HD video streamed with our backscatter camera. More specifically, our findings are:

- We stream 720p HD video at 10 frames per second up to 16 feet from the reader. The Effective Number of Bits (ENOB) received for each pixel at distances below six feet exceeds 7 bits. For all practical purposes, these results are identical to the quality of the source HD video.
- Our inter and intra-frame compression algorithms reduces total bandwidth requirements by up to two orders of magnitude compared to raw video. For example, for 720p HD video at 10 fps, our design uses a wireless bandwidth of only 0.98 MHz and 2.8 MHz in an average-case and worst-case scenario video, respectively.

We design and simulate an ASIC implementation for our system, taking into account power consumption for the pixel array. Our results show that power consumption for video streaming at 720p HD is 321 μ W and 252 μ W for 60 and 30 fps, respectively. Power consumption at 1080p full-HD is 806 μ W and 561 μ W at 60 and 30 fps, respectively. We run experiments with RF power harvesting from the reader which shows that we can support 1080p full-HD video streaming at 30 fps up to distances of 8 feet from the reader. *Our results demonstrate that we can eliminate bat-*

teries and achieve HD video streaming on battery-free devices using our analog video backscatter approach.

2.2 A Case for Our Architecture

Fig. 2.2(a) shows the architecture of a traditional wireless camera. Photodiodes' output is first amplified by a low noise amplifier (LNA) with automatic gain control (AGC). The AGC adjusts the amplifier gain to ensure that the output falls within the dynamic range of the analog to digital converter (ADC). Next, the ADC converts the analog voltage into discrete digital values. The video codec then compresses these digital values, which are transmitted on the wireless communication link.

Unfortunately, this architecture cannot be translated to an ultra-low-power device. Although camera sensors consisting of an array of photodiodes have been shown to operate on as low as $1.2 \mu\text{W}$ of power [104] at 128×128 resolution, amplifiers, AGC, ADC and the compression block require orders of magnitude higher power. Furthermore, power consumption increases as we scale the camera's resolution and/or frame rate. A low resolution 144p video recorded at 10 fps requires an ADC sampling at 368 KSPS to generate uncompressed data at 2.95 Mbps. With the latest advancements in ADCs [137] and backscatter communication [120], uncompressed data can be digitally recorded using low-power ADCs and transmitted using digital backscatter while consuming below $100 \mu\text{W}$ of power, which is within the power budget of energy harvesting platforms. However, as we scale the resolution to HD quality and higher, such as 1080p and 1440p, the ADC sampling rate increases to 10-100 MHz, and uncompressed data is generated at 100 Mbps to Gbps. ADCs operating at such high sampling rates consume at least a few mW [140]. Further, a compression block that operates in real time on 100 Mbps to 1 Gbps of uncompressed video consumes up to one Watt of power [3]. This power budget exceeds by orders of magnitude that available on harvesting platforms. Thus, while existing architectures might operate on harvested power for low-resolution video by leveraging recent advancements in low-power ADCs and digital backscatter, as we scale the resolution and/or the frame rate of the wireless camera, these architectures' use of ADCs and compression block drives up power consumption to levels beyond the realm of

Table 2.2: Intra-frame compression for average/worst-case scenarios across 100 videos.

Video	Frame Rate: 60 fps			Frame Rate: 30 fps			Frame Rate: 10 fps		
	Raw	Data	Analog BW	Raw	Data	Analog BW	Raw	Data	Analog BW
Resolution	Rate (Mbps)		(MHz)	Rate (Mbps)		(MHz)	Rate (Mbps)		(MHz)
1080p	995.3		10.8/30.5	497.6		5.6/18.7	165.8		1.9/9.3
720p	442.3		6.3/15.6	221.1		3.1/8	73.7		0.98/2.8
480p	147.4		3/6.7	73.7		1.6/3.3	24.5		0.53/1.4
360p	82.9		1.9/3.9	41.4		0.94/2.3	13.8		0.32/0.78

battery-free devices.

2.3 System Design

In the rest of this section, we first describe our analog video transmission scheme followed by the intra-frame compression and finally the interactive distributed inter-frame compression technique.

2.3.1 Analog Video Backscatter

Problem. At a high level, our design eliminates the ADC on the camera to significantly reduce power consumption. This limits us to working with analog voltage output from the photodiodes. A *naive* approach would leverage existing analog backscatter technique used in wireless microphones [180] to implement an analog backscatter camera. In an analog backscatter system, sensor output directly controls the gate of a field-effect transistor (FET) connected to an antenna. As the output voltage of the sensor varies, it changes the impedance of the FET which amplitude modulates the RF signal backscattered by the antenna. The reader decodes sensor information by demodulating the amplitude-modulated backscattered RF signal. However, this approach cannot be translated to a camera sensor. Photodiode output has a very limited dynamic range (less than 100 mV under indoor lighting conditions). These small voltage changes map to a very small subset of radar cross-sections at the antenna [86]. As a result, the antenna backscatters a very weak signal. Since wireless channel and receivers add noise, this approach results in low SNR signal at the




Figure 2.4: Architecture of the PWM Converter.

receiver, which limits the system to both poor signal quality and limited range. One can potentially surmount this constraint by introducing a power-hungry amplifier and AGC, but this would negate the power-saving potential of analog backscatter.

Our solution. Instead of using linear amplifiers to amplify the analog video stream, we use comparators to pulse width modulate [116] camera sensor’s analog output and convert into digital domain as shown in Fig. 2.4. At a high level, pulse width modulation (PWM) converts analog input voltage to different pulse widths. Specifically, the output of PWM modulation is a square wave, where the duty cycle of the output square wave is proportional to the analog voltage of the input signal.

As we show next, a PWM converter can be implemented with passive RC components and a comparator, thereby consuming very low power. Fig. 2.4 shows the PWM converter architecture. The input is a square wave operating at frequency f , which is determined by the camera’s frame rate and resolution. First, the square wave is low-pass filtered by an RC network to generate a triangular waveform, as shown in the figure. This waveform is then compared to the pixel value using a comparator. The comparator outputs a zero when the triangular signal is less than the pixel value and a one otherwise. Thus, the pulse width generated changes with the pixel value: lower pixel values have a larger pulse duration, while higher pixel values have a smaller pulse duration. We choose the minimum and maximum voltages for the triangular signal to ensure that the camera’s pixel output is always within this limit which also maximizes the gain (analog voltage to digital pulse width) of the PWM modulation operation.

In the frequency domain, in addition to the desired baseband signal, PWM modulation also

generates additional harmonics which contain redundant copies of the input analog video signal [116, 176] and can result in electromagnetic interference (EMI) from the antenna at higher frequencies. However, since the typical power re-radiated by the antenna in a backscatter system is significantly lower than the power transmitted by an active radio, the resulting EMI from the harmonics of the baseband PWM signal in a backscatter system are also negligible (less than <-20 dBm in all scenarios). There are different ways one can suppress the minuscule EMI generated by harmonics of baseband PWM signal. A PWM conversion with full dynamic range generates harmonics which are about 4 dB lower than the desired baseband signal. However, if we increase the dynamic range by a factor of 2, the PWM harmonics reduce to 10 dB lower than the baseband. So, we can tradeoff harmonic suppression for slightly reduced performance. This is especially useful since EMI due to harmonics is a potential issue when the camera is close to the reader where it can re-radiate higher power EMI. However, at close distances we can safely tradeoff a bit of performance for EMI rejection. At farther distances, since the backscattered power is low, EMI is not an issue and we don't need to tradeoff any performance. Finally, in addition to performance tradeoff, we can introduce additional low power pass filter after the sub-carrier modulation operation in the digital domain to suppress harmonics and transmit filtered signal using backscatter. However, the implementation of EMI suppression is beyond the scope of this work.

The final component in our hardware design is sub-carrier modulation, which addresses the problem of self-interference. Specifically, in addition to receiving the backscatter signal, the receiver also receives a strong interference from the transmitter. Since the sensor data is centered at the carrier frequency, the receiver cannot decode the backscattered data in the presence of a strong, in-band interferer. Existing backscatter communication systems, such as RFID and Passive Wi-Fi, address this problem using sub-carrier modulation. These systems use a sub-carrier to shift the signal from the carrier to a frequency offset Δf from the carrier frequency. The receiver can then decode the backscattered signal by filtering out of band interference. Another consideration in choosing sub-carrier frequency is to avoid aliasing; sub-carrier frequency should not be smaller than the effective bandwidth of the analog signal.

Our PWM-based design integrates subcarrier modulation. We implement this modulation with

a simple XOR gate. The sub-carrier can be approximated by a square wave operating at Δf frequency. We input sub-carrier and PWM output to an XOR gate to up-convert the PWM signal to a frequency offset Δf . Sub-carrier modulation addresses the problem of self-interference at the reader, and as a result, the PWM backscattering wireless camera can now operate at a high SNR and achieve broad operating ranges. We show in §2.5.1 that our PWM backscatter wireless camera can operate at up to 16 feet for a 10 fps monochrome video signal in HD resolution. We also show in §2.6 that our camera system can work at up to 150 feet for a 50 KHz video signal in 112×112 resolution, over a $4\times$ improvement relative to the 3 kHz bandwidth analog backscatter wireless microphone [182]. Subcarrier modulation with a square wave generates high order harmonics copies of the baseband PWM signal but these higher order components can be eliminated via harmonic cancellation techniques introduced in prior work [178].

2.3.2 *Intra-Frame Compression*

There is significant redundancy in the pixel values of each frame of uncompressed video. Redundancy occurs because natural video frames usually include objects larger than a single pixel, which means that the colors of nearby pixels are highly correlated. At the boundaries of objects (edges), larger pixel variations can occur; conversely, in the interior of an object, the amount of pixel variation is considerably less than the theoretical maximum. The net result of pixel correlations is a reduction in the information needed to represent natural images to levels far below the worst-case maximum. Traditional digital systems use a video codec to reduce pixel value redundancy. Specifically, the ADC first digitizes the camera’s raw output at the Nyquist rate determined by the resolution and frame rate. The raw digital data stream is then fed to the video codec, which implements compression algorithms.

In the absence of the ADC, our wireless camera transmits analog video directly. However, we note that the bandwidth of any analog signal is a function of the new information contained in the signal. Inspired by analog TV broadcast, which transmits pixel information in a raster scan (left to right), we introduce and implement a zig-zag pixel scanning technique: pixels are scanned from left to right in odd rows and from right to left in even rows. The intuition here is that neighboring

pixels have less variation, and the resulting signal would thus occupy less bandwidth.

We evaluate how well zig-zag transmission performs in terms of bandwidth reduction of the PWM modulated analog signal. We download one hundred 60 fps Full-HD (1080p) videos from [24] to use as our baseline. These videos range from slow to fast moving and contain movie action stunts, running animals, racing, etc. To create video baselines at different resolutions and frame rates, we resize and subsample these Full-HD video streams. For each of the resulting video resolutions and frame rates, we create an analog video by zig-zagging the pixels as described above. We then apply low-pass filters with different bandwidths on this analog signal and report the minimum bandwidth of the baseband component of the PWM modulated analog video stream at which the PSNR of the resulting signal exceeds 30 dB. The bandwidth requirement reported in Table 2.2 shows the average/worst-case scenario, i.e., the bandwidth that ensures the recovery of average/worst-case videos with minimal quality degradation. We outline the uncompressed digital data rate for reference. Compared to the digital uncompressed data sampled at the Nyquist rate, the analog signal occupies 17.7–32.6x less bandwidth for the worst-case and 43–92x for the average-case, demonstrating our intra-frame compression technique’s capability. Fig. 2.5(a) shows the CDF of effective bandwidth for our 30 fps 720p video dataset and demonstrates that an average-case 30 fps 720p video achieves up to a $71\times$ improvement compared to raw digital video transmission.

Finally, we note that, compared to raster, a zig-zag scan faces less discontinuity in pixel values: instead of jumping from the last pixel in a row to the first pixel in the next row, it continues at the same column in the next row, thereby taking greater advantage of the video’s inherent redundancy. This further lowers bandwidth utilization in the wireless medium. As an example, on average, zig-zag pixel scanning occupies $\sim 120\text{KHz}$ and $\sim 60\text{KHz}$ less bandwidth than raster scanning in a 60 fps and 30 fps 720p video stream, respectively. Fig. 2.5(b) shows the CDF of bandwidth improvement for zig-zag scanning over raster in our one hundred 30 fps 720p videos dataset. The plot makes clear that use of zig-zag pixel scanning provides greater bandwidth efficiency than its raster counterpart.




Figure 2.5: CDF of effective bandwidth for zig-zag pixel scanning (a), and improvement it provides over raster scanning (b).

2.3.3 Distributed Inter-Frame Compression

In addition to the redundancy within a single frame, raw video output also has significant redundancy between consecutive frames. Existing digital architectures use a video codec to compress the raw video prior to wireless transmission to remove inter-frame redundancy and reduce power and bandwidth. Our approach to address this challenge is to once again leverage the reader. Like backscatter communication, where we move power-hungry components (such as ADCs) to the reader, we also move compression functionality to the reader. Specifically, our design distributes the compression algorithm between the reader and the camera. We delegate power-hungry computation to the reader and leverage the fact that our camera system can transmit *super-pixels*. A super-pixel value is the average of a set of adjacent pixel values. A camera frame consists of an $N \times N$ array of pixels, which can be divided into smaller sections of $n \times n$ pixels. A super-pixel corresponding to each section is the average of all the pixels in the $n \times n$ frame section. The camera sensor, a photodiode, outputs a current proportional to the intensity of light. The camera uses a buffer stage at the output to convert the current into an output voltage. To compute the super-pixel, the camera first combines the current from the set of pixels and then converts the combined current into a voltage output, all in the analog domain. We note that the averaging of close-by pixels is supported by commercial cameras including the CentEye Stonyman camera [6].

Instead of transmitting the entire $N \times N$ pixel frame, the camera transmits a lower resolution frame consisting of $n \times n$ sized super-pixels (the average value of the pixels in the $n \times n$ block), called the low-resolution frame or L frame. Doing so reduces the data transmitted by the camera by a factor of $\frac{N^2}{n^2}$. The reader performs computation on L frames and implements a change-driven compression technique. At a high level, the reader compares the incoming L frame with the previous L frame. If a super-pixel value differs by more than a predetermined threshold between frames, then the super-pixel has sufficiently changed, and the reader asks the camera to transmit all pixels corresponding to it. If the difference does not cross the threshold, then the reader uses the pixel values corresponding to the previous reconstructed frame to synthesize the new frame without requesting new pixel values. We call the frame that contains the pixel values corresponding to the sufficiently changed super-pixels, the super-pixel frame or S frame. In addition to transmitting the S and L frames, the camera periodically transmits uncompressed frames (I) to correct for potential artifacts/errors that the compression process may have accumulated.

In streaming camera applications, the communication overhead of the reader requesting specific pixel values is minimal and is implemented using a downlink similar to prior 100-500 kbps designs [120, 167], where the receiver at the camera uses a simple envelope detector to decode the amplitude modulated signal from the reader. We note that prior designs can achieve Mbps downlink transmissions [157] as well as full-duplex backscatter [133], which can be used to remove the downlink as a bottleneck.

Fig. 2.6 shows the sequence of frames and pixels transmitted by our camera. Between two I frames, the camera transmits M low-resolution L frames and M super-pixel S frames which contain pixel values corresponding to super-pixels whose value differences exceed the threshold between consecutive frames. The number of L and S frames (M) transmitted between consecutive I frames trades off between the overhead associated with transmission of full resolution frames and the artifacts and errors the compression algorithm introduced. In our implementation of 10 fps HD video streaming, we transmit an I frame after a transmission of every 80 L and S frames.




Figure 2.6: **Distributed Compression.** The sequence of frames and pixels transmitted by the camera to the reader.

2.4 Implementation

We built our wireless camera using off-the-shelf components on a custom-designed Printed Circuit Board (PCB). We use the COTS prototype to evaluate the performance of the wireless camera in various deployments. We then present the application-specific integrated circuit (ASIC) design of the wireless camera that we used to quantify the power consumption for a range of video resolutions and frame rates.

COTS implementation. Our wireless camera design eliminates the power-hungry ADCs and video codecs and consists only of the image sensor, PWM converter, a digital block for camera control and sub-carrier modulation, a backscatter switch and an antenna. We built two hardware prototypes, one for the high definition (HD) and another for the low-resolution version of the camera.

We built the low-resolution wireless camera using the 112×112 grayscale random pixel access camera from CentEye [6], which provides readout access to individual analog pixels. We implement the digital control block on a low-power Igloo Nano FPGA by Microsemi [9]. The analog output of the image sensor is fed to the PWM converter built using passive RC components and a Maxim NCX2200 comparator [14]. We set $R_1 = 83K\Omega$, $R_2 = 213K\Omega$ and $C = 78pF$ in our PWM converter design to support video frame rates of up to 13 fps. PWM converter's output acts as input to the FPGA. The FPGA performs sub-carrier modulation at 1.024 MHz using an XOR gate and outputs the sub-carrier modulated PWM signal to the Analog Devices ADG919 switch which switches a 2 dBi dipole antenna between open and short impedance states. The FPGA in-




Figure 2.7: **HD Wireless Camera Prototype.** A laptop emulating an analog output camera feeds raw pixel data into our backscatter prototype. Then the PWM encoded pixels are backscattered to the reader.

jects frame and line synchronization patterns into the frames data before backscattering. We use Barker codes [4] of length 11 and 13 for our frame and line synchronization patterns, respectively. Barker codes have a high-autocorrelation property that helps the reader more efficiently detect the beginning of the frame in the presence of noise.

We use Verilog to implement the digital state machine for camera control and sub-carrier modulation. Verilog design can be easily translated into ASIC using industry standard EDA tools. We can further reduce our system’s power consumption by using the distributed compression technique. As described in §2.6, a camera deployed in a normal lab space can achieve an additional compression ratio of around $30\times$, which proportionately reduces wireless transmissions.

To develop an HD-resolution wireless camera, we need access to the raw analog pixel outputs of an HD camera. Currently, no camera on the market provides that access. To circumvent this constraint, we download from YouTube HD-resolution sample videos lasting 1 minute each. We output the recorded digital images using a USB interface to an analog converter (DAC) to simulate voltage levels corresponding to an HD quality image sensor operating at 10 fps. Given the USB speeds, we achieve the maximum frame rate of 10 fps. We feed the voltage output to our PWM converter. Fig. 2.7 shows the high-level block diagram of this implementation. For the high-resolution version of the wireless camera, we set $R_1 = 10K\Omega$, $R_2 = 100K\Omega$ and $C = 10pF$ and use an LMV7219 comparator by Texas Instruments [12] in our PWM converter. The digital block and other system components were exactly the same as for the low-resolution wireless camera described above except that sub-carrier frequency is set to ~ 10 MHz here to avoid aliasing.

Table 2.3: Power Consumption

	Frame Rate: 60 fps	Frame Rate: 30 fps	Frame Rate: 10 fps
Video Quality	Power (μW)	Power (μW)	Power (μW)
1080p (1920x1080)	806.50	560.63	167.77
720p (1280x720)	320.94	252.10	78.31
480p (640x480)	126.88	106.78	36.71
360p (480x360)	75.63	65.68	25.11

ASIC Design. As noted, our design eliminates the power-hungry LNA, AGC and ADC at the wireless camera by delegating them to the reader to reduce wireless camera power consumption by orders of magnitude. However, since commercially available Stonyman cameras (like CentEye) and components (such as FPGA) are designed for flexibility and ease of prototyping and are not optimized for power, our COTS implementation cannot achieve the full power savings from our design. Therefore, we analyze the power consumption of an application-specific integrated circuit (ASIC) implementation of our design for a range of video resolutions and frame rates. An ASIC can integrate the image sensor, PWM converter, digital core, oscillator and backscatter modulator onto a small silicon chip. We implement our design in a TSMC 65 nm LP CMOS process.

We use Design Compiler by Synopsis [20] to synthesize transistor level from the behavioral model of our digital core, which is written in Verilog. We custom-design the PWM converter, oscillator and backscatter modulator described in §2.3 in Cadence software and use industry standard simulation tools to estimate power. To support higher resolution and higher frame rate video, we simply increase the operating frequency of the oscillator, PWM converter and digital core. As an example, 360p at 60 fps requires a 10.4 MHz input clock, which consumes a total of 42.4 μW in the digital core, PWM converter and backscatter switch; a 1080p video at 60 fps requires an ~ 124.4 MHz input clock, which consumes 408 μW in the digital core, PWM converter and backscatter switch. To eliminate aliasing in all cases, we choose a sub-carrier frequency equal to the input clock of each scenario. Note that sub-carrier frequency cannot be lower than the effective




Figure 2.8: **Signal Processing Steps at the Reader.** Recovering video from a PWM backscatter signal.

bandwidth of the signal reported in Table 2.2.

We use existing designs to estimate the power consumption of the image sensor for different video resolutions. State-of-the-art image sensors consume $3.2\text{pW}/(\text{frame} \times \text{pixel})$ [184] for the pixels-only image sensor, which results in $33.2\ \mu\text{W}$ for 360p resolution; this increases to $398\ \mu\text{W}$ for 1080p resolution video at 60fps. Table 2.3 shows the power consumption of the ASIC version of our wireless camera for different video resolution and frame rates. Note that these results show power consumption before inter-frame compression distributed across the reader and camera, which could further reduce wireless bandwidth and power consumption.

Reader Implementation. We implement the reader on the X-300 USRP software-defined radio platform by Ettus Research [76]. The reader uses a bi-static radar configuration with two 6 dBi circularly polarized antennas [1]. Its transmit antenna is connected to a UBX-160 daughterboard, which transmits a single-tone signal. USRP output power is set to 30 dBm using the RF5110 RF power amplifier [16]. The receive antenna is connected to another UBX-160 daughter board configured as a receiver, which down-converts the PWM modulated backscattered RF signal to baseband and samples it at 10 Msps. The digital samples are transmitted to the PC via Ethernet.

Fig. 2.8 shows block diagram of the signal processing steps required to recover the transmitted video. For example, for low-resolution video, the received data is centered at an offset frequency of the 1.024 MHz; therefore, we first filter the received data using a 500 order bandpass filter centered at 1.024 MHz. Then, we down-convert the signal to baseband using a quadrature down-conversion mixer. Next, we correlate the received data with 13 and 11 bit Barker codes to determine the frame sync and line sync. After locating frame and line sync pulses, we extract the time periods corresponding to a row of PWM-modulated pixel values and low-pass filter the signal with a 500

order filter to remove out of band noise. We divide the row in evenly spaced time intervals that corresponded to the number of pixels in a single row of the image sensor. We recover the pixel value by calculating the average voltage of the signal, which corresponds to the duty cycle of the PWM-modulated signal. We sequentially arrange recovered pixel values into rows and columns to create video frames. Finally, we run a histogram equalization algorithm on the video to adjust frames intensity and enhance output video contrast [101].

2.5 Evaluation

We now evaluate various aspects of our wireless camera system. We start by characterizing the received video quality from the camera as a function of its distance to the reader. Next, we evaluate the performance of our wireless camera while it is worn by a user under different head positions and orientations. We then evaluate the power available by harvesting energy from RF signals transmitted by the reader to demonstrate the feasibility of a battery-free wireless camera. Finally, we evaluate the distributed interactive compression algorithm under two different scenarios.

2.5.1 Operational Range

We deploy our high-definition wireless camera device in a regular lab space. We use the USRP-based reader implementation (§2.4), which we set to transmit 23 dBm into a 6 dBi patch antenna. This is well below the 30 dBm maximum transmit power permitted by FCC in the 900 MHz ISM band. We vary the distance between the reader and the wireless camera prototype from 4 to 16 feet and configure the camera to repeatedly stream a 15-second video using PWM backscatter communication. The 720p resolution video is streamed at 10 fps, and the pixel values are encoded as 8-bit values in monochrome format. We record the wirelessly received video at the reader and measure the Signal to Noise Ratio (SNR). From that, we calculate the Effective Number of Bits (ENOB) [121] at the receiver.

We plot the ENOB of the video digitized by the reader as a function of distance between the reader and the wireless camera in Fig. 2.9. The plot shows that up to 6 feet from the reader, we achieve an ENOB greater than 7, which indicates negligible degradation in quality of the video,




Figure 2.9: **ENOB of the Received Video.** ENOB of the received video versus distance of our wireless camera prototype from the reader.

streamed using PWM backscatter. As the distance between the reader and the camera increase the SNR degrades, indicates a decrease in ENOB. Beyond the distance of 16 feet, we stop reliably receiving video frames. A separation of 16 feet between the reader and wearable camera is more than sufficient for wearable cameras, typically located a few feet away from readers such as smartphones. For reference, Fig. 2.10 shows frames from a 720p HD color video backscattered with our prototype at different distances from the reader; this resulted in different ENOB values.

We conclude the following: our analog video backscatter approach is ideal for a wearable camera scenario since the video is streamed to a nearby mobile device such as a smartphone. In this scenario, the SNR is high; hence, quality degradation due to an analog approach is not severe. Further, we gain significant power reduction benefits.

2.5.2 *Effect of Head Pose and Motion*

We ask a participant to wear the antenna of our wireless camera on his head and perform different head movements and poses while standing around five feet from a reader with fixed location on a table. The poses included standing still, moving head left and right, rotating head to the side, moving head up and down, and talking. Hence, our evaluation includes scenarios with relative mobility between the reader and camera; in fact, it also includes cases where no line of sight communication exists between the reader and camera. We next, evaluate our wireless camera



Figure 2.10: **Effective Number of Bits (ENOB) versus video quality.** The frame corresponding to an ENOB of 3 (d) shows video quality degradation.

for in-situ applications and assess how movements and antenna contact with a body affect video quality. We record the streaming video with the reader and measure the SNR and ENOB as we did in §2.5.1. Fig. 2.11 plots the ENOB of the received video for five different poses and movements. This plot shows that we can achieve sufficient ENOB while performing most of these gestures, resulting in a high-quality video compared to the original source video.

2.5.3 *RF Power Harvesting*

Next, we evaluate the feasibility of developing a battery-free HD streaming camera that operates by harvesting RF signals transmitted by the reader. We build an RF harvester for the 900 MHz ISM band based on a state-of-the-art RF harvester design [158]. The harvester consists of a 2 dBi dipole antenna and a rectifier that converts incoming RF signals into low-voltage DC output. The low-voltage DC is amplified by a DC-DC converter to generate the voltage levels to operate the

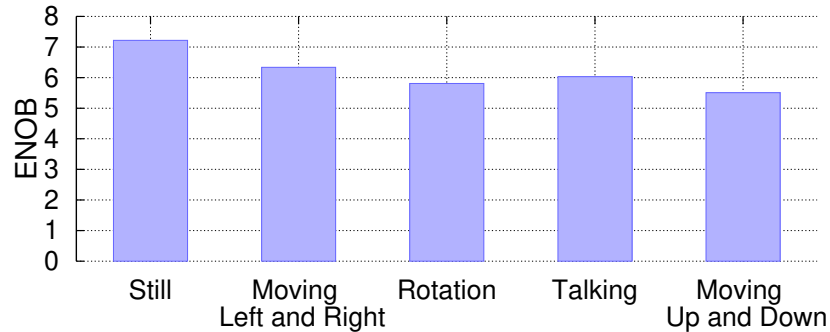


Figure 2.11: ENOB of Received Video Under Different Head Motions.

image sensor and digital control block. We measure the power available at the output of the DC-DC converter.

We configure the reader to transmit a single tone at 30 dBm into a 6 dBi patch antenna and move the RF harvester away from the reader. Fig. 2.12 plots available power at the harvester as a function of distance. Based on available power, we also plot in Fig. 2.12 the maximum resolution and frame rate of the video that could be transmitted by an RF-powered ASIC version of our wireless camera. At close distances of 4 and 6 feet, we see sufficient power available from RF energy harvesting to operate the wireless camera at 60 fps 1080p resolution. As the distance increases, available power reduces, which lowers resolution of video being continuously streamed from the wireless camera. At 16 feet, the wireless camera continuously streams video at 10 fps 480p resolution; beyond this distance, the harvester does not provide sufficient power to continuously power the wireless camera. Note that Fig. 2.12 shows camera performance without using the distributed inter-frame compression algorithm described in §2.3.3. That algorithm, distributed across the wireless camera and reader, reduces camera transmissions, which lowers power consumption and consequently increases operating distances.

2.5.4 Distributed Inter-Frame Compression Evaluation

We consider two scenarios to evaluate our distributed inter-frame compression (subsection 2.3.3). We analyze HD video streamed from a fixed highway monitoring camera and from an action cam-




Figure 2.12: **Power Harvesting.** We plot the average power harvested by the battery-free hardware over different distances from the reader.

era mounted on a user riding a motorcycle [13]. We evaluate the trade-off between the compression ratio and PSNR under both static and dynamic video feeds using our design. We measure the Peak Signal to Noise Ratio (PSNR) for different compression ratios by varying the threshold at which we consider the super-pixel (20 by 20 pixels) to have significantly changed. A higher threshold would result in higher compression ratios but at the cost of a degraded PSNR. Fig. 2.13 shows PSNR of the compressed video as a function of the compression ratio for our distributed inter-frame compression technique. For static cameras, we achieve a compression ratio of at least $35\times$ while maintaining a PSNR above 30 dB. For dynamic videos recorded from a motorcycle, we achieve a compression ratio of $2\times$ for a PSNR greater than 30 dB. This is expected since mobile scenarios significantly change majority of pixel values between frames, resulting in lower compression using our approach. We could address this by implementing a more complex compression algorithm at the reader to track moving objects in the frame and request specific pixel locations, achieving compression levels similar to video codecs. Implementing such complex compression algorithms, however, is beyond the scope of this paper.

2.6 Low-Resolution Security Camera

So far, we have demonstrated how high-resolution video offers a useful paradigm for a wearable camera system. However, various other applications, such as security systems and smart homes,

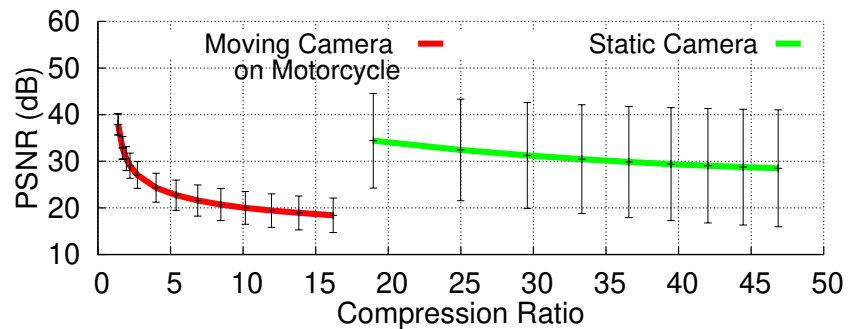


Figure 2.13: **Evaluation of Distributed Inter-frame Compression Algorithm for HD Video.** The plots show PSNR at different compression ratio for two HD videos.

do not require high-resolution video; lower resolution video would suffice for applications such as face detection. Specifically, wireless cameras are increasingly popular for security and smart home applications. In contrast to wearable camera applications, these cameras require low-resolution but much longer operating distances. To show that our design can extend to such applications, we evaluate the operation range at 13 fps 112×112 resolution. Our IC design at this resolution consume less than $20 \mu\text{W}$ without accounting for any distributed inter-frame compression saving.

To evaluate the range, we use a 13 fps 112×112 resolution, gray-scale random pixel access camera from CentEye [6] as our image sensor. The camera has a photodiode image sensor array, a trans-impedance buffer stage (to convert photodiode current to voltage) and a low-noise amplifier. It is extremely flexible, and the user can modify various settings, such as gain of the amplifier stage and, if desired, completely bypass the amplifier. We use this unique camera feature to prototype our wireless camera, which directly transmits analog values from the image sensor (sans amplification) using PWM backscatter. Fig. 2.14 shows photographs of our wireless camera prototype. The camera allows random access, i.e., any pixel on the camera can be accessed at random by setting the corresponding value in the row and column address registers. It can also be configured to output a single pixel, two adjacent pixels, or a super-pixel with sizes ranging from 2×2 to 7×7 . We use the camera's random access and super-pixel functionality to implement our distributed inter-frame compression algorithm. The power consumption of our low-resolution, off-the-shelf




Figure 2.14: **Prototype of Our Low-resolution, Video Streaming Home Security Camera.** Image of the analog camera, FPGA digital core and pulse width modulated (PWM) backscatter, all implemented using COTS components. The overall board measures 3.5 cm by 3.5 cm by 3.5 cm.

analog video streaming prototype is 2.36 mW. We emphasize that this off-the-shelf camera is used only to demonstrate operational range; to achieve the tens of microwatts power budget, we need to use our ASIC design.

Deployment results. We deployed our wireless camera system in the parking lot of an apartment complex. We use the USRP-based reader implementation set to transmit 30 dBm into a 6 dBi patch antenna. We vary the distance between the wireless camera prototype, and, the reader and at each separation, we stream 20 seconds of video from the camera to the reader. Simultaneously, we record camera output using a high input impedance National Instrument USB-6361 DAQ as the ground truth. We choose the popular PSNR metric commonly used in video applications to compare the video wirelessly streamed to the reader using PWM backscatter to the ground truth video recorded at the camera using an NI DAQ. PSNR computes the difference between the ground truth and wirelessly received video.

We measure the PSNR of the received video to evaluate the performance of the wireless camera under normal lighting conditions (below 300 lux) at a frame rate of 13 fps. To evaluate how much our sole wireless communication method affects the quality of received video, we consider PWM converter output as the ground truth for PSNR measurement. Also, to isolate the impact of AGC, which occurs at the reader, unaltered video received by the reader prior to applying any AGC is




Figure 2.15: Operational Range of Low-resolution Security Camera.

compared to the ground truth for PSNR measurement. Fig. 2.15 plots the PSNR of the received video at the reader as a function of the separation between the reader and the wireless camera. The plot shows that wireless camera streamed video at an average PSNR greater than 24 dB to a reader up to a distance of 150 feet away. Beyond 150 feet, the reader does not reliably decode the sync pulses, which limits the operating range of our wireless camera system. Thus, our analog backscatter approach achieves significantly longer range for low-resolution video compared to the HD version of the camera due to the trade-off between bandwidth/data rate and operating distances.

Applying our distributed inter-frame compression algorithm to low resolution videos. We deploy this security camera in a normal lab space and then implement our distributed inter-frame compression technique on the videos received. We evaluate the performance of our algorithm for three different super pixels sizes, 3×3 , 5×5 and 7×7 pixels, and plot results in Fig. 2.16. We achieve a $29.4\times$ data reduction using our distributed inter-frame compression technique while maintaining a PSNR greater than 30 dB.

Face detection accuracy. Next, we demonstrate that the quality of the video streamed from our low-resolution COTS implementation is sufficient for detecting human faces. Such a system can be used to detect human occupancy, grant access (such as Ring [17]), or set off an alarm in case of an intruder. To evaluate the system, we place the wireless camera at five different distances ranging from 16 to 100 feet from the reader. We ask ten users to walk around and perform gestures within 5 feet of the camera. We stream a 2 minutes video at each location at 13 fps and use the MATLAB

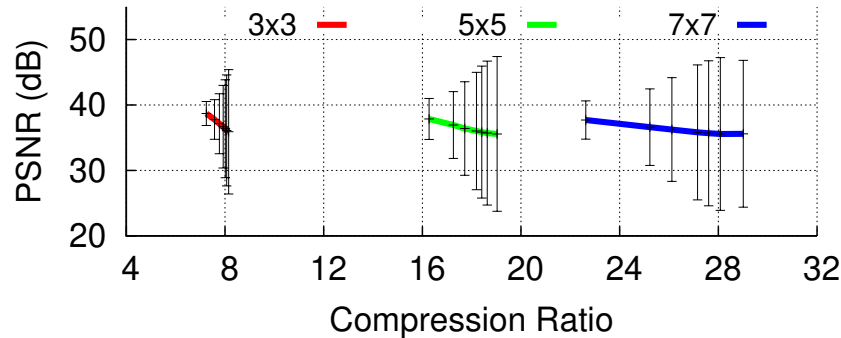


Figure 2.16: **Distributed Interactive Compression with Our Low-resolution Security Camera.** PSNR of the compressed video as a function of compression parameters in a typical lab setting.

implementation of the Viola-Jones algorithm to analyze the approximately four thousand video frames. Fig. 2.17 shows the accuracy of face detection as a function of the PSNR of the received video: as the quality (PSNR) of the video improves, the accuracy of face detection increases. We accurately detect up to 95% of human faces when the PSNR exceeds 30 dB.

2.7 Related Work

Prior work falls in two different categories.

Backscatter communication. An early example of analog backscatter was a gift by the Russians to the US embassy in Moscow, which included a passive listening device. This spy device consisted of a sound-modulated resonant cavity. The voice moved the diaphragm to modulate the cavity’s resonance frequency, which could be detected by analyzing the RF signals reflected by the cavity. [7]. A more recent example of analog backscatter is a microphone-enabled, battery-free tag that amplitude-modulates its antenna impedance using microphone output [182, 180]. In contrast, we design the first analog video backscatter system. Further, prior microphone designs had a low data-rate compared to video streaming. Our camera at 10 fps transmits about 9.2M pixels per second; for a microphone, a few kilo-samples of audio transmission is sufficient to fully recover the voice. In addition, our 13 fps 163K pixels per second camera operates at more than four times the range of the microphone in [182] due to pulse width modulation.

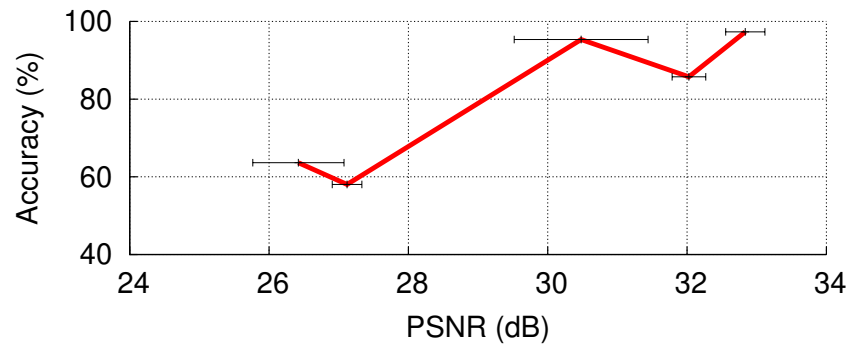


Figure 2.17: **Face Detection with Our Low-resolution Security Camera.** We show the accuracy of face detection on the video streamed from our wireless camera.

Ekhonet [201] optimizes the computational blocks between the sensor and the backscatter module to reduce the power consumption of backscatter-based wireless sensors. Our design builds on this work but differs from it in multiple ways: 1) prior work still uses ADCs and amplifiers on the cameras to transform pixels into the digital domain and hence cannot achieve streaming video on the limited harvesting power budget. In contrast, we provide the first architecture for battery-free video streaming by designing an analog video backscatter solution.

Recent work on Wi-Fi and TV-based backscatter systems [119, 120, 114, 133, 157] can achieve megabits per second of communication speed using a backscatter technique. Integrating these designs with our video backscatter approach would prove a worthwhile engineering effort.

Low-Power cameras. [152] introduces a self-powered camera that can switch its photo diodes between energy harvesting and photo capture mode. Despite being self-powered, these cameras do not have wireless data transmission capabilities. [181, 145, 147, 142, 143, 144] show that using off-the-shelf, low-resolution camera modules, one can build battery-free wireless cameras that will capture still images using the energy they harvested from RF waves, including Wi-Fi and 900 MHz transmissions. Despite their ability to transmit data wirelessly, they are heavily duty cycled and cannot stream video. In particular, these designs can send a new frame at most every ten seconds when they are very close to the RF power source (within about a foot) and once every few tens of minutes at longer distances [181].

[136] presents a 90×90 pixels image sensor with pixels that are sensitive to changes in the environment. If a pixel receives sufficient illumination variation, the pixel address will be stored in a FIFO, thus compressing the image to the pixels that have significantly changed. Despite enabling image compression to occur at the image sensor, this system does not stream live video. In addition, at this low resolution, it burns about 3 mW of power when running at 30 fps. [128] introduces a 128×128 pixel event-driven image sensor that emphasizes low latency for detecting very fast moving scenes, so its power consumption is orders of magnitude higher than our system's.

[130] addresses the problem of conventional image sensors' power consumption not scaling well as their resolution and frame rate increases. In particular, the authors propose to change the camera input clock and aggressively switch the camera to standby mode based on desired image quality. However, streaming video requires addressing the power consumption of multiple components, including camera, communication, and compression. Our work jointly integrates all these components to achieve the first battery-free video streaming design.

Finally, [146] shows that a regular camera on wearable devices burns more than 1200 mW, which limits the camera's operation time to less than two hours on a wearable device. They instead design a low-power wearable vision system that looks for certain events to occur in the field of view and turns on the primary imaging pipeline when those events happen. The power and bandwidth savings, however, are limited to the application and do not address communication. In contrast, we present the first battery-free video streaming application by jointly optimizing both backscatter communication and camera design and by eliminating power-consuming interfaces such as ADCs and amplifiers.

2.8 Limitations and Conclusion

This paper takes a significant first step in designing video streaming for battery-free devices. In this section, we discuss limitations and a few avenues for future research.

Security. Our current implementation does not account for security. However, to secure the wireless link between the camera and reader, we can leverage the fact that our digital core processes the PWM signal. Each wireless camera can be assigned a unique pseudo random security key. Based

on this key, the camera's digital core can modulate the width of the PWM-encoded pixel value using an XOR gate. The reader, which knows the security key, can map the received data to the desired pixel values by performing the analogous operation.

ASIC versus off-the-shelf. While existing works on backscatter cameras focus on using off-the-shelf components, they treat cameras and backscatter independently and just interface the two. Thus, these works cannot achieve video streaming and the low-power demonstrated in this paper. Our key contributions are to make a case for a joint camera and backscatter architecture and to design the first analog video backscatter solution. However, this holistic architecture cannot be achieved with off-the-shelf components. The cost of building ICs in a research environment is prohibitively high. We believe that we spec out the IC design (in §2.4) in sufficient detail using industry standard EDA tools to take it from the lab to industry.

Mobile device as reader. To support our design for HD video streaming from wearable cameras, the smartphone must support a backscatter reader. We can use RFID readers that can be plugged into the headphone jack of the smartphone [8] to achieve this. In the future, we believe that backscatter readers would be integrated into smartphones to support video and other applications [114].

Enabling concurrent streaming. Our analog backscatter camera design can use frequency division multiplexing techniques to share the wireless medium across multiple devices. The reader can coordinate communication by assigning different frequency channels to each camera. Multiple cameras can simultaneously backscatter on different channels by using different (assigned) frequency offsets in our sideband modulation. However, evaluating this approach is beyond the scope of this paper.

Chapter 3

LOW-POWER LONG-RANGE CONNECTIVITY FOR LOW-RATE APPLICATIONS

3.1 Introduction

Embedding cheap connectivity into billions of everyday objects has been a long-standing vision in the ubiquitous computing community. There is however a significant disconnect between the ubiquity articulated by this vision and the capabilities of today's communication technologies. Active radio technologies including Wi-Fi, ZigBee, SigFox [64], LoRa [46] and LTE-M [61] provide reliable coverage and long ranges but are power consuming and cost at least 4–6 dollars [55, 67]; making them too expensive for embedding into objects at scale. Further, given their high peak current and power requirements, active radios significantly deteriorate battery life and are incompatible with emerging small, flexible and ultra-thin printed batteries [57, 53, 58] that promise innovative applications across healthcare, wearable devices and cosmetics [71, 139, 135].

Backscatter promises to be an extremely low power, smaller and cheaper alternative to active radios. Given the absence of expensive radio analog components including RF oscillators, decoupling capacitors and crystals, backscatter designs including passive RFID and Wi-Fi backscatter [120, 114] cost only a few cents to manufacture at scale [45]. Further, they consume three to four orders of magnitude lower power and peak current than radios and hence can operate with emerging flexible, printed and ultra-thin printed battery technologies. However, despite all these benefits, current backscatter designs have seen very limited adoption beyond RFID applications. This is because current backscatter designs are unreliable, limited in operating range, and in fact today cannot achieve robust coverage across rooms [114, 97, 2, 185] as outlined in Fig. 3.2 and Table 3.3¹. Furthermore, since RF signals get significantly attenuated by the human body, backscatter range

¹We note that active RFID does not use backscatter. Instead, it uses radios, is power consuming and costs between




Figure 3.1: **LoRa backscatter deployment.** The LoRa Backscatter device consumes $9.25 \mu\text{W}$, operates at 100s of meters and can be powered by flexible printed batteries and button cells (10 cents), a capability that cannot be achieved with radios. The RF source transmits a single tone that the backscatter device uses to synthesize CSS signals. The challenge is that at the receiver, the backscatter signal is not only drowned by noise but also suffers interference from the RF source.

is limited to tens of centimeters in several healthcare and wearable device applications [114].

This paper questions the conventional wisdom that backscatter is a *short-range* system. Specifically, we ask if one can achieve wide-area backscatter communication with a range of hundreds of meters, if not kilometers. A positive answer would give us the best of both worlds: long-range reliable communication capabilities of radios at the low-power and cost of backscatter hardware. This enables, for the first time, wide area connectivity for everyday objects and opens applications in domains like smart cities [29], precision agriculture [35], industrial, medical and whole-home sensing [114], where backscatter is currently infeasible.

To appreciate why this is hard, consider the deployment in Fig. 3.1. Here the backscatter device reflects signals from an RF source to synthesize data packets that are then decoded by a receiver. The challenge is that, before arriving at the backscatter device, the signals from the RF source are already attenuated. The backscatter device can reflect these weak signals to synthesize data packets which get further attenuated as they propagate to the receiver. Our experiments show that with a separation of 400 m between the RF source and receiver, the backscattered signal is at -134 dBm. In contrast, the direct signal from the RF source at the receiver is more than a million times stronger at -45 dBm. Thus, the backscatter signal is not only drowned by noise but also suffers significant

interference from the RF source.

We present the first wide-area backscatter communication system. Achieving this requires us to satisfy two key constraints. First, the backscatter device should code information in a way that can be decoded at the receiver down to and below -135 dBm signal strength and reliably operate in the presence of strong out-of-band interference. Second, instead of using a custom receiver for the backscattered signal that can be prohibitively expensive (e.g., RFID readers), the backscattered signals should be decoded on readily and cheaply available commodity hardware that would expedite the adoption and development of our design.

To do so, we first profile existing radio technologies in Table. 3.3, which shows that LoRa provides the highest sensitivity of -149 dBm and supports bit rates of 18 bps to 37.5 kbps, which are sufficient for most IoT applications. Further, LoRa is resilient to both in-band and out-of-band interference [50]. Specifically, the Sx1276 receiver hardware from SEMTECH can reliably decode LoRa packets in the presence of 95 dB higher out of band interference [50]. This is significant because as described above, in typical backscatter deployments, the backscattered signal at the receiver can be as low as -135 dBm and this signal must be decoded in the presence of a strong single tone out of band interference as high as -45 dBm. 95 dB out of band interference specifications of Sx1276 Lora receivers by SEMTECH satisfy these stringent requirements. Motivated by this, we present the design and implementation of the first LoRa backscatter system. At a high level, the RF source transmits a single tone signal, which our backscatter devices use to synthesize LoRa compatible packets. To achieve this, we make two key technical contributions.

- *We introduce the first chirp spread spectrum (CSS) backscatter design [179].* LoRa uses CSS modulation where, as shown in Fig. 3.5(a), a ‘0’ bit can be represented as a continuous chirp that increases linearly with frequency, while a ‘1’ bit is a chirp that is cyclically shifted in time. Thus, CSS requires continuously changing the frequency as a function of time, which has not been demonstrated on backscatter hardware. This is challenging since while existing backscatter approaches can generate DBPSK/DQPSK (802.11b/ZigBee [120, 114]) and 2-FSK (Bluetooth [97]) transmissions, they are all limited to discrete digital values. Building on existing radio architectures, we design the first backscatter design in §3.3.2 that can synthesize continuous frequency modulated

chirps, while consuming as low as $9.25 \mu\text{W}$ which is 3 orders of magnitude lower power compared to LoRa radios. We also reverse-engineer the proprietary LoRa PHY layer in §3.3.4 to backscatter LoRa-compatible packets.

- *We present the first backscatter harmonic cancellation mechanism.* Backscatter uses a switch to either reflect or absorb the incident RF signals and create square waves. This however creates third and fifth harmonics in adjacent frequency bands when backscattering the signal from the RF source, resulting in interference and affecting network performance. State-of-the-art single-sideband backscatter designs [114] ignore these harmonics and hence create out-of-band interference. Since LoRa has a high sensitivity, this affects other LoRa devices operating in adjacent bands. In §3.3.3, we present a low-power backscatter design that cancels these sideband harmonics and improves spectral efficiency.

Building on the above techniques, we design a link-layer protocol that enables multiple long range backscatter devices to share the spectrum. We also design a LoRa backscatter IC and estimate the power consumption using Cadence and Synopsis software toolkits [54, 20]. Our results show that our IC design is comparable in area to RFID and consumes as little as $9.25 \mu\text{W}$ while generating continuous frequency modulated chirps and performing harmonic cancellation.

Below, we summarize our evaluation in various deployment scenarios.

- We evaluate our design in various line-of-sight scenarios. Our results show that even when the RF source and receiver are separated by 475 m, the backscatter device could operate at all locations between them. Further, when the backscatter device is co-located with the RF source, the LoRa receiver can decode transmissions from as far as 2.8 km from the backscatter device.
- We deploy our system in a one-acre vegetable farm, a $4,800 \text{ ft}^2$ (446 m^2) house spread across three floors and a $13,024 \text{ ft}^2$ (1210 m^2) office space spanning 41 rooms separated by wood, concrete walls and metal structures. Our results show reliable backscatter coverage across all these deployments, using only a single RF source and receiver.
- Finally, we show that our design can also enable backscatter in applications that are not favorable for RF propagation. Specifically, we evaluate a contact lens form factor antenna in-vitro and show




Figure 3.2: We compare the operating distance of existing backscatter systems.

Figure 3.3: **Communication Technologies.** We show the sensitivity and supported data rates for different communication technologies and the feasibility of powering them from different sources.

Technology	Sensitivity	Data Rate	Whole Home Coverage	Button Cell	Tiny Solar Cell	Printed Battery
Wi-Fi (802.11 b/g)	-95 dBm	1-54 Mbps	✓	✗	✗	✗
LoRa	-149 dBm	18 bps–37.5 kbps	✓	✗	✗	✗
Bluetooth	-97 dBm	1-2 Mbps	✗	✗	✗	✗
Sigfox	-126 dBm	100 bps	✓	✗	✗	✗
ZigBee	-100 dBm	250 kbps	✓	✗	✗	✗
Passive Wi-Fi	-95 dBm	1-11 Mbps	✗	✓	✓	✓
RFID	-85 dBm	40–640 kbps	✗	✓	✓	✓
LoRa Backscatter	-149 dBm	18 bps–37.5 kbps	✓	✓	✓	✓

that it can easily backscatter data from across a $3,328 \text{ ft}^2$ (309 m^2) room using only a single RF source and receiver. This is orders of magnitude larger than the 35 inch (89 cm) range achieved by prior designs [114]. We also build a flexible epidermal patch sensor that backscatters data in the above room, while being attached to human skin.

3.2 The case for LoRa backscatter and Related Work

To demonstrate that LoRa backscatter is the best fit for achieving ubiquitous connectivity, we first compare it with existing solutions and show how it fares across the spectrum of power consumption, cost, size, sources of power and operating range. We then discuss prior work and compare

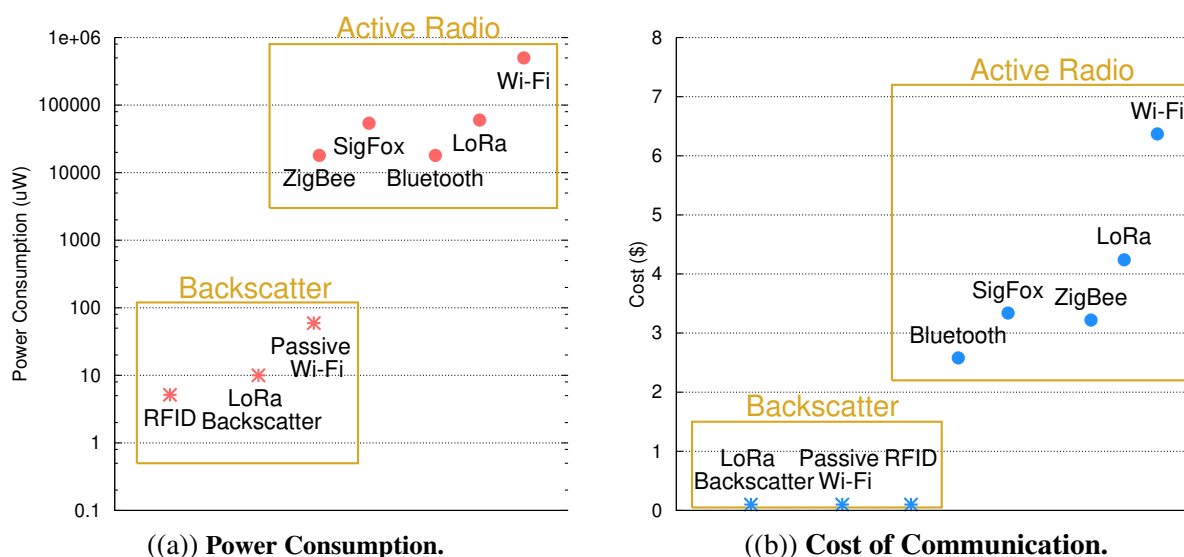


Figure 3.4: We compare power consumption and cost of different communication technologies including backscatter and radio techniques.

our LoRa backscatter system with state of the art backscatter approaches.

- *Operating range.* Fig. 3.2 shows the range of different radio and backscatter communication solutions. Radios including Wi-Fi, Bluetooth and ZigBee operate up to 100s of meters while wide area LoRa and SigFox deployments extend operation to kilometers. However, existing backscatter solutions such as RFID and Passive Wi-Fi are limited to tens of meters of operating distance in best-case scenarios. LoRa backscatter instead extends backscatter operation to 100s of meters. This achieves the whole home and office coverage of Wi-Fi and ZigBee radios while delivering the cost, size and power benefits of backscatter described below.

- *Power consumption.* Fig. 3.4(a) shows the power consumption of popular radio and existing backscatter solutions. Radios including Wi-Fi, BLE, ZigBee, Lora and SigFox all consume between 10 to 500 mW, which is 3–4 orders of magnitude higher than the power consumption of backscatter systems including RFID, Passive Wi-Fi and the LoRa backscatter system. LoRa backscatter consumes three orders of magnitude lower power than LoRa radios and would significantly extend the battery life. Further it can easily operate for more than 10 years on button cells

and printed batteries that are fraction of the size of batteries used with LoRa radios.

- *Cost and size.* Fig. 3.4(b) plots the cost of active radio and backscatter communication solutions. Radios are at least an order of magnitude more expensive compared to backscatter solutions. This is because an active radio requires analog RF components such as local oscillators, mixers and amplifiers that consume significant silicon area. Since, the cost of an IC is directly proportional to silicon area, the analog components increase the cost of radios. Additionally, radios require external components such as crystals, matching inductors and decoupling capacitors all of which increase the overall cost. In contrast, backscatter solutions are primarily digital in nature and scale with Moore's law that significantly reduce the area and consequently the cost of the backscatter IC to few cents. Also, backscatter communication modules can be built with just an IC, printed antenna and tiny battery and do not require external components like crystals, inductors and decoupling capacitors which reduces the component, manufacturing and assembly costs as well as the overall size.

- *Sources of Power.* Flexible thin film printed batteries are an emerging source of power which result in cheap, thin and small form factor solutions. However, a key challenge with printed batteries is that they have very limited capacity and peak current requirements. Radios due to their high active power consumption and peak current are incompatible with such technologies. Thin small inexpensive solar cells and button cells, which cost only 10 cents, also suffer from similar limitations. However, LoRa backscatter consume tens of microwatts of power and can operate for ten years on button cells and 10 cm^2 area printed batteries. Additionally, LoRa backscatter can be powered with a tiny 2 cm^2 size photodiode which makes the overall connectivity module small and inexpensive.

In summary, LoRa backscatter provides the first connectivity solution at a fraction of the cost, size and power consumption of radios while providing reliable and long ranges and being compatible with emerging thin film printed battery solution and inexpensive button cells and tiny solar cells.

Comparison to prior backscatter research. LoRa backscatter builds on recent work on backscat-

Table 3.1: **Comparison of LoRa Backscatter with Existing Backscatter Systems.** d_1 is the distance between the backscatter device and the RF source. d_2 is the distance between the backscatter device and the receiver.

Technology	Data Rates	Deployment # 1 (d_1, d_2)	Deployment # 2 (d_1, d_2)
Passive Wi-Fi [120]	1-11 Mbps	2 m, 30 m	4.5 m, 4.5 m
Interscatter [114]	2-11 Mbps	1 m, 27 m	N/A
BLE backscatter [97]	1 Mbps	0.9 m, 8.5 m	4.5 m, 4.5 m
BackFi [161]	1-5 Mbps	N/A	5 m (Full Duplex)
LoRea [187]	2.9 kbps	1 m, 225 m	N/A
HitchHike [200]	0–300 kbps	1 m, 50 m	N/A
FSK Backscatter [189]	1.2 kbps	3 m, 268 m	25 m, 25 m
LoRa backscatter	50 bps–37.5 kbps	5 m, 2.8 Km	237.5 m, 237.5 m

ter communication [197, 93, 198, 109, 154, 134, 98]. The closest to our work is recent work that backscatters ambient signals like TV [132, 157], Wi-Fi [119, 161, 120, 202, 200], Bluetooth [97, 202, 160] and ZigBee [160, 114]. A comparison of the LoRa backscatter system with existing backscatter systems in terms of data rates and ranges is shown in Table 3.1. In the table, deployment # 1 refers to the scenario where the backscatter device is placed near the signal source (d_1) and the receiver is placed at a distant location (d_2). In deployment #2, the backscatter device is equidistant to the signal source and receiver ($d_1 = d_2$) and demonstrates the practical operating range of backscatter systems.

As seen from Table 3.1, existing backscatter systems have a range of 10 m, when the backscatter device is 6 m away from the RF source and 50 m when the backscatter device is 1 m away [200]. [189, 187] uses FSK modulation to increase the backscatter range. Specifically, the authors show that when the backscatter tag is 3 m from the RF source, the receiver can be 246 m away. In more realistic scenarios, when the tag is around 25 m from the signal source and the receiver, their packet error rate was 100%. In a similar setup, we can deliver packets even when the receiver is 2.8 km away. Our design uses CSS modulation and achieves much higher sensitivities. As a result, we achieve 1–2 orders of magnitude higher communication ranges and thus enable wide-area backscatter.

3.3 System Design

LoRa backscatter uses chirp spread spectrum (CSS) to design a wide-area backscatter communication system. In this section, we first provide an overview of CSS. We then present our hybrid analog-digital backscatter design to create CSS transmissions as well as our harmonic cancellation mechanism. We then describe how to synthesize CSS packets that are compatible with the LoRa physical layer by reverse-engineering LoRa. Finally, we outline a link-layer protocol that enables multiple CSS backscatter devices to co-exist with each other.

3.3.1 Understanding CSS

CSS modulation and demodulation. Chirp spread spectrum (CSS) uses linear frequency modulated chirp pulses to convey information. A key characteristic that CSS leverages is that a time delay in the chirp signal translates to a frequency shift at the output of the FFT. CSS modulation uses this to encode data as cyclic time shifts in the baseline chirp. Figures 3.5(a) and 3.5(b) show two chirp symbols, where the first symbol is represented by the baseline chirp and the second is represented by a cyclic shifted chirp, offset by half the chirp duration.

The receiver demodulates these symbols by first multiplying the incoming signal with the baseline chirp and then performing an FFT. Since multiplication in the time domain is correlation in the frequency domain, the resulting operation results in a peak in the FFT frequency bin corresponding to the time delay in the received chirp. Figures 3.5(c) and 3.5(d) show the resulting FFT for the two chirps. The figure shows that the FFT has a peak in the first FFT bin for the first symbol (corresponding to zero-time delay) and has a peak in the middle FFT bin for the second symbol (corresponding to half a chirp delay). Thus, by tracking FFT peaks, we can decode the data.

Note that one can transmit multiple bits within each chirp symbol. Specifically, say the receiver performs a N point FFT. It can distinguish between N different cyclic shifts which result in a peak in each of the N FFT bins. Thus, we can transmit $\log_2 N$ bits within each chirp. In the figure, N is set to 2^{SF} , where SF is the spreading factor of CSS modulation, which we discuss next.

CSS parameters and bit rates. There are three parameters that determine the bit rate achieved




Figure 3.5: CSS Modulation & Receiver FFT Processing. The two figures on the left shows two chirp symbols, where the second symbol is a cyclic shifted version of the first, offset by half the chirp duration. BW and SF are the bandwidth and spreading factor respectively. The time delays in the chirp translate to a peak in the FFT domain. The two plots on the right show the FFT output corresponding to the two symbols.

while using CSS modulation: 1) chirp bandwidth, 2) spreading factor and 3) symbol rate. As shown in figs. 3.5(a) and 3.5(b), if BW denotes the bandwidth, the frequency of the baseline chirp increases linearly between $-\frac{BW}{2}$ and $+\frac{BW}{2}$. The data bits are encoded as cyclic shifts of this baseline chirp, where each cyclic shift represents a modulated symbol. The spreading factor, SF , is the number of bits encoded in each chirp duration. From earlier discussion, a chirp with N samples can encode $\log_2 N$ bits. Thus, a CSS chirp with a spreading factor SF has 2^{SF} samples. Finally, the symbol rate is the number of chirp symbols per second.

With a bandwidth of BW , the Nyquist sampling rate is $\frac{1}{BW}$ samples per second. Thus, given a spreading factor of SF , the length of each symbol is given by $\frac{2^{SF}}{BW}$ seconds and so the symbol rate is $\frac{BW}{2^{SF}}$ symbols per second. Since each chirp can represent SF bits, the bit rate can be written as, $\frac{BW}{2^{SF}} SF$. Thus, one can achieve different bit rates by either changing the bandwidth or the spreading factor. As we see in §3.3.4, LoRa also uses error coding codes on top of CSS modulation, giving it a third degree of freedom, in addition to chirp bandwidth and spreading factor, to adapt bit rate.

Case for using CSS modulation in backscatter systems. We outline the unique properties of CSS which make it a great fit for wide area backscatter application.

- CSS achieves high sensitivity by using an efficient tradeoff between bandwidth and data rates when the signal is drowned in noise [85]. Consequently, Sx1276 Lora receiver can decode CSS

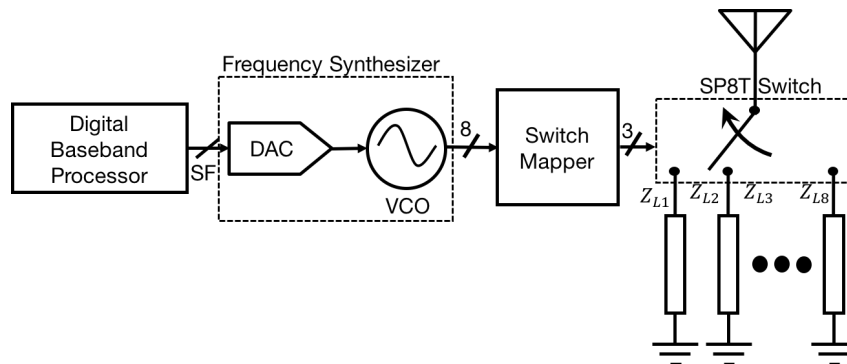


Figure 3.6: **Hybrid Analog-Digital Backscatter.** The digital baseband processor generates the frequency plan that is converted in the analog domain to control the output frequency of a VCO. The time shifted versions of the VCO output are mapped according to the dataset of the approximated exponential signal to their respective backscatter impedance values using a SP8T RF switch.

data packets down to -149 dBm.

- CSS is resilient to fading, Doppler and in-band interference that are common in wide-area deployments [170]. We demonstrate in §3.5.2 that CSS is indeed resilient to out-of-band interference. Specifically, CSS receivers can correctly decode packets in the presence of a 95 dB higher out of band single tone interference.
- Unlike DSSS techniques (e.g., GPS) that have a long signal acquisition time at low SNRs, CSS has significantly lower acquisition overhead [170, 87, 36]. This reduces the overhead of transmitting data. In addition, CSS transceivers do not require fine-grained frequency synchronization. This is because small offsets in the oscillators at the transmitter and receiver result in a frequency offset that which can be corrected during the FFT operation. As a result, CSS receiver hardware can be significantly cheap while operating at a high sensitivity.

3.3.2 Synthesizing CSS with Backscatter

Challenge. The key challenge is that the complexity of generating CSS signals in the digital domain scales exponentially with the spreading factor used in the CSS transmissions. To understand this, consider CSS modulation with a spreading factor of two. As described earlier, a CSS sig-




Figure 3.7: Four CSS symbols when spreading factor is 2.

nal with a spreading factor SF can have 2^{SF} cyclic shifts. Thus, the four cyclic shifts shown in Fig. 3.7 correspond to CSS modulation with a spreading factor of two. To create this signal, the backscatter device needs to generate at least the four frequencies, f_0, \dots, f_3 , shown in the figure. More generally, to synthesize a CSS modulation with a spreading factor of SF , the backscatter device must create signals at 2^{SF} frequencies. As we see in §3.3.4, LoRa receivers use spreading factors between 6 and 12. This translated to 64–4096 frequencies. Backscattering all these frequencies requires either using 64–4096 oscillators or running the digital clock at a frequency of $lcm(f_0, f_1, \dots, f_{4096})$. The first approach is expensive and power consuming while the latter requires using a clock frequency of GHz, which is power consuming and defeats the purpose of using backscatter.

Our solution. We present the first backscatter design that can generate CSS modulated signals. Our design uses a backscatter switch topology which is functionality equivalent to radios but consumes three orders of magnitude lower power [165, 195]. Specifically, we use a hybrid digital-analog backscatter design where we use the energy efficient digital domain to create a frequency plan for the continuously varying CSS signal and then map it to the analog domain using a low-power DAC. For example, to create the second cyclic shift in Fig. 3.7, the digital baseband creates the frequency plan f_1, f_2, f_3, f_0 , which the analog domain uses to create the desired frequencies.

Fig. 3.6 shows the architecture for our backscatter design. It has the digital baseband processor, digital to analog converter (DAC) and a voltage-controlled oscillator. The voltage-controlled oscillator (VCO) is a device that outputs a clock with a frequency that is proportional to the input

voltage. We vary the frequency output of the oscillator by using the DAC to generate the appropriate voltages. Specifically, the digital baseband processor outputs an SF bit number, where SF is the CSS spreading factor. This allows us to output 2^{SF} voltage levels at the output of the SF -bit DAC. The analog voltage output of the DAC controls the frequency of the VCO.

The challenge however, is that, as shown in figs. 3.5(a) and 3.5(b), in a CSS encoded packet, the frequency of the signal varies from a negative frequency ($-\frac{BW}{2}$) to a positive frequency ($\frac{BW}{2}$). A voltage-controlled oscillator however only outputs signals at positive frequency. So, we need a mechanism to synthesize negative frequencies using backscatter. From basic communication theory, negative frequencies essentially can be written as complex signals. Specifically, the complex exponent, $e^{j2\pi(\pm f)t}$ can be written as $\cos 2\pi ft \pm j\sin 2\pi ft$. Thus, generating negative frequencies requires us to generate both the in-phase cosine signal as well as the out-of-phase sine signal at the desired frequencies. To do this, existing solutions approximate the sine and cosine signals using the square wave output by the VCO. This however results in out-of-band harmonics. In the next section, we describe our harmonic cancellation mechanism.

We note that the receiver receives both the single-tone signal from the RF source as well as the backscattered LoRa packets. Since the backscatter signal is much weaker, the single-tone creates in-band interference. To address this, we shift the single-tone outside the desired band and transform it into out-of-band interference. At a high level, if the RF source transmits the single-tone signal $e^{2\pi f_c t}$ and the backscatter signal generates the complex signal, $e^{2\pi(\Delta f + f_{LoRa})t}$, then the resulting backscattered RF signal is $e^{2\pi((f_c + \Delta f) + f_{LoRa})t}$. Here Δf is a small fixed offset and f_{LoRa} is the varying frequency corresponding to the baseband LoRa modulation. The above equation shows that by using a small frequency offset, Δf , we can create the LoRa signals in a band centered at $f_c + \Delta f$ which is different from that of the single tone, f_c .

3.3.3 Backscatter Harmonic Cancellation

As described earlier, prior backscatter designs [114, 120] use square waves to approximate sine and cosine waves which results in harmonics. To understand why this happens, we recall that a square wave at a rate of Δf can be written as a sum of cosine waves using the following expression:

$$\text{Square}(\Delta ft) = \frac{4}{\pi} \sum_{n=0}^{\infty} \frac{1}{2n+1} \cos(2\pi(2n+1)\Delta ft)$$

If the RF source transmits $\cos(2\pi f_c t)$ and the backscatter device is switching with a square wave operating at Δf frequency, signal transmitted by the backscatter device can be written as $\cos(2\pi f_c t) \text{Square}(\Delta f)$. As a result, in addition to generating the desired signal at $f_c + \Delta f$, the above operation also generates the mirror copy at $f_c - \Delta f$, 9.5 dB lower harmonic at $f_c \pm 3\Delta f$, 15 dB lower harmonic at $f_c \pm 5\Delta f$ and additional lower power harmonics. Recent work [114] has demonstrated how one can eliminate the mirror copy being generated at $f_c - \Delta f$ using single side band backscatter technique. However, this technique still preserves the third, fifth and other odd order harmonics. The third and fifth order harmonics are only 9.5 and 15 dB lower than the desired backscattered signal and hence create interference on the wireless channel. More importantly, since the LoRa protocol has very low sensitivities, LoRa devices operating in channels overlapping with the third and fifth harmonics experience in-band interference from backscatter devices.

Our Solution. Our insight is to use a different signal from the square wave to approximate a cosine and sine wave. On a high level, one can think of an analog signal as a discrete signal with infinite distinct voltage levels and smooth transitions, which results in a clean spectrum without any harmonics. However, square wave has only two levels with discontinuous step transitions, which results in high frequency components. Our key idea with harmonic cancellation is to introduce additional voltage levels to better approximate a sinusoidal signal, by imitating radios [193, 163], and obtain a cleaner frequency spectrum.

Consider the approximation of a cosine wave in Fig. 3.8, using a signal with four voltage levels. This approximated cosine wave can be written as the sum of three square waves slightly shifted from one other, $S_0(t)$, $S_1(t)$ and $S_2(t)$, as shown in the figure. Here T is the time period.

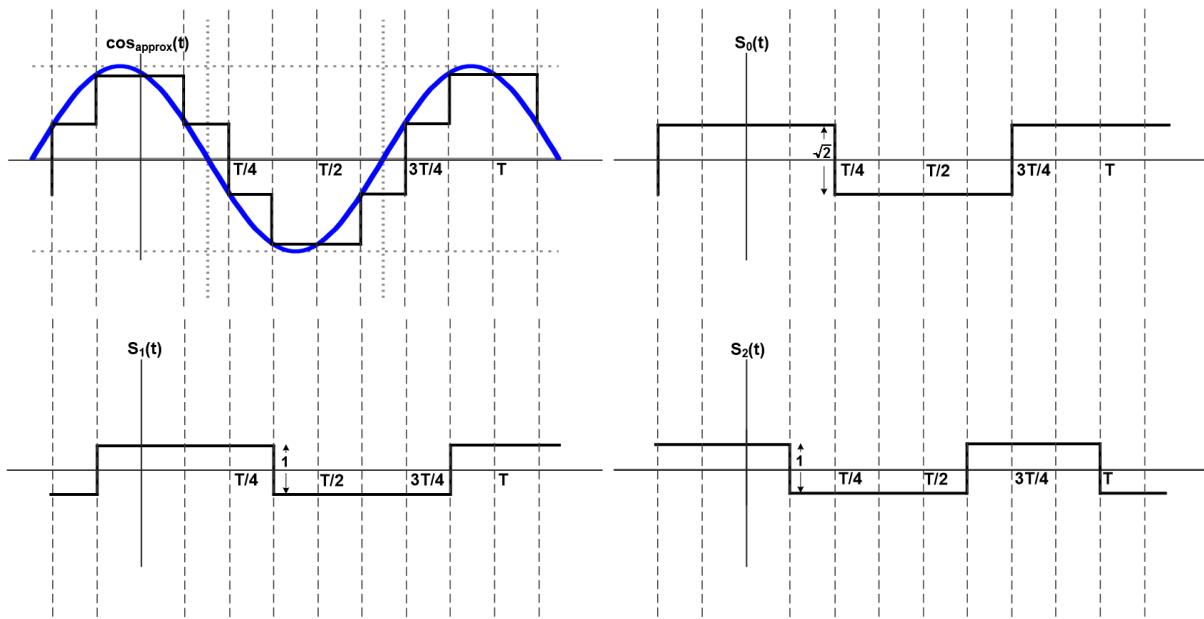


Figure 3.8: **Approximation of a cosine wave with multi-level signal.** We approximate the cosine wave as a sum of three digital signals $S_0(t)$, $S_1(t)$ and $S_2(t)$ resulting in a multi-level signal.

$$S_0(t) = \frac{4\sqrt{2}}{\pi} \sum_{n=0}^{\infty} \frac{\sin[(2n+1)2\pi\Delta f(t + \frac{T}{4})]}{2n+1}$$

$$S_1(t) = \frac{4}{\pi} \sum_{n=0}^{\infty} \frac{\sin[(2n+1)2\pi\Delta f(t + \frac{T}{8})]}{2n+1}$$

$$S_2(t) = \frac{4}{\pi} \sum_{n=0}^{\infty} \frac{\sin[(2n+1)2\pi\Delta f(t + \frac{3T}{8})]}{2n+1}$$

Using the above expression for the three signals, we can now express the approximated cosine waveform as,

$$\begin{aligned} \cos_{approx}(2\pi\Delta ft) &= S_0(t) + S_1(t) + S_2(t) \\ &= \frac{4}{\pi} \sum_{n=0}^{\infty} \frac{\sin[(2n+1)2\pi\Delta f(t + \frac{T}{4})][2\cos((2n+1)\frac{\pi}{4}) + \sqrt{2}]}{2n+1} \end{aligned}$$

The sine wave can now be generated by simply shifting the cosine waveform by quarter of the time period. Using these approximations for the sine and cosine parts of the waveform, the

exponential, $e^{j2\pi\Delta ft}$, can now be mathematically written as,

$$\begin{aligned}
e^{j2\pi\Delta ft} &= \cos(2\pi\Delta ft) + \sin(2\pi\Delta ft) \\
&\approx \cos_{approx}(t) + \sin_{approx}(t) \\
&= \frac{2}{\pi} \sum_{n=0}^{\infty} \frac{1}{2n+1} [2\cos((2n+1)\frac{\pi}{4}) + \sqrt{2}] \\
&[e^{j(2n+1)2\pi\Delta ft}((-1)^n + 1) + e^{-j(2n+1)2\pi\Delta ft}((-1)^n - 1)]
\end{aligned}$$

Let us now consider what happens with the above equation for different values of n . When n is zero, the term corresponding to the negative frequency in the second parenthesis computes to zero and only the positive frequency is preserved resulting in single side band generation.

$n = 1$ and 2 correspond to the third and fifth harmonic respectively. For these cases $\cos[(2n+1)\frac{\pi}{4}] = \frac{-\sqrt{2}}{2}$ and so the above equation computes to zero cancelling the third and fifth harmonics. Thus, by using the above four level signal, we can cancel the third and fifth harmonics as well as achieve single sideband modulation at the same time. More generally, when n is of the form $(8k+3)$ and $(8k+5)$, $\cos[(2n+1)\frac{\pi}{4}] = \frac{-\sqrt{2}}{2}$ and hence all the corresponding harmonics will be cancelled. In summary, the four-level approximated exponential signal cancels at least the third and fifth order harmonics.

If required, subsequent harmonics can be cancelled by adding more levels. Specifically, addition of each level cancels the next higher order harmonic. For example, five voltage levels cancel the seventh harmonic and ninth harmonic is cancelled with six voltage levels. Finally, every backscatter switch has a finite delay, which provided additional filtering and automatically suppresses higher order harmonics (greater than 9), without the need for additional levels.

In our implementation, we use four levels to cancel the third and fifth order harmonics. We generate the approximated signals on the backscatter device in the digital domain. The four-level cosine signal takes one of four values $\{0.9239, 0.3827, -0.3827, -0.9239\}$. Fixing the cosine value, lets the sine take one of two values. Thus, the exponential $e^{j2\pi\Delta ft}$ can take one of eight complex values. We create these complex values by leveraging existing backscatter techniques [185, 114]

that change the impedance connected to the antenna. Fig. 3.6 shows the architecture of our system where we switch the antenna between eight different complex impedance values to generate the eight complex values corresponding to our exponential signal. Specifically, we implement the exponential wave approximation in a digital logic block called the switch mapper. It takes as input the eight phases of the VCO output to correspond to the time instances when waveform S_0 , S_1 and S_2 and their time shifted versions (corresponding to the sine wave) undergo transitions. We generate the eight phases of the clock signal by just shifting the signal in one-eighth of the time period increments and the switch mapper outputs the 3 control bits which toggle the switch between 8 impedance values to generate the approximate exponential wave. Using this technique, we can successfully cancel the mirror image as well as the third and fifth harmonics, thereby improving the spectral efficiency of backscatter systems.

3.3.4 Synthesizing LoRa Packets

LoRa achieves its high sensitivity numbers using CSS modulation. The physical layer specification for LoRa however is proprietary and is not publicly available. So, we reverse-engineer the LoRa physical layer using the patents filed by Semtech [170, 87], which is the key LoRa chipset manufacturer. We also use the Semtech 1276 starter kit [50] that provides an interface to transmit LoRa packets with various bitrates and an arbitrary payload. Finally, we analyze the transmissions from the LoRa chipsets on a USRP.

Packet Structure. Fig. 3.9 shows the structure of a LoRa packet, in the form of a spectrogram. The figure shows a sequence of repeating chirps at the beginning to represent the preamble. LoRa supports a variable length preamble between 6 and 65535 chirp symbols. To convey the end of the preamble to the receiver, the preamble ends with synchronization symbols and two and a quarter down-chirp symbols where the chirp goes from the positive to negative frequency. After down-chirps, the packet has an optional header with information about the bit rate used. This is followed by a CSS-encoded payload. An optional 16-bit CRC is sent at the end of the packet.

Bit Rates. LoRa bit rates depend on three parameters: the error correction coding rate, chirp




Figure 3.9: LoRa Packet Structure.

bandwidth and spreading factor. LoRa supports four different hamming code rates and eight chirp bandwidths of 7.8 kHz, 10.4 kHz, 20.8 kHz, 31.25 kHz, 62.5 kHz, 125 kHz, 250 kHz and 500 kHz. Further, the spreading factor can be set independently to one of seven values: 6, 7, 8, 9, 10, 11 and 12. The LoRa hardware allows these three parameters to be independently modified resulting in a total of 224-bit rate settings between 11 bps and 37.5 kbps.

We use the above packet format to synthesize LoRa packets with backscatter. We note the following.

- LoRa bandwidth and spreading factor are set a-priori and assumed known at the transmitter and receiver. The header can include information about the bit rate and payload size used, but is optional and its overhead can be reduced by statically configuring these parameters, which we do in our backscatter system.
- To achieve a high sensitivity, the phase of the LoRa chirps has to change continuously with time and has the same value at the beginning and the end of the chirp [170, 87]. To achieve this with backscatter, at each frequency, we increase the phase by $\frac{2\pi}{SF}$. This ensures that the phase at the beginning and end of each chirp symbol is the same and so we can maintain phase continuity across chirp symbols.
- To comply with FCC regulations, LoRa uses frequency hopping while using lower-data rate transmissions that occupy significant amounts of time on the channel. Specifically, while using a chirp bandwidth of 125 kHz, LoRa divides the 900 MHz ISM band into 64 channels starting at 902.3 MHz, with increments of 200 kHz. Similarly, with a chirp bandwidth of 500 kHz, LoRa divides the band into 8 channels in increments of 1.6 MHz. The transmitter performs frequency hopping between these channels to transmit data to be compliant with FCC. For backscatter, FCC

only regulates the signal source and not the backscatter device [47]². Thus, we instead hop the frequency of the single-tone transmitter. The backscatter device however uses the same frequency offset, Δf and is oblivious to this frequency hopping mechanism. To ensure that the backscattered LoRa transmissions always lie in the LoRa channels, we hop the frequency of the single-tone source at a constant frequency offset of Δf from the LoRa channels. Specifically, to generate the backscatter signals at the LoRa channels, f_1, f_2, \dots, f_n , the single-tone source performs frequency hopping across $f_1 - \Delta f, f_2 - \Delta f, \dots, f_n - \Delta f$.

3.3.5 Link-Layer Protocol

We describe how the RF source arbitrates the channel between backscatter devices. Then we explore concurrent transmissions from multiple backscatter devices.

Arbitrating the channel between backscatter devices. At a high level, we use TDMA to allocate the wireless channel between different backscatter devices. Specifically, the RF source divides time into slots and transmits the single tone signal once in each slot. Each backscatter device only transmits during its assigned slot.

The above protocol requires the backscatter devices to detect the beginning of the single tone from the RF source. To do this, our design uses existing energy detector hardware circuits that consume between 98 nW and 2.4 μ W and can detect input signals as low as -71 dBm [162, 103]. Note that this is larger than the -148 dBm LoRa sensitivity. This is however the power of the RF source at the backscatter device, while the latter is the power of the backscattered signal at the LoRa receiver. In all our experiments, the signal strength at the backscatter device was at least -45 dBm. We note that to improve the accuracy of detecting the signal from the RF source, we can also use a preamble signal like an alternating ON-OFF keying sequence of energy and no energy. This reduces the probability of confusing random transmissions in the 900 MHz band for our RF source.

To synchronize the slots across different backscatter devices, the RF device uses a unique ON-

²As a result, RFID tags do not have an FCC ID, while RFID readers have to get approved by FCC and have an FCC ID.

OFF keying sync pattern at the beginning of the TDMA round robin. This allows devices to determine the slot boundaries for a whole round-robin duration. We note that the backscatter devices do not need to have their receivers ON all the time. In particular, depending on the application, the backscatter device only transmits when it has new data. Similarly, the RF source does not transmit the single-tone signal during a time slot, if the corresponding backscatter device is not scheduled.

Concurrent LoRa transmissions using backscatter. So far, we assume that only a single backscatter device can transmit at a time. However, LoRa divides the 900 MHz band into 64 125 kHz LoRa channels each of which can have a LoRa transmission. Thus, using a single RF source, we can enable up to 64 backscatter devices to transmit concurrently on different LoRa channels to their corresponding receivers. Taking it a step further, CSS transmissions that use different spreading factors on the same LoRa channel are uncorrelated with each other [50]. Thus, in principle, we can have multiple backscatter devices with different spreading factors use the same LoRa band and transmit concurrently to their receivers, increasing the overall network throughput.

We note that we can use the energy detector circuit to transmit ACKs on the downlink channel. However, given the resilience of CSS modulation and the high sensitivity values, performing application level error correction codes is sufficient to deliver data to the receivers, with a high probability, without the need for explicit ACKs.

3.4 Hardware Implementation

We first built a proof of concept prototype using commercial off-the shelf (COTS) components and used it to characterize the performance and operating range of our system. Then, we design an integrated circuit based on the hybrid analog-digital architecture for harmonic cancellation proposed in §3.3.2 to quantify its power consumption.

COTS implementation. Our COTS implementation consists of an RF and a baseband section. The RF section is implemented on a four layer FR4 substrate and consists of three cascaded ADG904 [42] switches to create a SP8T switch network. The switch toggles a 2 dBi whip antenna [39] across the eight impedance states required for the harmonic cancellation technique.

In our optimized implementation, we use 47 pF , $3.3\text{ nH}||82\ \Omega$, $21\text{ nH}||680\ \Omega$, $8.2\text{ nH}||330\ \Omega$, $1.8\text{ k}\ \Omega$, $1.5\text{ pF}||56\text{ k}\ \Omega$, $9.1\text{ pF}||560\ \Omega$ and 3.9 pF as our impedance values to achieve the desired complex values while incurring a loss of only 4 dB in our backscatter switch network.

We implement the baseband section using the DE0-CV development board for an Altera Cyclone V FPGA [52, 56]. We generate CSS modulated packets in digital domain using Verilog HDL and output square waves corresponding to the real and imaginary components of the signal to the RF section by interfacing the two digital I/Os on the FPGA to the SP8T switch network through level shifters.

IC design. We present a LoRa backscatter IC design simulated in TSMC 65 nm LP CMOS process [38] using industry standard EDA tools. Our IC is composed of three main components: digital baseband processor, frequency synthesizer and the backscatter switch network.

Baseband Processor. It takes payload data and packet specifications such as spreading factor, bandwidth and code rate as input and synthesizes the LoRa packet in accordance to the structure described in §3.3.4. Next, it maps the bits in the packet to a frequency plan that is used by the frequency synthesizer block to create the chirp spread spectrum signal. We describe the behavioral model for LoRa packet in Verilog and use Design Compiler by Synopsis [20] to synthesize the transistor level implementation. Our baseband processor consumes $1.25\ \mu\text{W}$ to generate a LoRa packet with spreading factor of 12, 31.25 kHz bandwidth and (8,4) hamming code.

Frequency Synthesizer. We integrate the DAC and VCO in Fig. 3.6 into a single frequency synthesizer block. A frequency synthesizer or phase locked loop (PLL) takes a low frequency clock source as input and up converts it to a higher frequency. The ratio of the output frequency to the reference frequency is set by a divide ratio. We use the baseband processor output to directly control the divide ratio of the frequency synthesizer and modulate the output frequency to generate CSS modulated data. We use Johnson counter to generate the four versions of the clock which are shifted by one eighth of the time period. The four shifted versions and their complements are used to represent the eight phases of the clock signal. The frequency synthesizer consumes $4.5\ \mu\text{W}$ to generate 31.25 kHz CSS packets with a spreading factor of 12 and a 3 MHz offset.


Backscatter Switch Network. The backscatter switch network takes the eight phases of the clock

(four time shifted versions and their complements) and maps the eight phases to RF switches corresponding to their respective impedance values. The switches are implemented using NMOS transistors that toggle the antenna between eight discrete impedance states consisting of resistors and capacitors. We limit the impedances to only resistors and capacitors since inductors consume huge area and are prohibitively expensive in IC's. This results in a more constraint constellation map and 3 dB loss in backscattered signal but is a reasonable compromise for low cost. During active operation, the backscatter switch network consumes $3.5 \mu W$ to backscatter CSS modulated packets at 3 MHz offset. In total, the IC consumes $9.25 \mu W$.

3.4.1 Cost Analysis

Backscatter communication systems consume 3–4 orders of magnitude lower power than their counterpart radios [120, 114, 97, 174, 2, 185]. This is because instead of using complex and power hungry analog front end components of radios, backscatter systems use switches to modulate reflections. As a result, backscatter systems are not only extremely low power, but also consume significantly smaller area. This is a key benefit of backscatter systems because for a given technology node, the cost is directly proportional to the die area. Since active radios (e.g. Wi-Fi, LoRa, BLE) are composed of complex and area intensive components such as power amplifiers, mixers, local oscillators operating at RF frequencies, the RF front end in these systems (excluding the digital baseband) typically occupies about 10 mm^2 of die area [151, 92, 80, 127, 126]. In contrast, backscatter systems such as RFID tags are much simpler and occupy significantly lower area. The typical die area for the communication module in RFID based backscatter systems is in the order of 0.15 mm^2 [197, 93, 198, 192, 148]. As a result, the cost of a backscatter chip is at least 60 times lower.

Our LoRa backscatter system uses a similar architecture to an RFID tag with some key additions. We use an All-Digital Phased-Lock Loop (ADPLL) and an eight-stage backscatter network implemented using area efficient resistor and capacitors (metal in metal capacitors are implemented on higher metal layers on top of existing semiconductor structures) which occupies less than 0.01 mm^2 . Therefore, the die area and the cost structure for our communication module would



((a)) **Receiver Characterization.** PER as a function of RSSI values reported by the receiver for different data rates.

((b)) **Resilience to Out-of-band Interference.** PER as a function of RSSI values for different interference power levels.

Figure 3.10: We characterize our receiver by measuring PER as a function of reported RSSI in presence and absence of out of band interference.

be similar to that of an RFID tag and would be at least 60 times less than active radios.

3.5 Evaluation

We first characterize the LoRa receiver at different bit rates and sensitivities and evaluate its ability to decode the backscattered signals in the presence of single tone interference from the RF source. Then we evaluate the operating range as well as our harmonic cancellation technique.

3.5.1 Receiver Characterization

As described in §3.3.1, the data rate and sensitivity of a CSS modulated packet is determined by the spreading factor, bandwidth and coding rate. In this section, we will use the Semtech SX1276 receiver and evaluate the trade-off between sensitivity and data rates supported by the hardware.

To do this, we use a wired experimental setup to ensure that variations due to multipath do not impact our results. We use two SX1276 chipsets and configure them to be both the CSS transmitter and receiver. We connect the antenna port on the two chipsets using variable attenuators and an RF cable. As described in §3.3.4, LoRa supports 224-bit rate configurations between 11 bps and 37.5 kbps. We implement seven rates spread across 21.8 kbps and 45 bps by picking the appropriate spreading factor, bandwidth and coding rate. We use variable attenuators to change the power of

the received packet. Specifically, for each of the data rates, we start with a low attenuation value where we receive all packets and increase the attenuation till we stop receiving all packets. We transmit 1000 packets with eight-byte payload and two byte CRC and measure the packet error rate (PER).

Fig. 3.10(a) plots the PER as a function of the RSSI values reported by the receiver chipset for the seven data rates. As expected, we observe that the sensitivity is inversely proportional to the data rate of the packet. For the highest data rate of 21.8 kbps, receiver can decode packets with PER of less than 1% up to an RSSI value of -119dBm while for the lowest tested data rate of 45 bps we receive packets down to an RSSI of -142 dBm. Below an RSSI value of -142 dBm, the receiver cannot correctly decode packets for any of the tested data rates. Finally, we note that similar to Wi-Fi, the RSSI values reported by the chipset do not correspond to the actual power level of the received packet. However, the reported RSSI value is a good indicator and is proportional (though not linear) to the power level of the received packet.

3.5.2 Resilience to Out-of-band interference

Next, we evaluate how well the receiver can decode backscattered packets in the presence of out-of-band interference from the RF source. Specifically, the RF source transmits at a frequency offset from the backscatter signal and hence, can create out-of-band interference. To check this, we setup the following test bench: we connect the RF source, backscatter prototype and the receiver using a circulator setup to isolate the results from multipath. Specifically, we use a variable attenuator and connect the RF source transmitting the single tone at 905 MHz to the first port of the circulator. The LoRa backscatter hardware prototype is connected to the second port using another attenuator. The SX1276 receiver configured to receive packets at 906 MHz is connected to the third port of the circulator with an RF cable. We set the LoRa backscatter device to backscatter packets with a spreading factor of 12, bandwidth of 31.25 kHz and a (8,4) hamming code with 8-byte payload and 2 bytes of CRC at a frequency offset of 1 MHz.

The out of band interference experienced by the receiver is a function of the distance between the RF source and the receiver and depends on the deployment scenario. To cover different sce-

narios, we use the attenuator corresponding to the RF source to vary the power of the interference experienced by the receiver. We set the power of the interference at the receiver to -30 dBm, -40 dBm, -45 dBm and -50 dBm which translates to distances of 50 m, 200 m, 300 m and 500 m respectively between the RF source and receiver in free space. Next, we change the attenuator corresponding to the backscatter device to decrease the power of the backscattered packet. We measure PER as a function of the RSSI values reported by the receiver for different interference power level. Fig. 3.10(b) plots the results which show the following:

- The receiver sensitivity is inversely proportional to out of band interference. When the interference is -30 dBm, the receiver can decode packets down to an RSSI value of -139 dBm. As this interference reduces to -45 dBm and -50 dBm, which is closer to values seen in our deployments, the receiver can correctly decode packets down to RSSI values of -146 and -148 dBm respectively.
- With out-of-band interference, the receiver can decode packets at lower RSSI numbers than in §3.5.1. This is because the RSSI values reported by the chipsets are just an approximate indicator of link quality and do not represent the absolute measure of power or sensitivity.
- In setup with -40dBm and -50 dBm out of band interference, an offset of 1 MHz at the backscatter device achieves sensitivities that were within a few dB of the theoretical limit. We can increase the offset for better sensitivities when facing a higher interference power.

3.5.3 Operational Range

We evaluate the operating range of the LoRa backscatter system in two different scenarios. We use the SX1276 LoRa development kit connected to a power amplifier as the RF source and configure it to transmit 30 dBm into a 6 dBi patch antenna at 915 MHz. This is the maximum power permitted by FCC on the 900 MHz ISM band. We configure the backscatter device to transmit LoRa packets at a frequency offset of 3 MHz with a spreading factor of 12, bandwidth of 31.25 kHz and a (8,4) hamming code with 3-byte payload and 2 byte CRC. The SX1276 transceiver chip decodes packets received at 918 MHz and we log the RSSI values for packets that pass CRC.

Deployment scenario 1. The first scenario we consider is a deployment where the RF source

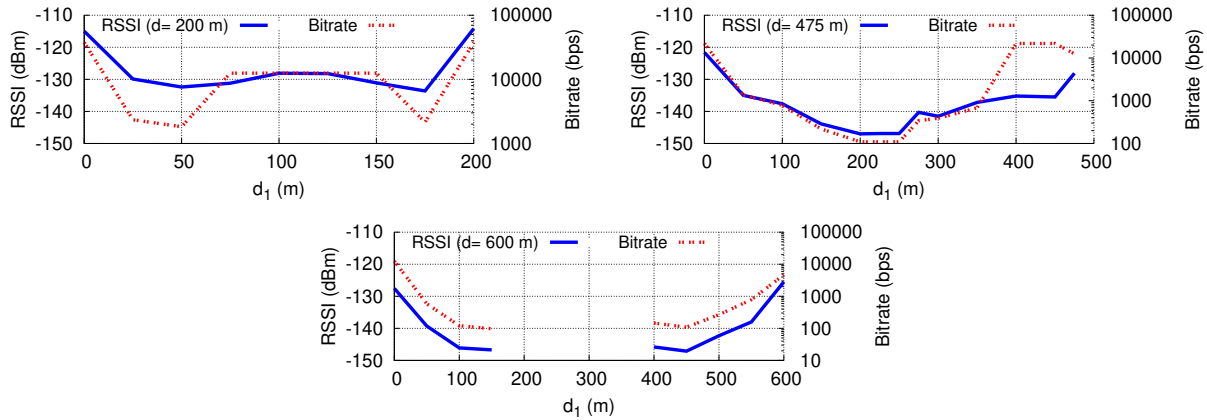


Figure 3.11: **RSSI in Deployment Scenario 1.** d is the distance between the RF source and receiver. We move the backscatter device along the line joining them. This figure also shows the corresponding LoRa bit rate at which we successfully receive all our ten packets from the backscatter device without any loss.

and the receiver are separated from each other and the backscatter device can be at any location between them. To test this, we place the RF source and the receiver at a distance d , as shown in Fig. 3.11, and move the backscatter device in a straight line between them. At each location of the backscatter device, we measure the RSSI value reported by the receiver. Due to the large operating range, we ran the experiments on a straight road next to open fields. Fig. 3.11 plots the RSSI values for three different distances between the signal source and the receiver. The x-axis represents d_1 , the distance between the RF source and the backscatter device. The plots show the following,

- We get a low RSSI when the backscatter device is at the midpoint between the RF source and the receiver. This is because the signal from the RF source attenuate as $1/d_1^2$ before arriving at the backscatter device. The signals generated by backscattering these attenuated signals further attenuate as $1/d_2^2$. Thus, the backscatter signal strength at the receiver scales as $1/d_1^2 d_2^2$ which is minimum when $d_1 = d_2$.
- Our system operates at all locations up to a maximum separation of 475 m between the RF source and the receiver. At 600 m, the backscatter device only works close to either the RF source or the receiver. This shows that our design has an operational range of 475 m.

Deployment scenario 2. In the second scenario, we fix the distance between the backscatter device

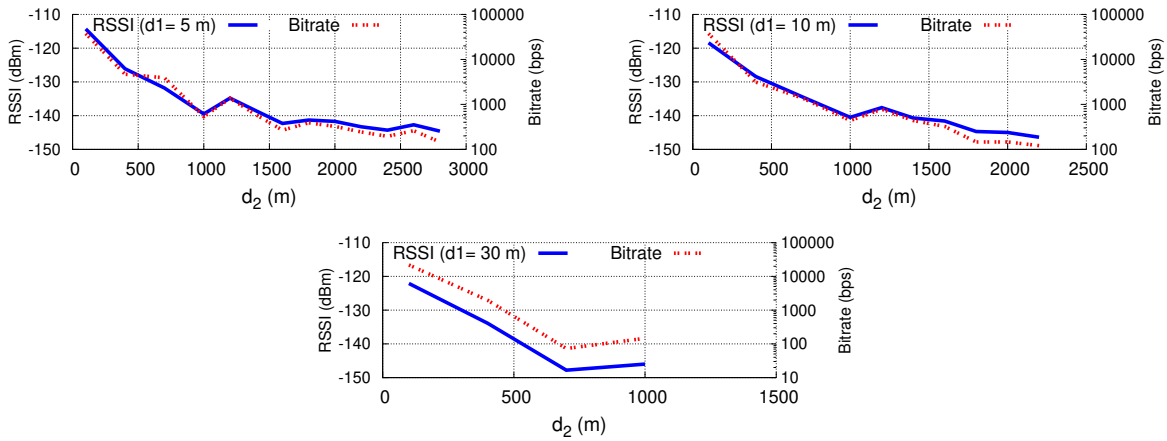


Figure 3.12: **RSSI in Deployment Scenario 2.** d_1 (d_2) is the distance between the backscatter device and RF source (receiver). We fix the location of the backscatter device and RF source and move the receiver away from the backscatter device.

and the RF source and move the receiver away from the backscatter device and measure RSSI. This is a scenario where the backscatter device is close to the RF source — an example deployment is one where the sensor is in the wall and the RF source is placed nearby where it can be plugged in. Our initial experiments in this setup showed that we could achieve very large ranges. So, we test this scenario on the same road, which spans multiple kilometers. Specifically, we set the RF source and backscatter device on the roadside. We change the separation between the RF source and the backscatter devices between three values of 5, 10 and 30 meters. For each separation value, we drive the receiver away from the backscatter device, alongside the road and measure the RSSI of CRC passed packets at the granularity of 200 m.

Fig. 3.12 plots RSSI as a function of the distance between the backscatter device and the receiver. It shows results for three different distances between the RF source and the backscatter device. The plots show that the receiver can receive packets at a distance of 2.8 km from the backscatter device, when the backscatter device and the RF source are separated by 5 m. We note that the receiver could decode all transmitted packet at all reported locations. When we increase distance between the RF source and the backscatter device, the operating distance reduces. At 30 m separation, the receiver could receive packets up to 1 km.

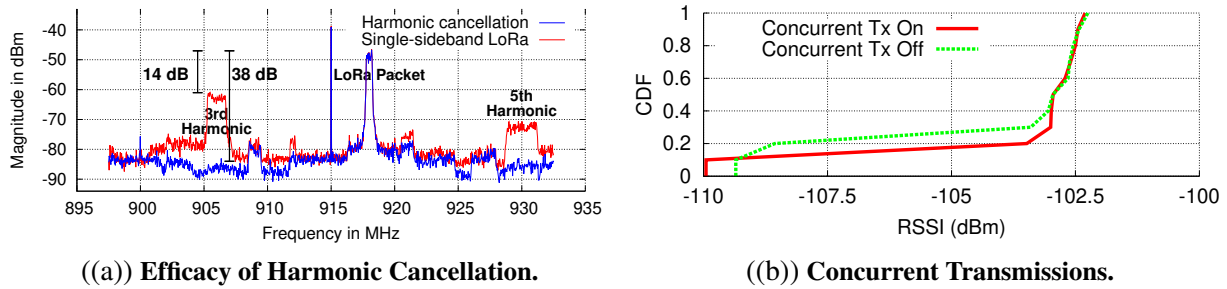


Figure 3.13: The plot on the left demonstrates the efficacy of our harmonic cancellation technique. The right plots CDF of RSSI of a LoRa backscatter device in the presence and absence of another LoRa backscatter device concurrently transmitting in adjacent band.

3.5.4 Efficacy of harmonic cancellation

Next, we evaluate how well our harmonic cancellation technique works in practice. To do this, we capture the spectrum of the signal generated using our harmonic cancellation backscatter hardware and compare it to a baseline hardware that does not implement harmonic cancellation. Specifically, for the baseline we replicate the state-of-the-art single-sideband Wi-Fi hardware from [114] and adapt it using the techniques in §3.3.2 to generate LoRa packets. We use the Semtech SX1276 development kit to generate the single tone signal at 915 MHz. Since the third and fifth harmonics span a wide band, we use a spectrum analyzer to capture the signal. To avoid interference from other wireless devices, we connect the Semtech kit to port 1 of a circulator, the backscatter hardware to port 2 and measure the spectrum of the backscattered signal by connecting a spectrum analyzer to the third port of the circulator. We set our backscatter hardware to transmit LoRa packets with spreading factor of 6, bandwidth of 500 kHz and a (8,4) hamming code with 8-byte payload and 2 byte CRC at a 3 MHz offset.

The red line in Fig. 3.13(a) shows the spectrum of the backscattered LoRa packet with the single-sideband hardware. The plot shows a single tone signal at 915 MHz, which is the RF source and the desired LoRa packet at 918 MHz. In addition, we see the third and fifth harmonics at 906 MHz and 930 MHz respectively. These are 14 dB and 22 dB lower than the desired backscat-

ter LoRa packet. These numbers are 4.7 dB and 7 dB lower than our analysis with square wave in §3.3.2 because in frequency modulated systems, the higher order harmonics are spread across frequency proportional to the order of the harmonic. Note that the third and fifth harmonics occur on opposite sides of the desired LoRa transmission at 918 MHz because of single-sideband architecture. The key observation however is that the third and fifth harmonics create interference in the 900 MHz band.

The blue line shows the spectrum of the backscattered LoRa packet with our harmonic cancellation backscatter hardware. The spectrum shows that the third (906 MHz) and fifth harmonics (930 MHz) are 38 dB lower than the backscattered LoRa packet at 918 MHz and is close to noise. This demonstrates the efficacy of our harmonic cancellation technique. We note that while in theory we should get perfect cancellation, in practice backscatter hardware built using switches and passive components have tolerances and variances that introduce small errors.

3.5.5 *Concurrent Backscatter Transmissions*

Finally, we present a proof-of-concept evaluation of LoRa backscatter when multiple backscatter devices concurrently transmitting signal from a single RF source. We use two backscatter devices and configure them to continuously transmit packets at different frequency offsets of 0.75 MHz and 1 MHz. Since the two devices create transmissions on different LoRa bands, they can concurrently transmit to their receivers. We set the RF source to transmit a single tone at 905 MHz and deploy the system in a large atrium measuring 104 by 32 feet. We place the RF source and the two receivers at either ends of the room. We fix the location of the first backscatter device at the center of the atrium. We move the second backscatter device (operating at 1 MHz offset) across ten different locations of the atrium. Fig. 3.13(b) shows the CDF of the RSSI values of packets transmitted by the first backscatter device in the presence and absence of concurrent transmissions by the second device. The plots show that concurrent transmissions have negligible impact on the performance of the first backscatter device.

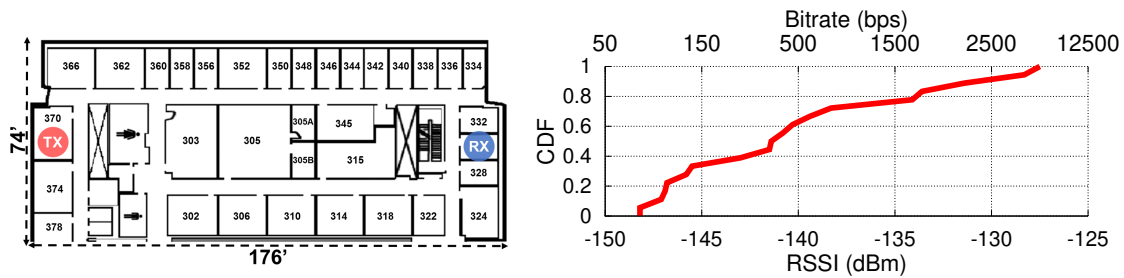


Figure 3.14: **Office Deployment.** We receive backscattered packets across one floor of an office building with $13,024 \text{ ft}^2$ (1210 m^2) area.

3.6 Application Deployments

We show two classes of applications enabled by LoRa backscatter. The first set of applications home/office sensing and precision agriculture leverage the large operational range of our communication system. We then demonstrate the use of LoRa backscatter in applications with unfavorable RF propagation conditions such as implantable and epidermal devices. In all our application deployments, we set the backscatter device to transmit with a spreading factor of 12, bandwidth of 31.25 kHz and a (8,4) hamming code with three-byte payload and two bytes of CRC. We use the SX1276 receiver to log the RSSI value of packets that pass the CRC.

3.6.1 Wide-Area Applications

We deploy our system in two wide-area applications: home sensing and precision agriculture and evaluated the system in real world deployment scenarios.

Whole-Home and Office Sensing

Prior backscatter techniques such as [120, 119] have tried to replace power-consuming radios in home sensing applications with backscatter but suffer from limited range and do not provide the coverage required for a practical home deployment. We deploy our system in a $4,800 \text{ ft}^2$ (446 m^2) home spread across three floors in a major metropolitan US city. The layout and floor plan of the

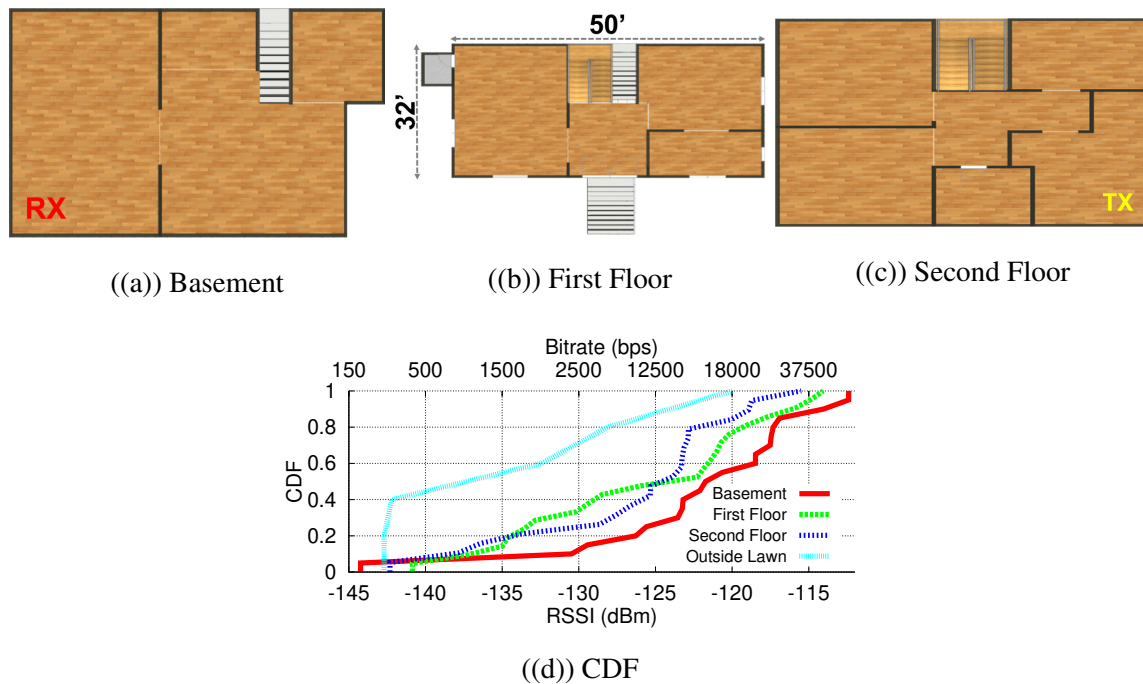


Figure 3.15: **Home Deployment.** We receive packets across the $4,800 \text{ ft}^2$ (446 m^2) house spread across three floors.

house is shown in Fig. 3.15. We put the RF source at one corner of the house on the third floor and place the receiver in the basement. We move the backscatter device across the three floors of the occupied house in 6×6 -foot grid increments and report RSSI of the received packets at each location.

Fig. 3.15 plots the RSSI of the received backscattered packets at each of the three floors as well as the outside lawn area. We use a 1 MHz offset on the backscatter device. Our results show that across the entire house, our system was able to achieve RSSI values greater than -144 dBm which translates to reliable wireless coverage across the entire house at 45 bps, with a single RF source and receiver. These rates are sufficient for most home automation devices such as temperature sensors that transmit small packets sporadically.

Next, we deploy our system with a 3 MHz offset on the backscatter devices in an office space spanning 13024 ft^2 (1210 m^2). As shown in Fig. 3.14, the deployment spans 41 offices that are separated by double sheet-rock (plus insulation) walls with a thickness of approximately 5.7 inch

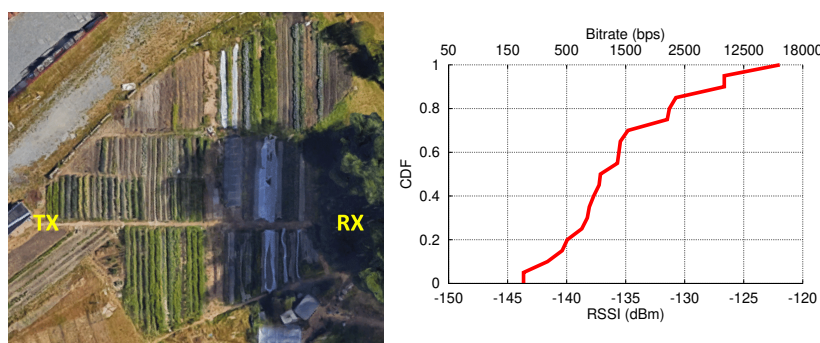


Figure 3.16: **Precision Agriculture.** We receive backscattered packets across a one-acre (4046 m^2) farm.

(14.5 cm). We place a RF source and receiver on the two end of the space as shown in the figure. We note that the receiver is behind a heavily insulated set of rooms including the restroom and supply closets with concrete and metal structures separating them. We move the backscatter device into different offices across the whole floor in 26 by 26-foot grid increments and report the RSSI of received packets at each location. Fig. 3.14 plots the CDF of the RSSI of received backscattered packets across the whole space, demonstrating wide-area backscatter coverage in significant multi-path environments.

Precision Agriculture

We deploy LoRa backscatter in a one-acre (4046 m^2) vegetable farm owned by our organization. We set the RF source and the receiver at the opposite ends of the farm, whose aerial view is shown in Fig. 3.16. We divide the farm into 45 ft by 45 ft grids and placed the backscatter device with a 1 MHz offset at ground level between plants and bushes. We measure RSSI values by placing the backscatter device across 20 locations at the center of the grids. These locations are a combination of line of sight and non-line of sight positions to the RF source and/or receiver due to presence of green houses and storage sheds across the farm.

Fig. 3.16 plots the CDF of the RSSI of the packets received across the entire farm. The plot shows that the median RSSI is -137 dBm and the minimum RSSI was -143 dBm . This demonstrates that using a single RF signal source and receiver, we could provide reliable backscatter

communication across a one-acre ($4046 m^2$) farm. We note that typical farms are much larger and might require more RF signal sources and receivers. However, since Semtech SX1276 transceiver is relatively cheap and the cost of the backscatter devices is in the order of a few cents, we could achieve cost and power saving by deploying backscatter in agriculture applications.

3.6.2 RF-Challenged Applications

A key advantage of CSS is its high sensitivity, which enables long-range operation. This sensitivity also enables long ranges in extremely challenging RF environments such as implantable and body worn devices.

Smart Contact Lens

Smart contact lens can measure vital indicators such as glucose, sodium and cholesterol in tears and enable long term unobtrusive real-time tracking of such vital parameters [129, 196]. Although progress has been made in miniaturization of the sensor IC, antenna and packaging, real time communication remains a bottleneck [114]. Radio communication is not applicable since it consumes orders of magnitude higher power than available on miniature batteries on contact lens form factor devices [155, 83, 129]. Since the environment is highly unfavorable to RF propagation, only extremely short range backscatter of tens of centimeters has been feasible using both custom hardware [129, 196] and standards-compliant Wi-Fi/BLE hardware [114]. Such short ranges limit applicability and renders backscatter less practical.

We leverage LoRa backscatter to demonstrate that a smart contact lens can communicate using CSS modulated backscatter at orders of magnitude larger distances than was feasible with prior approaches. We built a contact lens form factor antenna shown in Fig. 3.17(a). The antenna is a 1 cm diameter loop 30 AWG wire encapsulated between two soft contact lenses for biocompatibility and structural integrity. We immerse the antenna in contact lens solution and connect it to our COTS prototype.

We deploy the contact lens in a large atrium measuring 104 ft by 32 ft. We place the RF source

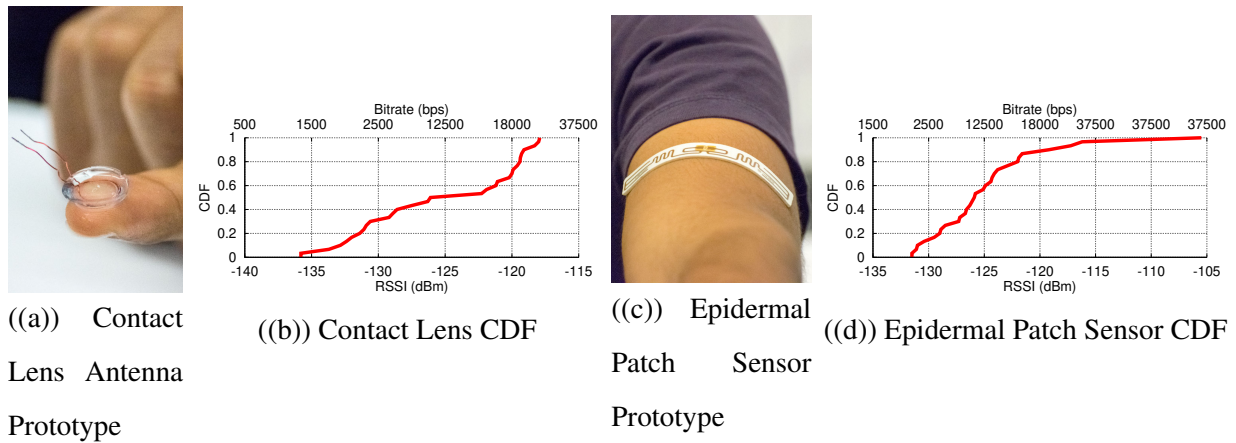


Figure 3.17: **Smart Contact Lens and Flexible epidermal patch sensor.** We show the RSSI and data rate distribution across an entire $3,328 \text{ ft}^2$ (309 m^2) atrium.

and the receiver at the two ends of the room. We set the backscatter device to transmit LoRa packets at a 1 MHz offset and place the contact lens form factor antenna at the center of 12 ft by 10 ft grids. We log the reported RSSI value of the received packets at each of the tested locations. Fig. 3.17(b) plots the CDF of the received RSSI. Our results show reliable connectivity across the entire atrium, which is order of magnitude larger range than prior designs [129, 114].

Flexible Epidermal Patch Sensor

Finally, flexible patch sensors can be used to monitor temperature, sweat, ECG and other vital signs in real time [83, 156]. These sensors are worn on the body and are in unfavorable RF environments since the antenna gets significantly detuned and degrades the link quality[105].

We use our system to prototype a form factor flexible epidermal patch sensor. To build our prototype, we use the sticker form factor antenna of an Alien swiggle RFID tag [2] and connect it to our COTS hardware. Using this approach, we can leverage the cost and size of RFID tags and create small, inexpensive and reliable epidermal patch sensors. We use a matching network to tune the impedance of the antenna when it's in contact with the human skin. We attach the sticker form factor antenna prototype on a subject's hand and test the device across the $3,328 \text{ ft}^2$ (309 m^2)

atrium using the same deployment as the contact lens prototype with a 1 MHz backscatter offset. Fig. 3.17(d) plots the CDF of RSSIs for the received packets. Our results show that we can provide reliable connectivity with an RSSI greater than -132 dBm.

3.7 Discussion and Conclusion

We present the first backscatter system that can achieve the ranges required to enable wide-area communication. In this section, we outline avenues for future research.

Achieving higher data rates. Our current design achieves all the bit rates supported by LoRa. However, an interesting research direction is to achieve higher data rates at the desired ranges which can be accomplished by using a combination of multiple frequency bands to concurrently transmit or multiple antenna designs.

LoRa-independent design. A reader might ask a very relevant question: does this technology depend on the wide adoption of LoRa? The answer is that if LoRa is successful, our technology can ride on top of its adoption. However, the key component of our wide area backscatter design is the CSS modulation, which can provide long ranges and high sensitivities independent of LoRa.

Sensitivity to orientation, multipath and obstacles. Unlike RFID and existing backscatter systems such as Passive Wi-Fi, LoRa backscatter can operate down to sensitivity of -149 dBm and in the presence of strong in-band and out of band interference. As a result, in typical applications, LoRa backscatter has sufficient wireless link budget margin which allows it to be robust to signal loss due to antenna orientation mismatch, multipath and obstacles in the path. We demonstrate this characteristic of our system by evaluating RF challenging applications such as contact lens, epidermal patch sensor and wide area deployments such as agriculture, home in §3.6.

Networking LoRa backscatter devices. The focus of this paper was to design and implement the LoRa backscatter physical layer. In §3.3.5 we described how LoRa backscatter uses TDMA to communicate across multiple-backscatter devices. However, since TDMA has an upper bound on throughput, we believe that there are several research opportunities in networking 100-1000s of tiny devices that use LoRa backscatter communication modules near each other. This is a topic of future research and out of scope of this paper.

RF Power harvesting. LoRa backscatter is independent of how it is powered. We noted that our current design can be powered using very small solar cells, button cells as well as printed batteries, given its low power consumption. One can also explore research opportunities for long-range power harvesting [181] from the RF source.

Chapter 4

SUPPORTING APPLICATIONS WITH LARGE-SCALE BACKSCATTER DEVICES

4.1 Introduction

The last few years have seen rapid innovations in low-power backscatter communication [132, 202, 120, 114, 119, 141], culminating in long range and reliable backscatter systems [178, 159, 187]. These designs enable wireless devices to communicate at microwatts of power and operate reliably at long ranges to provide whole-home or warehouse coverage. To achieve this, they employ low-power coding techniques such as chirp spread spectrum, to decode weak backscatter signals below the noise floor [159, 178] and deliver long ranges.

While these long range backscatter systems are promising for enabling power harvesting devices (e.g., solar and vibrations) as well as cheap and small Internet-connected devices that operate on button-cells or flexible printed batteries, they primarily work at the link layer and are not designed to scale with the number of devices — all these prior designs [178, 159, 187] are evaluated in a network of 1–2 devices.

Our goal in this paper is to design a network protocol that enables these low-power backscatter networks to support hundreds to thousands of concurrent transmissions. This is challenging because the resulting design must operate reliably with weak backscatter signals that can be close to or below the noise floor. To this end, we present NetScatter, the first wireless protocol that can scale to hundreds and thousands of concurrent transmissions from backscatter devices. Our design enables concurrent transmissions from 256 devices over a bandwidth of 500 kHz. Consequently, it can support transmissions from a thousand concurrent backscatter devices using a total bandwidth of only 2 MHz.

Our key innovation is a distributed coding mechanism that satisfies four key constraints: i) it enables hundreds of devices to concurrently transmit on the same frequency band, ii) it can operate




Figure 4.1: **Large-Scale Network Deployment of Backscatter Devices.** We deploy 256 backscatter devices across a floor of an office building covering multiple rooms.

below the noise floor while achieving reasonable bitrates, iii) its coding operation can be performed by low-power backscatter devices, and iv) it can decode all the transmissions at the receiver using a single FFT operation, thus minimizing the receiver complexity.

We introduce *distributed chirp spread spectrum coding*, which uses a combination of chirp spread spectrum (CSS) modulation and ON-OFF keying. In existing CSS systems (e.g., LoRa backscatter [178]), the AP transmits a continuous wave signal which each device backscatters and encodes bits using different cyclic shifts of a chirp signal. In contrast, in our distributed CSS coding, we assign a different cyclic shift of the chirp to each of the concurrent devices. Each device uses ON-OFF keying over these cyclic shifted chirps to convey bits, i.e., the presence and absence of the corresponding cyclic shifted chirp correspond to a ‘1’ and ‘0’ bit respectively, as shown in Fig. 4.2. Note that in comparison to existing CSS systems where each device transmits $\log_2 N$ bits using N cyclic shifts, our distributed design enables N concurrent devices, each of which transmits a single bit, using ON-OFF keying. Thus, our design transmits a total of N bits within a chirp duration, providing a theoretical gain of $\frac{N}{\log_2 N}$.

Our design leverages the fact that creating concurrent cyclic-shifted chirps at a single device requires distributing its transmit power amongst all the cyclic shifts, which reduces the ability of the receiver to decode each chirp. Instead we generate concurrent cyclic-shifted chirps across a

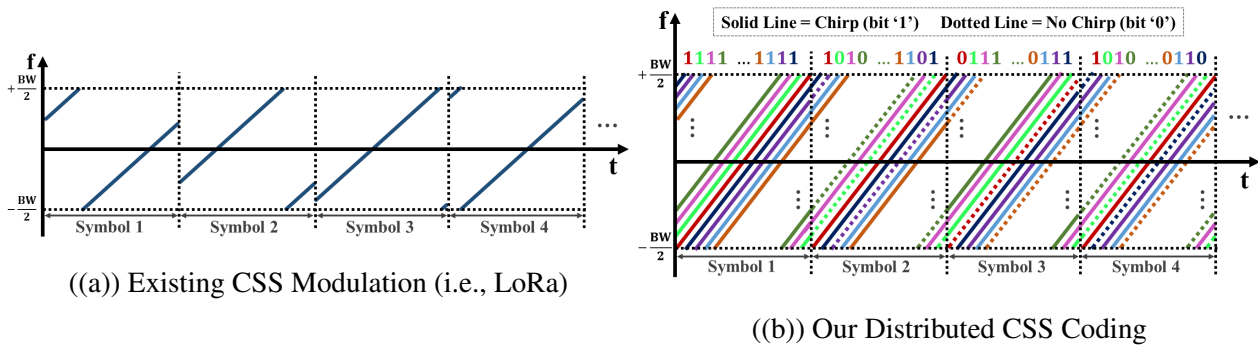


Figure 4.2: **NetScatter Overview.** In traditional CSS systems, a single device uses different cyclic shifts to convey bits. In distributed CSS coding, each cyclic shift is assigned to a different backscatter device. Each device then uses the presence and absence of cyclic shift to send ‘1’ and ‘0’ bits.

distributed set of low-power devices in the network. This allows us to efficiently leverage the coding gain provided by chirp spread spectrum under the noise floor [85]. Further, we can decode all the concurrent transmissions using a single FFT operation, since cyclic shifting the chirps in the time domain translates to offsets in the frequency domain.

Using the above distributed coding mechanism in practice, however, is challenging for two key reasons.

- *Near-far problem.* A fundamental problem with enabling concurrent transmissions is that signals from a nearby backscatter device can overpower a farther concurrent device. To address this issue, we introduce two main techniques. First, we present a power-aware cyclic shift allocation technique in §4.3.2, where lower SNR devices use much different cyclic shifts than higher SNR devices. We show that such an allocation can allow backscatter devices that have an SNR difference of up to 35 dB to be concurrently decoded. Second, to account for channel variations over time, we develop a zero-overhead power adaptation algorithm where backscatter devices use reciprocity to estimate their SNR at the AP, using the signal strength of the AP’s query message. The backscatter devices then adjust their transmission power to fall within the tolerable SNR difference. Since this calibration is done independently at each backscatter device using the AP’s query, it does not require additional communication overhead at the AP.

- *Timing synchronization.* The above design requires all the devices to start transmitting at the same time so as to enable concurrent decoding. However, hardware variations and propagation delays of different devices can make it challenging for hundreds of devices to be tightly synchronized in time. To avoid this coordination overhead, we leave gaps between cyclic shifts to ensure that concurrent devices can be decoded. We explore the trade-off between the required gaps and the chirp bandwidths in §4.3.2.

We implement NetScatter on a testbed of backscatter devices. We create backscatter hardware that implements NetScatter and includes circuits to perform automatic power adaptation before each transmission. We deploy our backscatter testbed with 256 devices in an office building spanning multiple rooms as shown in Fig. 4.1. We implement our receiver algorithm using USRP X-300 software-defined radios. Our results reveal that over a 256 node backscatter deployment, NetScatter achieves a 14–62x gain over prior long-range backscatter systems [178] for its end-to-end link layer data rates. The key benefit however is in the network latency which sees a reduction of 15–67x.

Contributions. Our paper demonstrates, to the best of our knowledge, the first network protocol that achieves orders of magnitude more concurrent transmissions than existing backscatter systems. The closest work to our design is Choir [96] *in the radio domain*, which decodes concurrent transmissions from 5–10 LoRa radios at a software radio. Choir leverages frequency imperfections to disambiguate between LoRa radios. However, backscatter devices achieve low-power operations by running at a lower frequency (1-10 MHz) than radios (900 MHz) and thus have much smaller frequency differences between backscatter devices. This severely limits the ability to rely on frequency imperfections to disambiguate between a large number of backscatter devices (see §4.2.2). In contrast, our distributed chirp spread spectrum coding mechanism provides a systematic approach to enable large scale backscatter networks.

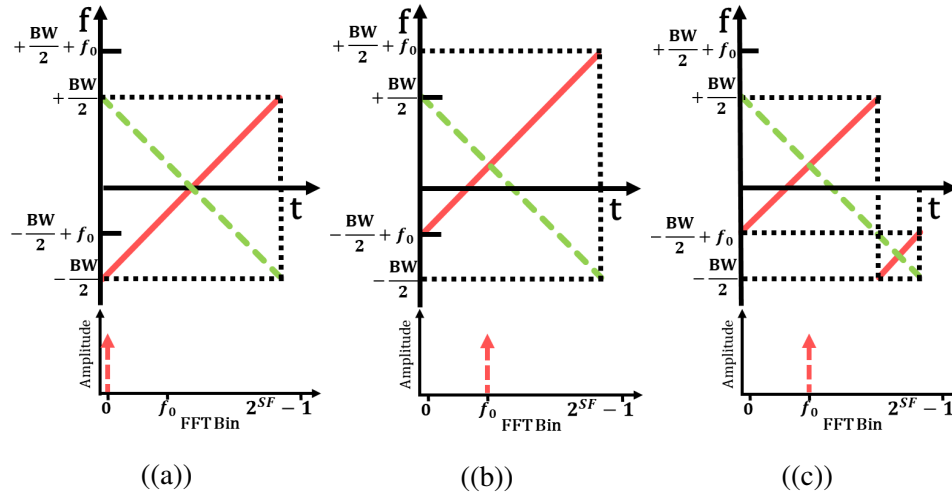


Figure 4.3: **CSS Primer.** We show upchirp and downchirp symbols and FFT results of their multiplication. (a) Baseline upchirp symbol, (b) frequency shifted upchirp symbol and (c) cyclically shifted upchirp symbol.

4.2 CSS Primer & Existing Approaches

4.2.1 Primer on Chirp Spread Spectrum

In CSS, data is modulated using linearly increasing frequency signals or upchirps. The receiver demodulates these symbols in a two step process. First, it de-spreads these upchirp symbols by multiplying them by a downchirp and it then performs an FFT on the de-spread signal. Since the slope of the downchirp is the inverse of the slope of the upchirp, multiplication results in a constant frequency signal, as shown in Fig. 4.3(a). Thus, taking an FFT on this will lead to a peak in an associated FFT bin. Changing the initial frequency of an upchirp will result in a change in the demodulated signal's FFT bin peak index which corresponds to the initial change in frequency, as shown in Fig. 4.3(b). This property is used to convey information. When the sampling rate is equal to chirp bandwidth (BW), frequencies higher than $\frac{BW}{2}$ will alias down to $\frac{-BW}{2}$ as shown in Fig. 4.3(c). This means cyclically shifting in time is equivalent to changing the initial frequency and thus to conserve bandwidth, CSS uses cyclic shifts of the chirp in the time-domain instead of frequency shifts. This means that to modulate the data we just need to cyclically shift the baseline upchirp in time. Note that one can transmit multiple bits within each upchirp symbol. In particular,

say the receiver performs an N point FFT. It can distinguish between N different cyclic shifts each of which corresponds to a peak in one of the N FFT bins. Thus, we can transmit $SF = \log_2 N$ bits within each upchirp symbol, where SF is called the spreading factor.

Based on above explanations, CSS can be characterized by two parameters: chirp bandwidth/sampling rate and spreading factor. Thus, each chirp symbol duration is equal to $\frac{2^{SF}}{BW}$ and the symbol rate is $\frac{BW}{2^{SF}}$. Since CSS sends SF bits per symbol, the bitrate is equal to $\frac{BW}{2^{SF}} SF$. This means increasing SF or decreasing BW decreases the bitrate. Further, the sensitivity of the system depends on the symbol chirp duration and increases with SF and decreases with BW .

4.2.2 Existing Collision Approaches

While existing CSS-based backscatter systems do not support collision decoding, we outline potential approaches to deal with collisions in CSS *radio* systems, i.e. LoRa, and explore whether they can be adopted for backscatter.

Using different spreading factors. One way to enable concurrent transmissions is to assign different spreading factors to each device. There are three problems with using multiple spreading factors in the same network: i) the receiver needs to use multiple FFTs and downchirps with different spreading factors to despread upchirp symbols of different devices, which increases the receiver complexity with the number of concurrent transmissions, ii) in LoRa, different BW and SF can be concurrently decoded without sensitivity degradation, only if the chirp slope is different [175]. Specifically, if two chirp symbols transmitted concurrently with different BW and different SF , which result in the same chirp slope, $\frac{BW^2}{SF}$ (shown in Fig. 4.6 as well), the receiver cannot decode their concurrent transmissions. This results in only 19 different BW and SF pairs that could be used concurrently, iii) further, requiring receiver sensitivity better than -123 dBm and bitrates of at least 1 kbps limits these concurrent configurations to only 8, which does not support hundreds of concurrent devices on a 500 kHz band. Note that ignoring the receiver complexity, this approach is orthogonal to our design since we could in principle run multiple concurrent NetScatter networks with the above 8 SF and BW pairs. Evaluating this is not in the scope of this paper.

Choir [96]. Recent work on decoding concurrent LoRa transmissions leverages the hardware




Figure 4.4: **Choir Approach.** We evaluate FFT variation of chirp symbols when $BW = 500 \text{ kHz}$ and $SF = 9$ for both active LoRa radios and backscatter devices.

imperfections in radios to disambiguate between multiple transmissions. Specifically, radios have slight variations which result in timing and frequency offsets, which translate to fractional shifts in the FFT indexes. Choir [96] uses these fractional shifts, with a resolution of one-tenth of an FFT bin, to map the bits to each transmitter. However, as demonstrated in [96], in practice this approach does not scale to more than 5 to 10 concurrent devices. To understand this limitation in theory, consider N concurrent devices. The probability that each of these transmitters has a different FFT peak index fraction, given the resolution of one-tenth of an FFT bin, is equal to $\frac{10!}{(10-N)!10^N}$. When N is 5 this probability is only 30%. Moreover, if any two transmitters use the same cyclic shifted upchirp symbol at the same time, it will result in a collision that cannot be decoded. In the case of LoRa modulation, if there are N transmitters and assuming each device transmits a random set of bits during each symbol interval, the probability of two transmitters using the same cyclic shift is equal to: $1 - \prod_{i=1}^N (1 - \frac{i-1}{2^{SF}})$ which is approximately $\frac{N(N-1)}{2^{SF+1}}$.

For $SF = 9$ and $N = 10$, this probability is around 9%. This means that there is around 9% probability that within each CSS symbol, two transmitters will use the same upchirp cyclic shift, which the receiver cannot disambiguate. This probability increases to 32% with 20 devices, preventing concurrent decoding of a large number of transmitters.

Moreover, Choir is based on oscillator imperfection causing frequency variation on different devices, and Choir cannot differentiate two concurrent transmissions if both transmissions fall into same FFT bin fraction. Choir uses an active radio system which generates frequencies in 900 MHz




Figure 4.5: **Bandwidth Aggregation.** Here we use an aggregate bandwidth of $2BW$ but each device transmits only using BW . Upchirps with different cyclic shifts shown in different colors. Each upchirp is assigned to a device.

band. However, since backscatter systems are designed to consume less power and only generate baseband signals, their output frequency is less than 10 MHz. Now, in the ideal scenario where the same crystal oscillator is used for both radios and backscatter devices, the frequency variation of the backscatter devices is 90 times smaller than radios and can be even less than 1 FFT bin depending on the SF and BW . This means a backscatter network cannot use all 10 different FFT bin fractions that Choir have used. Fig. 4.4 shows CDF of FFT bin variation for our actual backscatter hardware which are recorded over time. This results show that FFT variation is always less than a third of an FFT bin. Thus, Choir cannot enable large concurrent transmissions with backscatter.

In conclusion, the desired solution must satisfy three constraints: 1) ability to differentiate between FFT peaks corresponding to different backscatter devices, 2) ability to associate the FFT peaks to the corresponding devices, and 3) ensure that two devices do not use the same FFT peak at the same time. NetScatter design satisfies all these constraints.

4.3 NetScatter Design

4.3.1 Distributed CSS Coding

Our approach is to take advantage of low-power and high sensitivity of CSS modulation to design a communication and networking system that enables hundreds of backscatter devices to transmit at the same time.

At a high level, we use a combination of CSS modulation and ON-OFF keying to enable con-

current transmissions. Our intuition is as follows: if we look at the FFT plots of Fig. 4.3, all the FFT bins except one bin are empty; however these empty bins could be utilized for orthogonal transmissions. While it is difficult to design low-power backscatter devices that can transmit multiple cyclic shifts at the same time, we can leverage all these empty bins by having different devices transmit different shifts and make use of the unused FFT bins. In particular, each device is assigned to a particular cyclic shifted upchirp symbol. It sends data by either sending the upchirp symbol or not sending it, i.e., by using ON-OFF keying of its assigned cyclic shifted chirp. Since, there are 2^{SF} FFT bins, ideally we can support 2^{SF} concurrent transmissions. This modulation will satisfy the above three requirements. The peaks can be differentiated and assigned to their corresponding devices. Moreover, none of them will use the same FFT bin at the same time.

We note the following about our distributed design.

- **Receiver complexity.** The received signal is composed of multiple transmissions. They can be demodulated by despreading with a baseline downchirp multiplication and performing an FFT operation. Then, we can determine the presence and absence of a peak in each FFT bin and find if the corresponding backscatter device is sending '0' or '1'. The key point is that the process of despreading and performing FFT, which are the major contributors of the demodulation process and provide a coding gain for each of the backscatter devices enabling them to operate below the noise floor, are being done once and do not depend on the number of concurrent transmissions. This means that the receiver complexity is nearly constant with the number of devices.
- **Throughput gain.** In our approach, ideally there can be as many as 2^{SF} transmissions at each symbol period. Since each backscatter device uses ON-OFF keying over a symbol, their individual data rate is $\frac{BW}{2^{SF}}$. Thus, the aggregate network throughput is equal to BW . In comparison, LoRa have a throughput of $\frac{BW}{2^{SF}}SF$. Thus, we can achieve a throughput gain of $\frac{2^{SF}}{SF}$, which shows that the gain exponentially increases with the SF value used in the system. This is expected since the number of concurrent devices we can support is an exponential function of SF, i.e., 2^{SF} .
- **NetScatter and CDMA.** Our distributed CSS coding can be thought of as code-division multiplexing mechanism that is low-power and where each of the 2^{SF} cyclic shifts is in an orthogonal set




Figure 4.6: **Timing Mismatch**, in detecting beginning of a chirp symbol and its translation to FFT bin variation.

of codes in a CDMA system. These orthogonal codes are then assigned to 2^{SF} different backscatter devices which enables 2^{SF} concurrent transmissions.

- **Gain in the context of Shannon capacity.** A key gain we are achieving in our design stems from using the power across all the concurrent backscatter devices. Specifically we note that the Shannon capacity of a multi-user network that operates under the noise floor linearly increases with the number of devices. Said differently, the multi-user capacity of an access point network is given as [186], $C = BW \log_2(1 + \frac{NP_S}{P_N})$. Here BW is the channel bandwidth, P_N and P_S are the noise and signal power and N is the number of concurrent devices. At SNRs below the noise floor, the above equation can be approximated as $\frac{BW}{\ln(2)} \frac{NP_S}{P_N}$, since $\ln(1+x) \approx x$ when x is small. This means that for systems that operate below the noise-floor, the network capacity scales linearly with the number of users. This linear increase stems from the fact that the N backscatter devices put in N times more power back to the AP than a single backscatter device.

- **Bandwidth aggregation.** The bitrate achieved by each backscatter device in our distributed design is given by $\frac{BW}{2^{SF}}$ and the number of concurrent devices is 2^{SF} . Thus, while we can increase the number of devices by increasing SF, it would decrease the bitrate of each device. Thus, to increase both the bitrate and the number of device we should increase the bandwidth, BW . Say, we want to support twice the number of devices while maintaining the same bitrate by using twice the bandwidth. This can be achieved in two ways. First, we can use two filters and independently

operate two sets of devices across the two bands. This approach requires two different FFTs to be performed independently across the bands. The second approach is to use one aggregate band with twice the bandwidth, $2BW$, but use the same SF and chirp BW as before and alias down to $-BW$ whenever the chirp frequency hits the maximum as shown in Fig. 4.5. To demodulate this signal, we just need to multiply the signal which is composed of the aggregate band by the downchirp and perform 2×2^{SF} FFT operation once. The complexity of this method is lower than the former since there is no need to use filters and separate the bands.

4.3.2 Addressing Practical Issues

Timing Mismatch

The above design requires all the backscatter devices to be time synchronized. To understand why, consider two consecutive upchirps being sent by a device, as shown in Fig. 4.6. Now say that we demodulate the signal in these two timing durations, shown in blue and red, we will get different FFT peak locations. Specifically, with a Δt time difference between these durations, the corresponding FFT bin peak location would change by, $\Delta FFT_{bin} = \Delta t BW$. When this change is greater than a single FFT bin, backscatter devices that are assigned to consecutive cyclic shifts interfere with each other and cannot be decoded. Thus, all the devices should be time synchronized. In our design the access point sends a query message telling devices to transmit concurrently. The devices use this query to synchronize and respond concurrently. First, we explain the sources of time delay in our system and then we explain our solutions. There are multiple factors that can contribute to time delays introduced in practice and can be different for different backscatter devices.

- **Hardware delay.** Unlike Wi-Fi devices which use much higher clock frequencies for processors, backscatter devices use low-power microcontrollers (MCUs) that can introduce a variable delay into the system. For backscatter devices, the source of these hardware delay variations come from the time the envelope detector receives the query message from the access point, communicates it to the MCU and then the device backscatters the chirp signal. As we show in §4.4.2, this

hardware delay variations can be as high as $3.5 \mu\text{s}$, which can translate to more than one FFT bin at 500 kHz bandwidth.

- **Propagation delay and multipath.** Since backscatter devices can be at different distances to the access point, their time of flight (TOF) can be different. However, since our target application is for whole-home or warehouse sensing, the propagation distance is less than 100 m which translates to a $ToF < 666\text{ns} = \frac{2 \times 100}{3 \times 10^8}$ and corresponds to only a 0.33 FFT bin change, assuming a bandwidth of 500 kHz. The multipath delay spread for indoor environments is between 50 to 300 ns [166, 94]. For 500 kHz, this delay spread translates to less than 0.15 FFT bin change, which is negligible.

Our solution: Bandwidth-based cyclic-shift assignment. Hardware delay variations over time are hard to correct for. As described above, by nature of operating on MCUs and other low-power computational platforms, these devices have a hardware delay variation over time that changes between packets. Our solution to this problem is to put a few empty FFT bins adjacent to each FFT bin assigned to a device. That is, if FFT bin i is assigned to a device, the adjacent $SKIP - 1$ FFT bins are empty and not assigned to any device. This can be done by using only every $SKIP^{th}$ cyclic shift of the chirp. This ensures that the hardware delay does not result in interference between adjacent devices.

Achieving such an assignment requires us to answer the following key question: how do we pick the value $SKIP$? As described earlier, given the hardware delay variation Δt , the shift in the number of FFT bins is $\Delta t BW$. This means that there is a trade-off in our system regarding the total network throughput, bitrate for each device and sensitivity. In particular, increasing BW increases the number of FFT bins that have to be left empty and decreases the total network throughput. On the other hand, decreasing BW reduces the number of FFT bins but decreases the bitrate per device with the same SF . To compensate for the decreased device's bitrate, we can decrease the SF . Note that, we can choose total bandwidth, chirp BW and SF of the system by considering the hardware delay variations, required bitrate per device, sensitivity for each device and total number of devices. For our implementation, we pick the same total bandwidth and chirp BW of 500 kHz and $SF = 9$ which supports around 1 kbps (976 bps) bitrate at each device while ensuring that the

Table 4.1: **NetScatter Different Modulation Configurations**, with maximum time/freq. mismatch that can be tolerated.

BW [kHz]	SF	Time Variation	Frequency Variation	Bit Rate [bps]	Sensitivity [dBm]
500	9	2 μs	976 Hz	976	-123
500	8	2 μs	1953 Hz	1953	-120
250	8	4 μs	976 Hz	976	-123
250	7	4 μs	1953 Hz	1953	-120
125	7	8 μs	976 Hz	976	-123
125	6	8 μs	1953 Hz	1953	-118

number of empty bins between devices, *SKIP*, is two.

Frequency Mismatch

The devices experience frequency offsets because of hardware variations in the crystals used in their oscillators. As explained in §4.2.1, change in frequency translates to FFT bin change of the demodulated device packet. This again, causes one device to be misinterpreted as other device. Considering a bandwidth of BW and spreading factor of SF , the frequency difference between FFT bins is equal to $\frac{BW}{2^{SF}}$. This means that a Δf frequency offset results in a change in the FFT bin of $\Delta FFT_{bin} = \frac{2^{SF} \Delta f}{BW}$. Therefore, either increasing the spreading factor SF or decreasing the BW can increase the shift in the FFT bin. Crystals' frequency tolerance can be as high as 100 *ppm* [37]. Since backscatter devices run at a few MHz frequencies, this frequency variation translates to less than one FFT bin for the bandwidths and spreading factors in this paper which makes it negligible for our backscatter network.

Table 4.1 shows the timing and frequency mismatch that can be tolerated for different modulation configurations. As can be seen, there are multiple options for achieving the same bitrate and sensitivity. These options will result in different tolerable timing and frequency mismatch, requiring a different *SKIP* value; this is validated using experiments in §4.4.2.




Figure 4.7: **Power Adjustment for Backscatter.** (a) gain normalized to maximum power as a function of Z_0 impedance and (b) switch network to support multiple power levels.

Near-Far Problem

Since our network are designed to work in below-noise conditions, we need to address the near-far problem in our decoding process at the receiver. Specifically, to account for the residual timing and frequency offsets, a CSS receiver has to achieve a sub-FFT bin resolution. To do so without increasing the sampling rate, the receiver uses zero-padding which adds zeros at the end of the time domain samples of the single chirp [96]. Zero-padding operation in the time domain is effectively a multiplication operation with a pulse which translates to convolution with a sinc function in the FFT domain. This makes it easier to locate the FFT peak location. However, convolving with a sinc function introduces side lobes as shown in Fig. 4.8. Assume that there are two devices with cyclic shifts $C_1 = 0$ and C_2 . If the power of C_2 is lower than power of C_1 's side lobes, it cannot be decoded.

Our solution. To address this issue, we propose two techniques that work together to increase our dynamic range.

Coarse-grained power-aware cyclic shift assignment. Our intuition here is as follows: Fig. 4.8 suggests that we should assign adjacent FFT bins to devices that have a small SNR difference. In particular, when $SKIP$ is 2, for two neighboring backscatter devices with an SNR difference greater than 13 dB, the lower power device cannot be decoded. Further, it shows that the side-lobe




Figure 4.8: **Normalized Power Spectrum.** We show power spectrum of an upchirp multiplied by a baseline downchirp in FFT domain. This plot shows the main lobe and side lobes of a single chirp transmission. We assign devices to high and low power regions based on their power level.

power of a high SNR device decreases as we go to farther FFT bins. Thus, we need to ensure that a lower SNR device has to correspond to FFT bins that are farther from the FFT bins corresponding to higher SNR devices. This ensures that the side-lobes of the high-SNR device do not affect the decoding of the low-SNR devices. Specifically, we assign different cyclic shifts to different devices at association phase to ensure that the FFT bins corresponding to the lower-SNR devices are close to each other and are far from higher-SNR devices. To do this, the AP computes the signal strength of the incoming device in the association phase (see §4.3.3) and assigns its cyclic shift based on its signal strength and also the strengths of the devices that are already in the network.

We run simulations to understand the benefits of this allocation. Specifically, we assign two devices to FFT bins 2 and 258, with $SF = 9$ and $BW = 500 \text{ kHz}$. To be realistic, we added Gaussian frequency mismatch with variance of 300 Hz to each device to account for timing and frequency mismatches between them. We change the power of the second device and measure the bit error rate (BER) for the first device. Fig. 4.12 shows the BER over 10^4 symbols, for different power differences between the two devices. As can be seen, the BER remains unaffected even when the second device is around 40 dB stronger than the first device. This shows that our power-aware allocation can in theory tolerate power difference of 40 dB between devices. In practice however this is a little lower at 35 dB (see §4.4.3).




Figure 4.9: **Backscatter Devices SNR Variance.** CDF of SNR variance of backscatter devices in an office environment, when people were walking around, over 30 mins.

Fine-grained self-aware power-adjustment. While the above assignment is determined at association, mobility in the environment and fading will change the SNR of each of the devices over time (see Fig. 4.9). To address this, each device adjusts its power over time using the signal strength of the query message from the AP, using three different levels. We define the maximum power of the device as 0 dB power gain. First, during association, we consider two cases for the associating device. If it sees a low received signal strength for the AP’s query packet, it sets its power gain to the maximum. Otherwise, it sets its gain to the middle level. This gives the higher signal strength backscatter devices leeway to both increase and decrease their power, after association. The AP uses the resulting backscatter signal strengths during association to assign a corresponding cyclic shift. The backscatter devices use the signal strength at association as a baseline and either increase or decrease their power gains for the rest of the concurrent transmissions, i.e., if the signal strength for the AP’s query message increases (decreases), the backscatter devices decrease (increase) their power gain. If the device cannot meet its expected SNR requirements given its limited power levels and assigned cyclic shift, it does not join the concurrent transmissions. If this happens more than twice, the backscatter device re-initiates association after which the AP reassigns the cyclic shifts to account for the new significantly different power value (see §4.3.3).

The key question however is: how can a low-power backscatter device change its transmission power gain? This is interesting since power adaptation has not been used before in the network of backscatter devices. In backscatter, the transmit power gain, $Gain_{power}$, is equal to $\frac{|\Gamma_0 - \Gamma_1|^2}{4}$. Here Γ_0 and Γ_1 are reflection coefficients for switching between two impedance value, Z_0 and

Z_1 . Backscatter hardware is designed to maximize the difference between reflection coefficients to maximize their transmission power. This corresponds to $Gain_{power} = 0 \text{ dB}$. One way to achieve this is to switch between extreme impedance values, $Z_0 = 0\Omega$ and $Z_1 = \infty\Omega$. To achieve power adaptation, in contrast, we pick impedance values that correspond to multiple power settings. In particular, as shown in Fig. 4.7(a), instead of switching from $Z_0 = 0\Omega$, we switch from intermediary impedances and hence achieve lower power gains. Our hardware implementation achieves three power gains of 0 dB, -4 dB and -10 dB to achieve power adaptation. Note that [178] uses a similar circuit structure as Fig. 4.7(b) to cancel higher order harmonics. We instead design this circuit structure to control the power.

Design trade-off. Readers might wonder if reducing the power of high SNR devices would decrease the network throughput, since high SNR devices in traditional LoRa backscatter designs can achieve a higher bitrate. In contrast, by reducing their power we are enabling a large number of concurrent transmissions with a fixed bitrate. Thus, we are encouraging concurrency by reducing the bitrate of high SNR devices. §4.4.4 compares the results for NetScatter with one where each backscatter device uses rate adaptation to pick its ideal bitrate, while transmitting alone using LoRa backscatter [178]. The results show that the network throughput and latency gains due to large scale concurrency outweigh the reduction in the power for high SNR devices.

4.3.3 NetScatter Protocol & Receiver Details

Putting it together, the AP transmits an ASK modulated query message which is used to synchronize all the participating concurrent devices. This message conveys information about cyclic shift assignment which are based on the devices' signal strength at the AP. The devices measure the query message's signal strength using the envelope detector and use it to fine-tune their transmit power gain. In the rest of this section, we describe various protocol details required to make our design work in practice. Note that our focus in the protocol design is about scheduling a set of concurrent transmissions. Typically networks could have more devices than concurrent transmitters supported by our design. Since the AP knows the duty-cycle of each device from the association phase (see §4.3.3), it can i) assign the cyclic shifts and ii) schedule the devices involved in

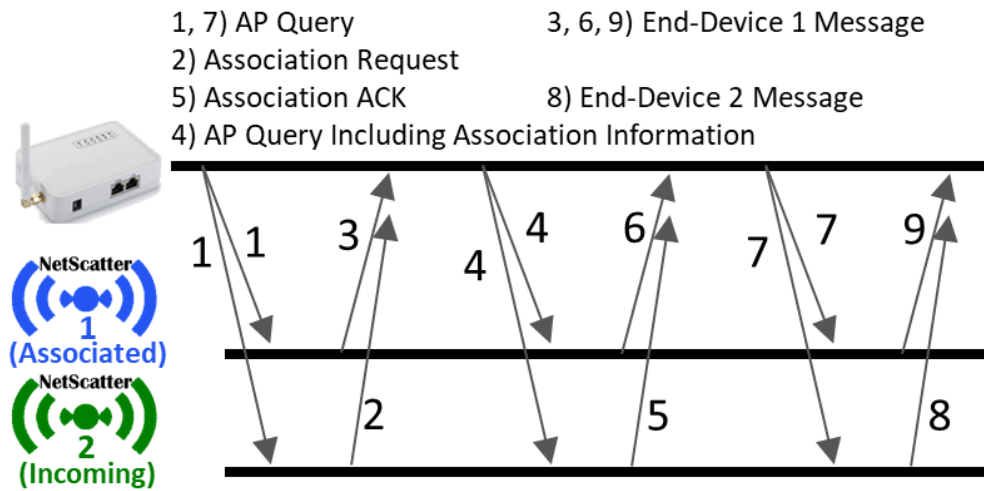


Figure 4.10: **NetScatter Network Association Process.** We show the association process of an incoming NetScatter device (#2) to the network, while there are existing devices associated with the network (i.e., device #1).

concurrent transmissions.

Link-Layer Backscatter Packet Structure

Similar to LoRa, the device packet starts with upchirp and downchirp preambles. They are designed to serve two purposes: i) finding the start of the packet and ii) detecting the transmissions. We emphasize here that the device transmits the same assigned cyclic shift for both upchirps and downchirps in the preamble as well as the payload. The preamble consists of six upchirps followed by *two* downchirps. This is then followed by the payload and the checksum. We note that in our design, all the devices send their preambles concurrently. This reduces the overhead of transmitting preambles for each device, which in turn increase the end-to-end throughput gain achieved by NetScatter. The AP uses the above structure to achieve two goals.

i) Finding the exact packet start. We use the downchirp in the preamble to find the start of the packet transmission. Specifically, we use the middle point between an upchirp and downchirp and switch by six upchirp symbols, number of preamble symbols in our implementation, to the left to find the packet beginning. We suspect that the LoRa preamble has a downchirp for this exact

purpose. We note that in our case, since the upchirp and downchirp in the preamble from each of the devices uses the same cyclic shifts, they are symmetric around the middle point and hence the same algorithm for estimating the packet beginning is applied.

ii) Detecting and decoding each concurrent transmitter. Now that we found the packet start, we need to find out which transmitters are in the network. To do so, for each preamble symbol, we demodulate it and look at the peaks in FFT domain. If there is an FFT peak in the demodulator output which repeats in all the preamble symbols, we conclude that the device corresponding to that cyclic shift is sending data. After finding current devices in the network, we compute the average power over the six preamble symbols for each device. This average power is used as a threshold to demodulate the payload of each device. In particular, if the power of the device's FFT peak for each payload symbol is more than half this average, we interpret that as 1 and 0 otherwise.

Network Association

Say the network already has N devices associated to the AP and the $N + 1^{th}$ device wants to join the network. A naïve approach is to periodically dedicate time periods for association. This however can lead to high association delays depending on the frequency of the association periods. Our approach instead is to reserve N_{assoc} cyclic shifts and the corresponding FFT bins for association and use the rest for communication. In other words, all the devices transmit at the same time but the ones who want to enter the network transmit with the N_{assoc} association cyclic shifts.

To address the near-far problem, we reserve two cyclic shifts, one in high-SNR and the other one in the low-SNR cyclic shift regions. The incoming device would choose which association region to transmit based on the signal strength of the AP's query message, calculated using the envelope detector. However to account for the hardware delay variations, as before, we skip two cyclic shifts to ensure that the association packets from the devices can be decoded and will not interfere with communication cyclic shifts. Finally, to support scenarios where more than one device want to associate at the same time, one can use Aloha protocol with binary exponential back-off in the association process. Our deployment does not implement this option and turns ON the backscatter devices one at a time and runs the network only after all the devices are associated.




Figure 4.11: Structure of AP's Query Message.

After the incoming device sends its packet to the AP in association process using the association cyclic shifts, the AP computes its signal strength and decides which cyclic shift and timing schedule it should be assigned to. The AP piggybacks these assignments in its query messages.

AP Query Message

Fig. 4.11 shows the ASK-modulated query message that the AP sends. The message has a group ID which identifies the set of 256 devices that should concurrently transmit. In our implementation, since there are only 256 devices, we set this group ID to 0. In a larger network, the AP can assign different sets of devices to different groups depending on their signal strengths, i.e., devices that have a similar signal strength are grouped into the same group to enable concurrent transmissions while further minimizing the near-far problem.

This is then followed by an optional association response payload that assigns an 8-bit network ID and a 8-bit cyclic shift. Note that prior LoRa backscatter designs are request-response systems that query each backscatter device sequentially and need most of the fields in Fig. 4.11 other than the group ID and cyclic shift assignment. Since these additional 12 bits is transmitted using 160 kbps ASK downlink, the overhead is negligible compared to the 1 kbps backscatter uplink. Finally, we note that if the AP is unable to assign a new device given the existing assignments, the AP updates the cyclic shift assignments for all the devices in the network. It does so by transmitting the identifier for one of the the 256! orderings, which requires $\log_2(256!) (\leq 1700)$ bits. This occupies less than 11 ms using our 160 kbps downlink.

Network Protocol

Fig. 4.10 summarizes our network protocol. First the AP broadcasts its query. device 1, which is




Figure 4.12: **Near-Far BER Results.** We show the effect of the second device’s power on the first device’s BER vs. SNR for different ratios of the second device’s to first device’s power with power aware cyclic shift assignments.

already associated to the network receives the query and sends its data using its assigned cyclic shift after performing any necessary power control. Concurrently, device 2 sends a *Association Request* using one of the N_{assoc} cyclic shifts. The AP receives these two messages and broadcast another query which includes association information for device 2. Upon receiving this query, device 1 continues to send its data, however, device 2 extract cyclic shift assignment from the query and then transmits *Association ACK* to the AP in the assigned cyclic shift. If AP receives *Association ACK*, it adds device 2 to associated devices. Otherwise, it will repeat the association information in the following queries. After association, each device uses its assigned cyclic shift for sending data.

4.4 Evaluation

4.4.1 Hardware Implementation

Implementation Using Discrete Hardware Components. Our discrete hardware implementation shown in Fig. 4.13 consists of RF section and baseband section, both implemented on a four layers FR4 PCB. On RF receive side, we implemented envelope detector similar to [120] but at 900 MHz and it has a sensitivity of -49 dBm to receive downlink query messages from AP.¹ RF transmit side

¹Note that since ASK-modulated AP query received by backscatter device experiences one-way path loss, its required sensitivity is only -44 dBm in contrast to the -120 dBm sensitivity for the backscatter signals.




Figure 4.13: **Our Backscatter Devices.** They are arranged closely for this picture. They are spread out across more than ten rooms in our deployment.

consists of five ADG904 [42] switches cascaded in three levels to build an impedance switch network for backscatter, power gain control and also switching between transmit and receive modes. Our backscatter device uses a 2 dBi whip antenna to transmit packets and receive query messages in the 900 MHz ISM band.

The baseband side is implemented using an IGLOO nano AGLN250 FPGA [34] and an MSP430FR5969 [62]. We generate CSS packets on the FPGA and output real and imaginary components of the square wave signal to the backscatter switch network. The envelope detector is controlled by the MCU. Downlink receiver algorithm is implemented on MCU. To be resilient to self-interference caused by the AP's single-tone, the baseband at the backscatter device shifts the AP's signal by 3 MHz. Note that the discrete implementation is for prototyping and proof-of-concept; an ASIC is typically required to achieve the orders of magnitude power benefits of backscatter communication. We use a battery to power each backscatter device for our evaluations.

IC Simulation. We design and simulate an IC for our backscatter device using TSMC 65nm LP process. It consists of four blocks with total power consumption of $45.2 \mu W$: i) An envelope detector that demodulates the APs ASK query messages and consumes less than $1 \mu W$. ii) Baseband processor for processing and extracting AP data from envelope detector, interfacing with sensors and sending the chirp specifications and sequence of data to chirp generator consuming $5.7 \mu W$ of




Figure 4.14: **Frequency Offset FFT Bin Variation.** (a) frequency offset of backscatter devices, and (b) effect of residual time and frequency offset for different configurations.

power. iii) A chirp generator that takes SF, BW, cyclic shift assignment and data sequence from the baseband processor to generate the sequence of ON-OFF keying chirps. We use Verilog code to describe the baseband signal's phase behavior and generate assigned cyclic-shift with required frequency offset. We use Synthesis, Auto-Place and Route (SAPR) to simulate Verilog code on chip. The power consumption of this block is $36 \mu\text{W}$. iv) We simulate a Switch network including three resistors that are connected to NMOS switches to generate backscatter signal with three power gain levels. Note that since these resistors and NMOS switches consume minimal area, more power gain levels can be added at almost no cost. The power consumption of the switch network is $2.5 \mu\text{W}$ with 3 MHz frequency offset.

Access Point Implementation. We implement the access point on the X-300 USRP software-defined radio platform by Ettus Research [76]. We use a mono-static radar configuration with two co-located antennas separated by 3 feet. The transmit antenna is connected to a UBX-40 daughterboard, which transmits the query message and the single-tone signal. The USRP output power is set at 0 dBm and we use an RF5110 RF power amplifier [74] to amplify the transmit signal to 30 dBm. The receiver antenna is connected to another UBX-40 daughterboard, which down-converts the NetScatter packets to baseband signal and samples them at 4 Msps.

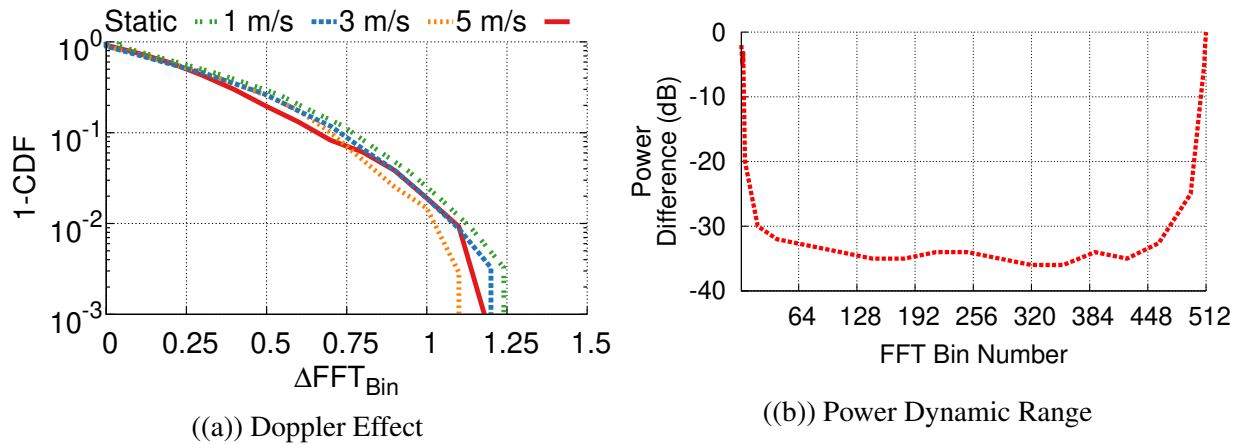


Figure 4.15: **Doppler Effect and Power Dynamic Range Evaluation.** We evaluate (a) Doppler effect, and (b) we show power difference between two concurrent transmissions at different locations of FFT domain. One transmission is fixed and the other is sweeping across different chirp symbols.

4.4.2 Frequency and Timing Mismatch

Measurements 1: Hardware frequency variations. We measure the frequency offsets of our hardware by recording thousand packets for each device. Using the method described in §4.3.3, we compute the frequency offset for the 256 backscatter devices in our network deployment which we show in Fig. 4.14(a). The variations of backscatter devices are less than 150 Hz which is nearly 0.15th of one FFT bin when $BW = 500kHz$ and $SF = 9$. Therefore, our system is not affected by frequency variation of different devices.

Measurements 2: Timing offsets. Next, we characterize how the timing offsets affect ΔFFT_{bin} . This helps us understand how many empty cyclic shifts, $SKIP - 1$, we need to put for each occupied cyclic shift. To do this, we setup a wireless experiment sending query messages from the AP and receiving transmissions from the backscatter devices deployed in our system. By decoding these transmissions and comparing the received cyclic shifts with what we have programmed the devices to send, we can find the ΔFFT_{bin} for each device; this measurement is a combination of both timing and the small frequency variations on the hardware.

Fig. 4.14(b) shows residual ΔFFT_{bin} for backscatter devices. The plots show that the ΔFFT_{bin}

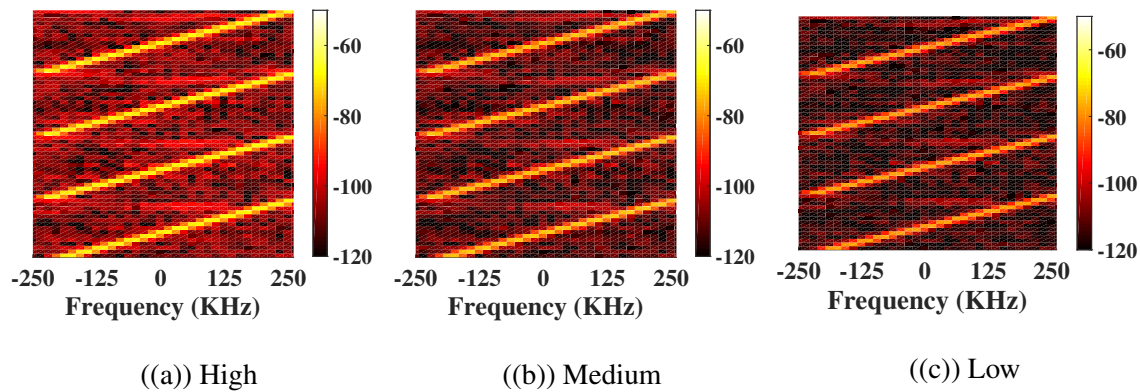


Figure 4.16: Spectrogram of Backscattered Signal at the Different Power Levels.

is considerable. This is because in backscatter devices, the energy detector receives the amplitude modulated query message and sends interrupt to initiate backscatter transmission. Both these steps add to the timing variations. Specifically, the hardware delay variation comes from variation in receiving query message and initiating the transmission on FPGA which can vary from packet to packet. In our deployment in §4.4.4 with backscatter devices, we use $BW=500$ kHz, $SF=9$ and leave one FFT bin between occupied cyclic shifts ($SKIP = 2$). This translates to supporting 256 devices with an aggregate throughput of around 250 kbps and bitrate per tag of around 1 kbps.

Measurements 3: Doppler effects. Other than hardware frequency offsets, Doppler effect can cause changes in frequency as well. However, this effect will be much less than 1 FFT bin, $\frac{BW}{2SF}$, for most cases. As an example, assume a backscatter device is moving with a speed of 10 m/s. Considering the carrier frequency is 900 MHz, the Doppler effect induced frequency change would be 30 Hz which is much less than 1 kHz, the FFT bin frequency, assuming $BW=500$ kHz and $SF=9$. To confirm this, we run various mobility experiments where a subject holds a backscatter device and moves with different average speeds which we measure using an accelerometer. We receive transmissions from the device and compute the ΔFFT_{bin} for different motion scenarios. Fig. 4.15(a) shows ΔFFT_{bin} for various speeds, which confirms that these speeds do not have an effect on ΔFFT_{bin} .

4.4.3 Near-Far Problem

Measurements 1: Power-aware cyclic shift assignment. As mentioned in §4.3.2, we assign cyclic shifts to devices depending on their signal strength values. To evaluate the effectiveness of this technique, we run experiments with two devices where one of them transmits at a high power (equivalent to being near the AP) with a cyclic shift corresponding to the beginning of the FFT spectrum. Then, we sweep the cyclic shift of the second device from small FFT bin difference cyclic shifts to high FFT bin difference ones. At each cyclic shift, we decrease the power of the second device using an attenuator up to when it has packet error rates less than one percent. Fig. 4.15(b) shows the maximum power difference that can be tolerated between these two devices versus the assigned FFT bin difference. As can be seen, as we go further in FFT bin difference, we can tolerate more power difference between the two devices. Note that, because of aliasing Fig. 4.15(b) is symmetric around the center. The maximum happens in middle and is equal to 35 dB. This is the dynamic range that our system can support in practice. We also note that when the second device is assigned to an FFT bin 2 cyclic shifts away from the first device, it can be up to 5 dB below the latter and still correctly decoded. This means there is an in-built 5 dB dynamic range resilience to channel variations between devices that have close cyclic shifts.

Measurements 2: Self-aware power-adjustment. The second method to address the near-far problem and also increase the dynamic-range is power adjustments at the devices using the signal strength of the AP's query message. To evaluate this, we first measure how well we can adjust power on the devices. We evaluate its efficacy in practical deployments. We use three different backscatter impedance values to be able to transmit packets in three different power gains. Fig. 4.16 shows the spectrum of backscattered signal at different power levels. These plots show that the hardware creates spectrum that is clean and does not introduce noticeable non-linearities into the backscattered signal. Furthermore, we can achieve three different power levels: 0, -4, and -10 dB.




Figure 4.17: **Network Physical Rate.** We evaluate NetScatter network physical rate and compare it with other schemes.

4.4.4 Network Deployment

We evaluate three key network parameters:

- **Network PHY bitrate.** This is the bitrate achieved across all the devices during the payload part of the packet.
- **Link-layer data rate.** This is the data rate achieved in the network which is defined as the data rate for sending useful payload bits, after considering overheads including the AP’s query message and the preamble of the packet transmission.
- **Network latency.** This is the latency to get the payload bits from all the backscatter devices in the network.

We compare three schemes: i) LoRa backscatter [178] where all devices use a fixed bitrate of 8.7 kbps, ii) LoRa backscatter with rate adaptation where each device uses the best bitrate given its channel conditions and iii) NetScatter. Note that the authors of [178] did not publicly release the code and so, we replicate the implementation adding the missing details and using $BW = 500 \text{ kHz}$ and $SF = 9$. We also note that [178] is not designed to work with more than one to two users. Here, we use query-response design with scheduling when there are more users where the AP queries

each device. While LoRa backscatter does not support rate adaptation, we want to compare with an ideal approach that maximizes the bitrate of each device by picking the optimal SF and BW . To do so, we measure the signal strength from each of the backscatter devices and compute the bitrate using the SNR table in [51]; this is the ideal performance a single-user LoRa backscatter design achieves with rate adaptation.

Network PHY bitrate. We set each device bitrate to 976 bps, $BW_{agg} = 500 \text{ kHz}$, $SF = 9$ and a payload size of five bytes. We deploy 256 backscatter devices across the floor of an office building with more than ten rooms. Fig. 4.1 shows our deployment in an office. We hard-code cyclic shift assignment on each device. Therefore, we skip the association phase in this deployment. Fig. 4.17 shows the results of network physical rate for our backscatter network deployment. The plot highlights the following key observations.

- The network data rate scales with the number of concurrent backscatter devices. When the number of concurrent devices is less than 128, the variance in the throughput is small. This is because in these scenarios effectively the backscatter devices are separated from each other by more than 2 cyclic shifts ($SKIP \geq 3$). As a result, the devices do not interfere with each other and hence can concurrently operate. As we increase the concurrent devices to 256, we are pushing the system to its theoretical limit (with $SKIP = 2$) and thus, we see larger variances in the network data rate.
- With 256 backscatter devices, NetScatter increases the PHY bitrate by 6.8x and 26.2x over LoRa backscatter with and without rate adaptation. The gains are lower with the ideal rate adaptation since with rate adaptation high-SNR devices could pick the maximum LoRa bitrate of 32 kbps.

Link-layer data rate. While the above plots measure the data rate improvements for the message payload, it does not account for the end-to-end overheads including preambles and the AP's query message to coordinate the concurrent transmissions. To see the effect of the AP query packet overhead for NetScatter, we consider two configurations.

- *NetScatter Config#1.* In this scenario the cyclic shifts are all assigned during the association phase and the AP query packet coordinating the concurrent transmissions is 32 bits long without




Figure 4.18: **Link-layer Data Rate.** We evaluate link-layer data rate for NetScatter and compare it with other schemes.

the optional fields in Fig. 4.11.

- *NetScatter Config#2.* In this scenario, the AP query packet contain cyclic shift assignments for all the devices in the network and has a length of 1760 bits.

The above two configurations represent the two extremes of our deployment. We set the backscatter payload and CRC to 40 bits and use the total 8 upchirps and downchirps for preamble. For LoRa backscatter which queries each individual device sequentially, the AP query is 28 bits long.

Fig. 4.18 shows that the gains at the link-layer are higher for NetScatter over LoRa backscatter without and with rate adaptation by 61.9x (50.9x) and 14.1x (11.6x) respectively for config#1 (#2). This is because, in NetScatter, the added overhead of devices' preambles happen once and at the same time for all devices. But the other schemes need to do TDMA which means that sending preamble will not happen concurrently for all devices and these have to be sent individually for each backscatter device since in traditional designs the AP querying each of them sequentially. Further, in LoRa backscatter which queries sequentially, the AP query message is transmitted once for each device in the network versus being transmitted once for all the devices in our design. Finally, since the downlink uses ASK at 160 kbps, the overhead of transmitting 1760 bits in config#2, while reducing the link-layer data rate over config#1, is still low because the backscatter links can only achieve a much lower bitrate.




Figure 4.19: **Network Latency.** We evaluate the latency of NetScatter and compare it with other schemes. We define latency as total time for transmitting all the devices’ data.

Network latency. Finally, Fig. 4.19 shows that NetScatter has a latency reduction of 67.0x (55.1x) and 15.3x (12.6x) over prior LoRa backscatter without and with rate adaptation respectively in network config#1 (#2). This is the key advantage of using concurrent transmissions in low-power backscatter networks. It is noteworthy that since the downlink AP query bitrate is 160 kbps, AP query duration is negligible compared to duration of backscatter devices’ preamble for prior backscatter methods and also for config#1. For config#2, the AP query duration is significantly higher than the config#1. However, the total duration is still dominated by the backscatter payload + CRC and preamble. As a result, AP query is not the dominant factor in link-layer latency.

4.5 Related Work

Recent systems use backscatter with Wi-Fi signals [120, 202], have a receiver sensitivity of only -90 dBm and hence have a limited range and cannot work across rooms unless the RF source is placed close to the backscatter tag [123, 120]. LoRa backscatter [178] can achieve long ranges by generating LoRa-compliant chirp signals at the backscatter device. pLoRa [159] backscatters ambient LoRa signals in the environment in contrast to the single tone used as the RF source in NetScatter as well as [178]. We note that all SemTech LoRa chipsets have the capability in software to transmit single tone signals. All these prior long range systems are evaluated in a network of only 1–2 devices and propose to use time-division to support multiple backscatter devices. In

contrast, our design enables large-scale concurrent transmissions and can achieve much higher link-layer data rates as well as lower latencies. We also note that these long range backscatter systems [178, 159] claim a kilometer range in outdoor scenarios such as open fields. This however requires placing the RF source close to the backscatter devices. In indoor environments where the signal propagates through walls and the RF source is not placed close to the backscatter devices, our network operational range across ten different rooms is consistent with these prior work. Finally, we note that while prior work [178, 159] decodes the backscatter signal on Semtech LoRa chipsets, our distributed CSS protocol is decoded on a software radio. We however note that SemTech LoRa SX1257 [75] chipsets provide I-Q samples and hence our approach could also be implemented on these off-the-shelf chipsets together with a low power FPGA for baseband processing; this however is not in the scope of this paper.

In addition, prior work [89, 188] use FMCW techniques to multiplex sources of FMCW reflectometry. Specifically, [89] uses FMCW to multiplex Fiber Bragg Grating (FBG) sensors at different positions which results in different delays and different beat frequencies corresponding to each sensor’s reflection. In contrast, NetScatter generates chirp signals with different cyclic-shifts to modulate information on each backscatter device at the same time.

Finally, recent work on decoding concurrent transmissions from RFID tags, does not achieve the long range operations and below-noise operations of CSS based systems. Buzz [191], LF-Backscatter [109], and others [110, 154, 115] leverage the differences in the time domain signal transitions and changes in the constellation diagram to decode multiple RFIDs. However, the number of concurrent transmissions in the above designs is limited — the latest in this line of work, Fliptracer [115], can reliably decode up to five concurrent RFID tags. Further, these systems were tested with ranges of 0.5 to 6 feet [191, 109, 115] and in the same room. Finally, receiver sensitivity of even battery-powered backscatter tags for RFID EPC-GEN2 readers is around -85 dBm. So it cannot support the long ranges and whole-home deployments that CSS modulation based backscatter achieves.

4.6 Conclusion

We present a new wireless protocol for backscatter networks that scales to hundreds of concurrent

transmissions. To this end, we introduce, distributed chirp spread spectrum coding, which uses a combination of chirp spread spectrum (CSS) modulation and ON-OFF keying. Further, we address practical issues including near-far problem and timing and frequency synchronization. Finally, we deploy our system in an indoor environment with 256 concurrent devices to demonstrate its throughput and latency performance.

Chapter 5

SUPPORTING APPLICATIONS WITH PROGRAMMABLE RADIO NETWORKS

5.1 Introduction

Recent years have seen development of numerous wireless protocols for Internet of Things (IoT) devices. In addition to longtime standards such as Bluetooth and Zigbee, a number of new protocols including LoRa, Sigfox, NB-IoT and LTE-M have been developed that achieve long ranges of more than a few kilometers. Due to the lack of a de-facto standard, this space remains an active area of research for both industry and academia. The rapid advances in this space however present practical challenges for researchers: each of these protocols requires a dedicated radio chipset to evaluate, and these proprietary solutions often leave little room for protocol modification. The academic community is therefore severely handicapped by the lack of a *flexible* platform, as even a complex multi-radio prototype cannot adapt to evaluate new protocols or even customize existing solutions. The current ecosystem therefore discourages researchers from investigating the important questions that arise when scaling up IoT networks, and more importantly taking a systematic approach to developing new protocols from the ground up.

Ideally, we would like a large scale IoT network testbed with the flexibility to run *any* IoT protocol at the PHY and MAC layers. Further, since many of these IoT testbeds can span hundreds of endpoints across a large campus or even a city, we need the ability to push changes to the PHY and MAC layers, using simple over-the-air software updates. This would allow for performance comparisons on a single testbed to investigate the trade-offs between existing standards as well as showcase the advantages of an entirely new custom protocol. Moreover, to make such a system representative of real-world deployments, individual network nodes should model the constraints of IoT endpoints. Specifically, these devices should have appropriate power controls and options to




Figure 5.1: TinySDR Hardware Platform. It has two antenna ports for running IoT PHY and MAC protocols at 2.4 GHz and 900 MHz. This image is the actual size of the platform on printed paper.

duty cycle transmissions, have an ultra-low power sleep mode and also have interfaces to connect sensors. Finally, the ability to run these endpoints on batteries would also allow for flexibility of deployment in spaces without dedicated power access, or even in mobile scenarios.

Realizing this vision however is challenging with existing software defined radio (SDR) platforms. Specifically, we require an SDR for the flexibility of implementing different PHY protocols; but there is currently no SDR platform that meets the requirements of IoT endpoints (see Table 5.1). Existing SDR systems consume large amounts of power for transmitting data, do not support ultra-low power sleep modes, require wired infrastructure and often a dedicated computer and furthermore, are expensive. More importantly, none of the existing SDR platforms support over-the-air programming to update PHY or MAC protocols. Finally, IoT devices prioritize power consumption and communication range and hence use limited radio bandwidth — LoRa, Sigfox, NB-IoT, LTE-M, Bluetooth and ZigBee use only 500 kHz, 200 Hz, 180 kHz, 1.4 MHz, 2 MHz and 2 MHz respectively. In contrast, existing SDR platforms focus on achieving high performance in terms of bandwidth because *they are tailored to the needs of gateway devices and not for IoT endpoint devices*.

Driven by a need for such a platform in our own research, we design tinySDR as shown in Fig. 5.1, the first SDR platform tailored to the needs of IoT endpoints. TinySDR provides an entirely standalone solution that incorporates a radio front-end, FPGA and microcontroller for custom processing, over-the-air FPGA and microcontroller programming capabilities, a micro SD

Table 5.1: Comparison Between Different SDR Platforms. Costs are based on sale prices for commercial products without a public bill of materials (BOM) and published BOM prices for research prototypes. OTA refers to over-the-air programming capabilities.

Platform	Sleep Power	Standalone	OTA	Cost	Max BW (MHz)	ADC (bits)	Frequency Spectrum (MHz)	Size (cm)
USRP E310 [22, 40]	2820 mW	✓	✗	\$3000	30.72	12	70~6000	6.8×13.3
USRP B200mini [21, 30]	N/A	✗	✗	\$733	30.72	12	70~6000	5×8.3
bladeRF 2.0 [5, 40]	717 mW	✓	✗	\$720	30.72	12	47~6000	6.3×12.7
LimeSDR Mini [10, 11, 59]	N/A	✗	✗	\$159	30.72	12	10~3500	3.1×6.9
Pluto SDR [41]	N/A	✗	✗	\$149	20	12	325~3800	7.9×11.7
μSDR [27, 68, 26]	320 mW	✓	✗	\$150	40	8	2400~2500	7×14.5
GalioT [150, 18]	350 mW	✓	✗	\$60	14.4	8	0.5~1766	2.5×7
TinySDR	0.03 mW	✓	✓	\$55	4	13	389.5~510, 779~1020, 2400~2483	3×5

card interface for storage, ultra-low power sleep modes and highly granular power management options to enable battery-powered operation. It is capable of transmitting and receiving in both the 900 MHz and 2.4 GHz ISM bands, supports 4 MHz of bandwidth which is sufficient for most IoT protocols including Bluetooth, Zigbee, LoRa, Sigfox, NB-IoT and LTE-M, and can achieve the high sensitivities of commercial solutions such as LoRa chips [51]. Additionally it includes

multiple analog and digital I/O options for connecting sensors.

Designing such an SDR platform required addressing multiple systems, architecture, power and engineering challenges:

- **Low-power hardware architecture.** Achieving a small form-factor, low-power SDR requires a minimalist design approach that can satisfy the real-time needs of IoT protocols and ensure flexibility at the PHY and MAC layers. To do this, we exploit recent advances in small, low-power microcontrollers, FPGAs and flash memory to pick the right components for our platform (see §5.3.1). We use a low-power FPGA to run the PHY layer while the microcontroller runs the MAC protocols as well as handles the I/O operations between the FPGA, radio, memory and sensor interfaces (see §5.3.2).
- **Efficient power management.** Achieving highly granular power management needed for battery-powered operation and enabling ultra-low power sleep modes requires shutting down parts of SDR when not in use. This is important for IoT endpoints that perform duty-cycled operations and require an ultra-low power sleep mode to achieve a long battery life. This presents a design tradeoff between the complexity of toggling the power of each hardware component ON and OFF, and the cost of additional circuitry to do so. We address this challenge in §5.3.3 and achieve sleep power as low as $30 \mu\text{W}$.
- **Over-the-air SDR programming.** Enabling a truly scalable system requires the ability to update the PHY and MAC layers on the platform, over-the-air, in a testbed deployment. This however also introduces the challenge of over-the-air FPGA and microcontroller programming as well as communicating these updates robustly to each device in the network while minimizing power consumption and network utilization. We use a dedicated wireless backbone subsystem complete with a MAC protocol and its own flash memory to program both the microcontroller and FPGA. Additionally we leverage compression and low-power decompression algorithms to minimize network downtime during the updates (see §5.3.4)

Fig. 5.2 shows the power consumption of the radio module in tinySDR compared to existing SDR platforms. We evaluate tinySDR's performance by presenting case studies of two common

protocols: LoRa and BLE beacons, and also evaluate tinySDR in a campus-testbed of 20 devices.

- LoRa modulation and demodulation use 4% and 11% of the FPGA resources respectively and achieve a sensitivity of -126 dBm for 3.12 kbps, which is similar to an SX1276 [51] LoRa chip with the same configuration. Further, the FPGA supports real-time modulation and demodulation of all LoRa spreading factors from 6 to 12. A LoRa MAC implementation on our MCU is also compatible with the *The Things Network*.
- TinySDR supports 2.4 GHz BLE beacon transmissions. The full baseband packet generation on the FPGA uses 3% of its resources. The platform can perform frequency hopping with a delay of 220 us and achieves a sensitivity of -94 dBm which is comparable to the commercial BLE chipsets [44].

Finally, we present a case study of how the unique capabilities of tinySDR could be used to answer new research questions. Recent work has explored techniques to enable concurrent transmissions in LoRa networks [96, 106]; however these solutions were prototyped on USRPs and it is unclear if IoT endpoints can decode concurrent transmissions in real-time within their power and resource constraints. We implement a custom decoder on tinySDR to demonstrate for the first time that IoT endpoints *can* receive concurrent transmissions.

Contributions. To summarize, we design the first SDR platform tailored to the needs of IoT endpoint devices. By making careful design and architectural choices, our platform achieves low power, supports IoT protocols at both 900 MHz and 2.4 GHz and has computation resources to do on-board processing. We present a highly granular power management scheme that enables duty-cycled operation and 10,000x lower power sleep modes. We also develop the first over-the-air SDR programming capability to support PHY and MAC updates in a wireless testbed. We characterize and evaluate our platform with case studies of LoRa and BLE beacons. Finally, we present a research exploration of concurrently receiving multiple LoRa transmissions on our SDR platform.

5.2 SDR Requirements for IoT Nodes

To motivate the need for tinySDR and inform our design decisions, we begin by identifying the key requirements for an IoT endpoint. These include 1) operation in the 900 MHz and 2.4 GHz bands, 2) low power operation which requires the ability to transition to ultra-low power sleep mode, 3) standalone operation which requires an on-board control unit to duty cycle the radio, 4) over-the-air programming capabilities for large scale IoT testbeds, 5) low cost per node, and 6) at least 2 MHz bandwidth to support IoT protocols including LoRa, SIGFOX, LTE-M, NB-IoT, ZigBee and Bluetooth. While there are a number of commercially available SDRs such as the USRP, BladeRF, Pluto SDR, and LimeSDR [22, 5, 76, 11, 69] on the market and SDR research prototypes such as WARP, Argos, SORA, SODA, KUAR, Tick, μ SDR, OpenMili, and GalioT [81, 122, 82, 183, 102, 150, 124, 125, 95, 199, 131, 194, 138, 177, 153, 172, 171], all of them are designed as *gateway devices* and do not satisfy many of the above constraints. Here, we analyze the shortcomings of these platforms in the context of these requirements.

- **Low power operation and sleep mode.** Fig. 5.2 compares the power consumption of the radio module *alone* in existing SDR platforms, since each one has different peripherals. We find that most SDR platforms consume 200-300 mW in receive mode, but a lot more power when transmitting. While this may be acceptable for a gateway devices that are more often receiving, typical IoT endpoints do the opposite and are required to transmit data like sensor information. Moreover, real IoT nodes spend a very short time transmitting before transitioning to ultra-low power sleep modes. Although IoT radios often consume tens to hundreds of milliwatts of power, the key to achieving long battery lifetimes is exploiting their microwatt power sleep modes. Table 5.1 shows that none of the other platforms can benefit from duty cycling as they consume more power in sleep mode than tinySDR does when transmitting; tinySDR's microwatt power consumption in sleep mode enables dramatic power savings with duty cycling.

- **Standalone operation and cost.** We observe that some of these platforms do not allow for standalone operation, i.e., they cannot be used in a testbed deployment without an external computer. Among the ones that do, the Embedded USRP and bladeRF cost \$700 or more per unit making




Figure 5.2: **Radio Module Power Consumption for Each Platform.** The TX output power of each radio module is shown on top of it.

large scale deployments expensive. μ SDR allows for standalone operation but only operates at 2.4 GHz and cannot support protocols like LoRa. GalioT [150] uses the low cost RTL2832U radio [18] connected to a Raspberry Pi computer which allows for standalone operation, however it does not support 2.4 GHz band. Moreover, this platform is *receiver only* and cannot be used to prototype a typical IoT node that transmits data.

- **Over-the-air (OTA) programming.** As shown in Table 5.1, all existing SDR platforms rely on wired interfaces for programming. This means that even if one of these systems were connected to a battery, running an experiment would require either tethering each one to a wired network or individually programming them. An OTA programming system is crucial to realizing the goal of a large scale wide area testbed as without it, researchers have to decide between limiting themselves to deployment scenarios with wired infrastructure that are not representative of real IoT use cases or traveling over *kilometer* distances to update individual nodes for *each* minor protocol modification, which would be unmanageable at scale.

5.3 TinySDR Platform

We first describe our design choices for the different components of our hardware shown in Fig. 5.3 and explain the interfaces between them. Next we present the power management module which enables our ultra-low-power sleep mode. Finally, we describe our over-the-air update protocol

including decompression algorithms and over-the-air reprogramming.

5.3.1 *Hardware Design*

We seek to minimize power consumption and cost while offering the flexibility of an SDR to process raw samples.

Designing the Software Radio

The core block on our platform is the software-defined radio, a programmable PHY layer that processes and converts bits to radio signals and vice versa. We begin by explaining our choices for the primary components of an SDR which are a radio chip that provides an interface for sending and receiving raw samples of an RF signal as well as an FPGA that can process these signals in real time. We then discuss the supporting peripherals for these devices such as a power amplifier (PA) to boost the output of the radio chip and non-volatile memory for the FPGA to read and write data from.

Choosing a radio chip. We begin by choosing a radio chip as its specs define the requirements for the FPGA and other blocks. Our primary requirement is that the chip supports reading and writing raw complex I/Q samples of the RF signal. As shown in Table 5.2, current SDR systems use I/Q radio chips that are designed to cover a multi-GHz spectrum and have high ADC/DAC sampling rates to support large bandwidth. For example, the AD936x [40] series which is used in USRP and Pluto SDR can transmit up to 3.8 GHz and supports sampling rates as high as tens of MHz. Each of these specs such as wide bandwidth, low noise, and high sampling rate represent fundamental trade offs of power for performance, and therefore these chips consume watts of power. Moreover, some of these radio chips costs more than \$100.

We instead take a different approach: identify the minimum required specs and find a radio that supports them. Specifically, a radio chip for an IoT platform must be able to operate in at least the 900 MHz and 2.4 GHz ISM bands, have 4 MHz of bandwidth, while otherwise minimizing power and ideally costing less than \$10. We analyze all of the commercially available radio chips

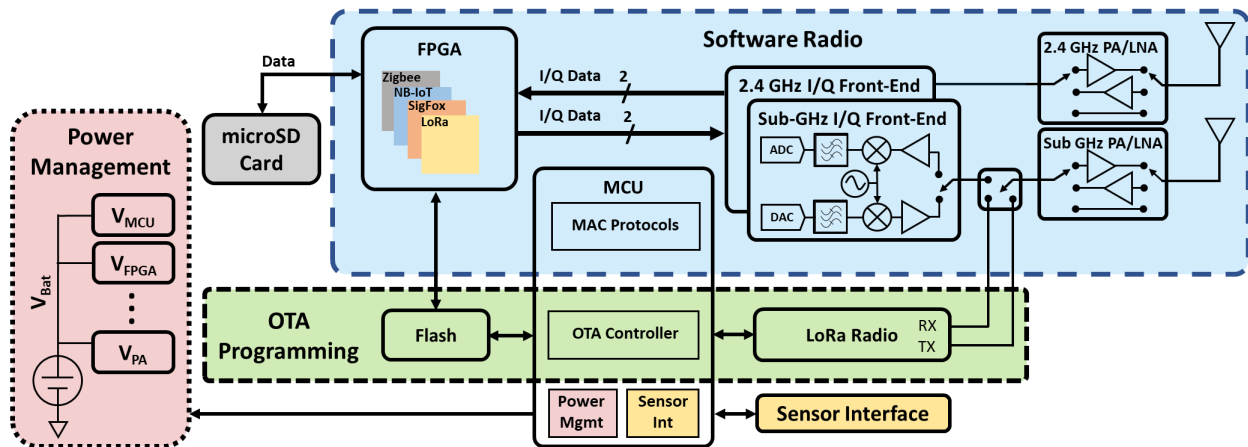


Figure 5.3: **TinySDR System Block Diagram.** A complete system diagram showing all of the components of tinySDR. This includes the software radio consisting of the radio, amplifiers, and FPGA, OTA programmer which uses a LoRa radio and flash memory to store programs, and a power management system with the flexibility to turn off power consuming components. Each of these subsystems are controlled in software running on the MCU.

that provide baseband I/Q samples and list them in Table 5.2, where only the AT86RF215 supports all of our requirements. In addition to its lower cost and support for both frequency bands, it also consumes less power than the MAX2831 and the SX1257. Moreover, the AT86RF215 integrates all the necessary blocks including an LNA, programmable receive gain, automatic gain control (AGC) and low pass filter, ADC on the RX chain, as well as a DAC and programmable PA with a maximum power of 14 dBm on the TX side. In terms of noise, the RF front-end has a 3-5 dB noise figure which is even better than the noise figure of the front-end used in Semtech SX1276 LoRa chipset, suggesting it should be able to achieve long range performance. It consumes 5x less power than the radios used on other SDRs as shown in Fig. 5.2 and has built in support for common modulations such as MR-FSK, MR-OFDM, MR-O-QPSK and O-QPSK that can save FPGA resources or power by bypassing the FPGA entirely.

Picking an FPGA. Now that we have chosen a radio chip, the next step in our design process is to find an FPGA that can interface with it. Aside from minimizing power and cost, we would also like to maintain a small form factor and short wake-up time. Although flash-based FPGAs are capable of fast wake-ups, they are more expensive compared to SRAM-based FPGAs with the

Table 5.2: Existing Off-the-Shelf I/Q Radio Modules.

I/Q Radio	Frequency (MHz)	RX Power (mW)	Cost
AD9361 [40]	70~6000	262	\$282
AD9363 [41]	325~3800	262	\$123
AD9364 [30]	70~6000	262	\$210
LMS7002M [59]	10~3500	378	\$110
MAX2831 [27]	2400~2500	276	\$9
SX1257 [75]	862~1020	54	\$7.5
AT86RF215 [43]	389.5~510 779~1020 2400~2483	50	\$5.5

same number of logic elements. We use LFE5U-25F [72] FPGA from Lattice Semiconductor for baseband processing which is SRAM-based and has 24k logic units. This chip provides a greater number of look up tables (LUTs) than the FPGAs on the Pluto SDR and LimeSDR mini, and at a lower cost. Moreover, it is significantly cheaper than the flash-based FPGA used in uSDR[124].

Adding a power amplifier (PA). AT86RF215 only supports a maximum transmit power of 14 dBm which is traditionally used by IoT radios but is less than the 30 dBm maximum allowed by the FCC. To provide flexibility, we add optional PAs. Given the high cost and power requirements of wide-band PAs that could operate at both 900 MHz and 2.4 GHz we instead select two different chips: the SE2435L [49] for 900 MHz and SKY66112 [65] for 2.4 GHz. Our 900 MHz PA supports up to 30 dBm output power, and the 2.4 GHz PA can output up to 27 dBm. Both chips also include an LNA for receive mode and a built in circuit to bypass either of these components for power savings. In receive mode, we can either pass the incoming signal through the LNA and then connect it to the radio or completely bypass the LNA and connect the signal directly. The maximum bypass current is 280 uA and the sleep current of both power amplifiers is only 1 uA. In transmit operation we can pass the signal through the PA and amplify the signal or turn off the PA and pass the signal directly to the antenna for transmit power < 14 dBm.

Picking the microcontroller. We use a microcontroller to control all the individual chips and

toggle all of these power saving options. In addition to having a low sleep current it must be able to support multiple control interfaces, have enough memory resources to support IoT MAC protocols and also be able to run a decompression algorithm for our OTA system. We select the MSP432P401R [63] a 32-Bit Cortex M4F MCU which meets all of our requirements with less than 1 μ A sleep current, has 64 KB of onboard SRAM and 256 KB of onboard flash memory. In addition to controlling the I/Q and backbone radio parameters, and reprogramming of the FPGA, the MCU performs the important function of power management. It is responsible for toggling ON and OFF the power amplifiers, as well as performing power-gating by turning ON and OFF different voltage regulators in §5.3.3.

Designing OTA Update Hardware

While the above discussion enables a small, low power, low cost SDR for easy deployment, FPGAs and microcontrollers typically require a wired interface for reprogramming. There are two options for enabling wireless reprogramming: i) using the existing I/Q radio and FPGA and ii) using a dedicated wireless communication chipset. We chose the second option since it provides a fail-safe mode for updating the firmware. Moreover, a dedicated wireless communication chipset would consume less power compared to the first option.

OTA wireless chipset. A key question when designing an OTA update system is, what wireless protocol should be used? To support wide area networking, we focus on protocols designed for long range operation. We analyze all of the available long range protocols and select LoRa for our OTA system for a number of reasons. First, LoRa receivers have a high sensitivity which enables kilometer ranges. LoRa also support a wide range of data rates from 11 bps to 37 kbps which allows us to trade off rate for range depending on the deployment scenario. Moreover, LoRa is becoming more and more wide-spread in the US. We use the SX1276 Semtech chipset [51] which is available for \$4.5, minimizing cost.

Flash Memory. Our FPGA is SRAM based and does not include on-chip non-volatile memory for storing programming data. We instead store the firmware bitstream on a separate flash memory chip. The FPGA programming bitstream is 579 KB and the MCU programs require a maximum

I_SYNC (0b10)	I_Data (13 bits)	Control (1 bit)	Q_SYNC (0b01)	Q_Data (13 bits)	Control (1 bit)
------------------	---------------------	--------------------	------------------	---------------------	--------------------

Figure 5.4: I/Q Word Structure Used by I/Q Radio.

of 256 KB. We chose the MX25R6435F flash chip with 8 MB memory. Although this is far more than the size required, it allows tinySDR to store multiple FPGA bitstreams and MCU programs to quickly switch between stored protocols without having to re-send the programming data over the air.

5.3.2 Interfacing Between Blocks

Reading and Writing I/Q Samples

The AT86RF215 radio chipset samples baseband signals at 4 MHz with a 13 bit resolution for both I and Q. Operating at the full rate therefore requires an interface which can support a throughput of over 100 Mbps without consuming a large amount of power to meet our design objectives. To do this we use low-voltage differential signaling (LVDS) [15] which is a high-speed digital interface that reduces power by using lower voltage signals but maintains good SNR by sending data over two differential lines to reduce common mode noise.

Receiving serial I/Q data. Our system communicates over LVDS to the FPGA in serial mode to transfer I/Q data with a physical interface consisting of 4 I/O lines, pairs of which are used to send data and clock signals. The radio outputs 32-bit serial data words at 4 Mwords/s using the format in Fig. 5.4. Each data word starts with the *I_SYNC* pattern which indicates the start of the *I* sample which we use for synchronization. Next, it has 13 bits of *I_Data* followed by a control bit. The same format follows for *Q*, beginning with a synchronization pattern *Q_SYNC* and then 13 bits for *Q_Data* and the final control bit. The required 128 Mbps data rate is achieved using a 64 MHz clock provided by the radio operating at double data rate by sampling at both the rising and falling edges of the clock. We implement an I/Q deserializer on the FPGA to read the data which samples the input at both the rising and falling edges of the clock, uses the *I_SYNC* and *Q_SYNC* to detect the beginning of the data fields and loads the *I* and *Q* values into 13 bit registers

for parallel processing.

Transmitting I/Q samples. In TX mode we need to do the opposite of the above sequence to convert from the parallel representation on the FPGA to a serialized LVDS stream. To do this, we use the FPGA's onboard PLL to generate the 64 MHz clock signal. Next to create our double data rate output signal that varies on both the positive and negative edges of this clock signal using a dual-edge D flip-flop design [108] resulting in the desired 128 Mbps data rate. We use this to generate the same I/Q word structure described above.

Memory Interfaces

After reading the raw data from the LVDS lines using the I/Q deserializer described above, we store the samples into a FIFO buffer implemented using the FPGA's embedded SRAM. We implement a simple memory controller to write data to the FIFO which generates the memory control signals and writes a full data word on each cycle. The embedded memory can run at rates significantly greater than 4 MHz meaning it is not a limiting factor for real-time processing. The SRAM can buffer up to 126 kB. The data stored in the FIFO can then be sent to signal processing blocks to implement filters, cryptographic functions, etc. or to non-volatile flash memory. For flash memory, we use a micro SD card which enables us to collect raw I/Q data and analyze the spectrum. The micro SD card supports two modes: native SD mode and standard SPI mode. In native SD mode, micro SD card's interface uses 4 parallel data lines to read/write data to/from the micro SD card. This mode supports a higher data rate compared to the SPI mode which only supports a 1-bit serial interface. However, we implement SPI mode since it supports the 104 Mbps data rate which we need to write data in real time. This allows us to re-use the same, simpler SPI block for multiple functions and save resources on the FPGA.

RF, Control and Sensor Interfaces

The AT86RF215 provides differential RF signals for both 900 MHz and 2.4 GHz and has an integrated TX/RX switch for both. At 2.4 GHz, the differential signal is transformed to a single-ended

output using the 2450FB15A050E [25] balun and fed to the SKY66112 [65] front-end with the bypassable LNA and PA. Finally, after passing through a matching network, the 2.4 GHz signal is connected to an SMA output.

On the 900 MHz side, the differential output of the AT86RF215 is connected to 0896BM15E0025E [70] to convert it to a single-ended output. This must be shared between the backbone radio's two separate RF paths for transmit and receive and AT86RF215's 900 MHz single-ended signal. We choose between them using a ADG904 [42] SP4T RF switch. The single port side is connected to the SE2435L [49] 900 MHz front-end which is similar to the 2.4 GHz front-end. The MCU communicates with the I/Q radio, backbone radio, FPGA and Flash memory through SPI which it uses to send commands for changing the frequency, selecting the outputs, etc. It also has control signals for FPGA programming, 900 MHz and 2.4 GHz front-end modules, RF switch and voltage regulators for active power control. TinySDR supports common digital interfaces like SPI and I2C to communicate with digital sensors and two analog to digital converters (ADC) for interfacing with analog sensors. We leverage the internal flash memory of the MCU (≈ 256 kB) and external flash memory (≈ 8 MB) to store sensor data.

5.3.3 *Power Management Unit*

Next, we present the design of our power management unit which seeks to maximize the system lifetime when running off of a 3.7 V Lithium battery. To enable long battery lifetimes we need to be able to duty-cycle our system and allow the MCU to toggle each of the above blocks ON and OFF when they are not in use. Further, different components have different supply voltage requirements and we wish to provide each one with the lowest voltage possible to minimize power usage.

Ideally we would want separate controllable voltage regulators for each component in the system. However, having many different regulators with individual controls significantly increases the complexity, number of components, and price. Moreover, it complicates the PCB design by requiring many control signals and a multitude of power planes. Therefore, there exists a trade-off between the granularity of power control and the price/complexity of a design. We outline the sup-

Table 5.3: Power Domains in TinySDR.

Component	Voltage [V]	Power Domain
MCU	1.8V	V1
FPGA	1.1, 1.8, 2.5, V _{lvs}	V2, V3, V4, V5
I/Q Radio	1.8 < V5 < 3.6	V5
Backbone Radio	1.8 < V5 < 3.6	V5
sub-GHz PA	3.5V	V6
2.4 GHz PA	1.8, 3.0	V3, V7
FLASH Memory	1.8	V3
Micro SD Memory	3.0	V7

ply voltages needed for each component and the power domain supporting it in Table 5.3. Below, we show how we group components to balance power and complexity.

- **Power domain V1 (MCU).** Since the MCU is the central controller that implements power management, it needs to be powered at all times and therefore has its own power domain. To minimize its sleep current we need to use a voltage regulator with a low quiescent current. Although switching voltage regulators have higher conversion efficiency when active, they also have high quiescent currents so we instead select the TPS78218 linear regulator.
- **Power domains V2, V3, V4, V6 and V7.** These power domains provide power to blocks such as the FPGA, memory blocks, and PAs. Since these components can all be turned off when not operating, the voltage regulators for these domains should have low shut-down current during sleep and high efficiency when active. We therefore choose the TPS62240 which has a shutdown current of only 0.1 uA. It is highly efficient and is rated to support the current draw required by all components except the 900 MHz PA. To support this PA at its maximum output power we use the TPS62080 switching regulator which supports the required current.
- **Power domain V5.** V5 is a shared power domain for I/Q radio, backbone LoRa radio and FPGA I/O bank. This power domain is initially set to 1.8V to minimize power consumption, however components such as the radio chips can require higher voltage to achieve maximum output power. Therefore, in addition to high efficiency and low shut-down current like the others, this domain




Figure 5.5: LoRa Packet Structure.

should be programmable. To do this, we use Semtech SC195ULTRT [33] which provides an adjustable output that can be set from 1.8 V to 3.6 V.

5.3.4 *Over-the-Air Programming protocol*

OTA AP and MAC protocol. To update a network of tinySDR devices, we use an AP with a LoRa radio to communicate with each device sequentially. In order to propagate updates throughout a testbed or to specific tinySDR nodes, we design a MAC layer for the LoRa PHY. We pre-program a timer on the MCU to periodically turn off the FPGA and switch from IQ radio mode to the backbone radio to listen for new firmware updates. If there is an update, the AP sends a programming request as a LoRa packet with specific device IDs indicating the nodes to be programmed along with the time they should wake up to receive the update. Upon processing this packet and detecting its ID, the tinySDR node switches into update mode and sends a ready message to the AP at the scheduled time. Then, the AP transmits the firmware update as a series of LoRa packets with sequence numbers. Upon receiving each packet, the tinySDR node checks the sequence number and CRC. For a correct packet it writes the data to its flash memory and transmits an ACK to indicate correct reception. In the case of failure no ACK is sent and the AP re-transmits the corrupted packet after a timeout. After sending all the firmware data, the AP sends a final packet indicating the end of firmware update which tells the tinySDR node to reprogram itself and switch back to normal operation.

To maintain the OTA timing, we use a programmable timer that operates with a 20 PPM low-frequency crystal oscillator source. This will result in timing drift between the tinySDR node and the AP over time. However, with each update we compensate the error by sending the correct time to the tinySDR.

Compressing and decompressing the bitstream. Our system compresses data to reduce update

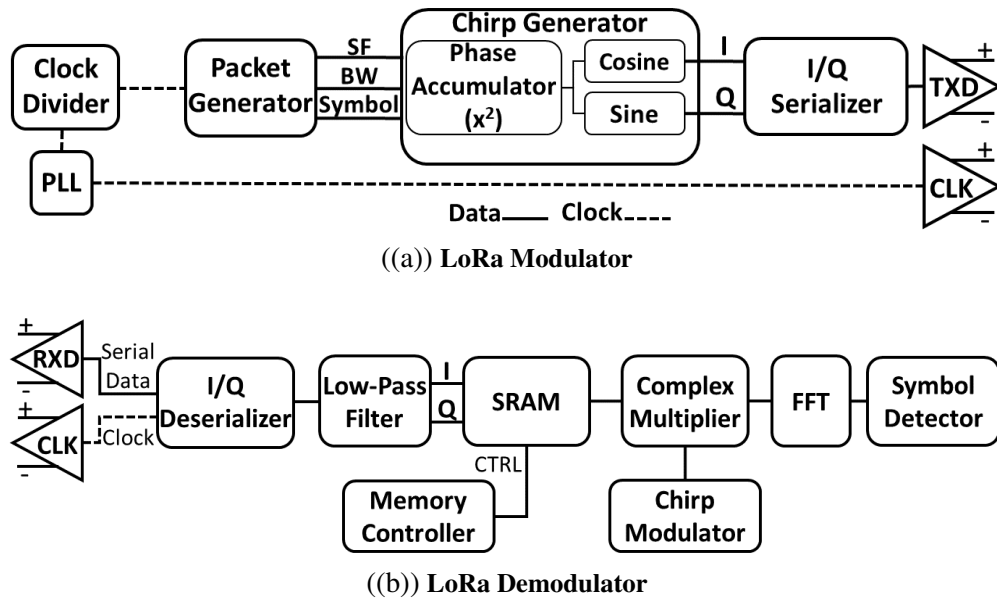


Figure 5.6: LoRa Implementation Block Diagrams.

times, however this compression must be compatible with the resources available on tinySDR. We choose the miniLZO compression algorithm [73], which is a lightweight subset of the Lempel–Ziv–Oberhumer (LZO) algorithm. Our implementation of miniLZO only requires a memory allocation equal to the size of the uncompressed data. We perform compression on the AP. The compression ratio of bitstream file varies based on the content of the bitstream, and in the worst case the compressed file could have almost the same size of the original file. This would require a maximum memory allocation of 579 kB which we cannot afford on a low-cost MCU. Instead, we first divide the original update file into blocks of 30 kB that will fit in the MCU memory. Then we compress each block separately and transmit them to the tinySDR node one by one. Considering the LoRa radio takes more power than the MCU, we immediately write the data to our dedicated programming flash memory using an SPI interface.

After receiving all the data we turn off the LoRa radio and decompress data. First, we allocate memory on the MCU’s SRAM equal to the block size and load a block of data from flash. Next, we perform decompression and write the data in the allocated SRAM memory. Finally, we write the

decompressed data back to the flash beginning at the corresponding address of the programming boot file. We repeat these steps until we decompress the full firmware update.

Over-the-air FPGA programming. After storing uncompressed programming data in flash memory, we program the FPGA. We use the MCU to set the FPGA into programming mode. When the FPGA switches to programming mode, it automatically reads its firmware directly from the flash memory using a 62 MHz quad SPI interface and programs itself. Reading from flash using quad SPI achieves programming times of 22 ms which is similar to FPGAs with embedded flash memory and results in minimal system down time. After programming is complete, it resumes operation and begins running the new firmware.

5.3.5 *TinySDR's Architectural Considerations*

We design tinySDR to achieve three main goals: i) low-power ii) low-cost and iii) over-the-air programmability. To do this, we use a low-power I/Q radio with lower bandwidth support compared to previous platforms. Since this radio is optimized for low bandwidths and in turn low sampling rates, it consumes less power during TX/RX operations. Previous platforms [69, 124] use hardware architectures that support high-bandwidth protocols such as Wi-Fi. However, we use a low-power radio and build our hardware architecture for IoT protocols around it. In addition, we design a power management system to be able to power cycle different parts of tinySDR's architecture in each operation to further reduce the power consumption. To achieve this, we use an MCU chip that enables full control of the tinySDR's blocks and power domains. We achieve minimum power consumption during sleep mode by turning off power-hungry components.

In contrast, previous architectures such as uSDR [124] use an MCU integrated with an FPGA which forces the FPGA to be always ON and increases the power consumption during sleep mode. This is because uSDR is not designed for IoT endpoints but for gateways and hence does not provide any of the architectural optimizations required for the low-power sleep operation required by IoT devices. Furthermore, by using a lower bandwidth radio and also a low-cost SRAM-based FPGA, we minimize the cost compared to platforms such as uSDR [124], Pluto SDR [69] and LimeSDR [10]. Finally, we design tinySDR to operate completely standalone on battery without

the need for a wired connection to a network or a computer. To do this, we design an OTA system on tinySDR to be able to re-program the FPGA and MCU on tinySDR wirelessly. This capability does not exist on prior SDR platforms.

5.4 Case Studies: LoRa and BLE Beacons

5.4.1 LoRa Protocol with tinySDR

We choose LoRa as it is gaining popularity for IoT solutions due to its long range capabilities. Since LoRa is a proprietary standard, we begin by describing the basics of its modulation and packet structure followed by the implementation details of our modulator, demodulator and MAC protocol.

LoRa Protocol Primer. LoRa achieves long ranges by using Chirp Spread Spectrum (CSS) modulation. In CSS, data is modulated using linearly increasing frequency upchirp symbol. Each upchirp symbol has two main features: Spreading Factor (SF) and Bandwidth (BW). SF determines the number of bits in each upchirp symbol [96, 106, 178] and BW is the difference between upper and lower frequency of the chirp which together with SF determines the length of an upchirp symbol. SF and BW trade data rate for range. Data is modulated by 2^{SF} cyclic-shifts of an upchirp symbol. The starting point of the symbol in frequency domain, which is the cyclic shift of the upchirp symbol, determines its value [36]. LoRa uses SF values from 6 to 12 and BW values from 7.8125 KHz to 500 KHz to achieve PHY-layer rates of $\frac{BW}{2^{SF}} \times SF$.

Fig. 5.5 shows the LoRa packet structure which begins with a preamble of 10 zero symbols (upchirps with zero cyclic-shift). This is followed by the Sync field with two upchirp symbols. Next, a sequence of 2.25 downchirp symbols (chirp symbol with linearly decreasing frequency) indicate the beginning of the payload. The payload then consists of a sequence of upchirp symbols which encode a header, payload and CRC.

LoRa Modulator. Fig. 5.6(a) shows the block diagram of our LoRa modulator. We use our FPGA to implement a LoRa modulator in Verilog and stream data to AT86RF215 in I/Q mode. The modulator begins with the *Packet Generator* module which reads data either from FPGA memory

for transmitting fixed packets or from the MCU, as well as LoRa configuration parameters such as SF, coding and BW. This module determines each symbol value and its corresponding cyclic-shift. Next, the *Packet Generator* sends these parameters along with the symbol values to the *Chirp Generator* module, which generates the I/Q samples of each chirp symbol in the packet using a squared phase accumulator and two lookup tables for *sine* and *cosine* function [178]. We then feed these I/Q samples into an I/Q Serializer to stream them over the LVDS interface to the I/Q radio. We generate the 64 MHz transmission clock using internal PLL of the FPGA.

LoRa Demodulator. Fig. 5.6(b) shows the block diagram of our LoRa demodulator. It begins by reading data from the I/Q radio into the *I/Q Deserializer* module on the FPGA which converts the serial I/Q stream to parallel I/Q for further signal processing. Next, we run the data through a 14 tap FIR low-pass filter to suppress high frequency noise and interference. We store the filtered samples in a buffer implemented using the FPGA's memory blocks. To decode the data, we use the *Chirp Generator* module from the *LoRa Modulator* described above to generate a baseline upchirp/downchirp symbol, and then we multiply that with the received chirp symbol using our *Complex Multiplier* unit. The output of the multiplication then goes to an FFT block implemented using a standard IP core from Lattice. Finally the *Symbol Detector* scans the output of the FFT for peaks and records the frequency of the peak to determine the symbol value. To detect the chirp type (upchirp/downchirp), we multiply each chirp symbol with both an upchirp and downchirp and then compare the amplitudes of their FFT peaks. The higher peak in the FFT shows higher correlation which indicates the chirp type.

LoRa MAC Layer. To demonstrate that our LoRa implementation on tinySDR is compatible with existing LoRa networks such as the LoRa Alliance's [60] The Things Network (TTN) [78], we adopt their LoRa MAC design from TTN's Arduino libraries [79] and implement it on tinySDR's MCU. TTN uses two methods for device association; Over-the-air activation (OTAA) and activation by personalization (ABP). In OTAA, each node performs a join-procedure during which a dynamic device address is assigned to a node. However, in ABP we can hard-code the device address in the device which makes it simpler since the node skips the join procedure. Our platform can support both OTAA and ABP methods.




Figure 5.7: Evaluation Testbed Map.




Figure 5.8: TinySDR Single-Tone Frequency Spectrum.

5.4.2 BLE Beacons with tinySDR

To demonstrate tinySDR's 2.4 GHz capabilities we implement Bluetooth beacons which are commonly used by IoT devices.

BLE Beacon Primer. We implement non-connectable BLE advertisements (`ADV_NON_CONN_IND`) which are broadcast packets used for beacons. These packets allow a low power device to broadcast its data to any listening receiver within range without the power overhead of exchanging packets to setup a connection. These packets have a bit rate of 1 Mbps in Bluetooth 4.0 or up to 2 Mbps in Bluetooth 5.0 and are generated using GFSK with a modulation index of 0.45-0.55. The GFSK modulation is binary frequency shift keying (BFSK) with the addition of a Gaussian filter to the square wave pulses to reduce the spectral width.

Generating a BLE Packet. Bluetooth advertisements consist of 6-37 octets, beginning with fixed

Table 5.4: Different Operation Timing for TinySDR.

Operation	Duration (ms)
Sleep to Radio Operation	22
Radio Setup	1.2
TX to RX	0.045
RX to TX	0.011
Frequency Switch	0.220

preamble and access address fields indicating the packet type set to 0xAA and 0x8E89BED6 respectively. This is followed by the packet data unit (PDU) beginning with a 2 byte length field and followed by a manufacturer specific advertisement address and data. The final 3 bytes of the packet consist of a CRC generated using a 24-bit linear feedback shift register (LFSR) with the polynomial $x^{24} + x^{10} + x^9 + x^6 + x^4 + x^3 + x + 1$. The LFSR is set to a starting state of 0x555555 and the PDU is input LSB first. The final LFSR state after inputting the PDU becomes the CRC. Data whitening is then performed over the PDU and CRC fields to eliminate long strings of zeros or ones within a packet. This is also done using a 7-bit LFSR with polynomial $x^7 + x^4 + 1$. The LFSR is initialized with the lower 7 bits of the channel number the packet will be transmitted on, and each byte is input LSB first. We implement both these blocks in Verilog on the FPGA.

Packet Transmission and MAC Protocol. From this bitstream, we need to generate the I/Q samples to feed to the I/Q radio. First, we upsample and apply a Gaussian filter to the bitstream. This gives us the desired changes in frequency which we integrate to get the phase. We then feed the phase to *sine* and *cosine* functions to get the final I and Q samples, which are passed to I/Q serializer and sent to the I/Q radio. BLE divides the 2.4 GHz band into channels, each spaced 2 MHz apart, but BLE beacons are only transmitted on three advertising channels without carrier sense, typically in sequential order separated by a few hundred microseconds. This sequence is re-transmitted every advertising interval [77].

5.5 Evaluation

We deploy a testbed of 20 tinySDR devices across our institution's campus as shown in Fig. 5.7. To see if tinySDR meets the requirements for IoT endpoint devices, we characterize its power, computational resource usage, delays and cost when operating in different modes and running different protocols.

5.5.1 Benchmarks and Specifications

Sleep mode power. Many IoT nodes perform short, simple tasks allowing them to be heavily duty cycled which allows them to achieve battery lifetimes of years. We design tinySDR with this critical need in mind such that the MCU can actively toggle on and off power consuming components such as the radio, PAs, and FPGA to enter a low power sleep mode.

We do this by first turning off the the I/Q transceiver and LoRa radios. To reduce the static power consumption of the FPGA, we shut it down by disabling the voltage regulators that provide power to its I/O banks and core voltage. Similarly, we also turn off the PAs. Finally, we put the MCU in sleep mode LPM3 running only a wakeup timer. The measured total system sleep power in this mode was 30 uW.

The low sleep power allows for significant power savings, but also introduces latency. Table 5.4 shows the time required to wake up from sleep mode until the radio is active. Because we can perform the I/Q radio setup in parallel with booting the FPGA, the total wakeup time for RX and TX is 22 ms. The I/Q radio setup takes 1.2 ms, so the wakeup time is dominated by booting up the FPGA which itself takes 22 ms. We compare this to a SmartSense Temperature sensor [32] and find that tinySDR has only a 4x longer wakeup time even though it requires programming unlike commercial products that use a custom single protocol radio. Additionally many IoT devices operate at low duty cycles waiting in sleep mode for seconds or more making tinySDR's wakeup latency insignificant.

Switching delays. We also measure the switching delays for different operations on the I/Q radio as this is an important parameter for meeting MAC and protocol timing requirements. Table 5.4

Table 5.5: TinySDR Cost Breakdown for 1000 Units.

Components		Price
DSP	FPGA	\$8.69
	Oscillator	\$0.9
IQ Front-End	Radio	\$5.08
	Crystal	\$0.53
	2.4 GHz Balun	\$0.36
	Sub-GHz Balun	\$0.3
Backbone	Radio	\$4.5
	Crystal	\$0.4
	Flash Memory	\$1.6
MAC	MCU	\$3.89
	Crystals	\$0.68
RF	Switch	\$3.14
	Sub-GHz PA	\$1.54
	2.4 GHz PA	\$1.72
Power Management	Regulators	\$3.7
Supporting Components	–	\$4.5
Production	Fabrication [48]	\$3
	Assembly [48]	\$10
Total	–	\$54.53

shows that it takes $45 \mu\text{s}$ and $11 \mu\text{s}$ to switch from TX to RX mode and RX to TX mode respectively. As we see later, this is sufficient to meet the timing requirements of IoT packet ACKs and MAC protocols. Further, the delay for switching between different frequencies is only $220 \mu\text{s}$. To measure this number, we switch between 2.402 GHz, 2.426 GHz and 2.480 GHz. This switching delay is again sufficient to meet the requirements of frequency hopping during Bluetooth advertising.

Transmitter performance. First, we implement a single-tone modulator on the FPGA that generates the appropriate I/Q samples and streams them over LVDS to the radio. We connect the output to an MDO4104b-6 [28] spectrum analyzer and observe a single tone, shown in Fig. 5.8, with no




Figure 5.9: **Single-Tone Transmitter Power Consumption.** We show the total power consumption of tinySDR including I/Q radio, FPGA, MCU and regulators at different transmitter output power. This is 15-16 times lower power consumption than the USRP E310 embedded SDR.

unexpected harmonics introduced by the modulator.

Next we measure the end-to-end DC power consumption of our system including the I/Q radio, FPGA, MCU and regulators to see how it scales with RF output power. We vary our radio output power while transmitting a single tone and use a Fluke 287 multimeter to measure its DC power draw. Fig. 5.9 shows the power consumption of tinySDR for 900 MHz and 2.4 GHz operation. Interestingly, we observe the DC power is constant at low RF power but increases as expected beyond some RF power level. TinySDR consumes 231 mW when transmitting at 0 dBm, and for comparison the end-to-end power consumption of the USRP E310 is 16x higher under the same conditions. Similarly tinySDR consumes 283 mW at its 14 dBm setting while the USRP E310 is 15x higher. In addition, we measure the peak current consumption of tinySDR while transmitting a single-tone on the I/Q radio. The peak current consumption is 105 mA when tinySDR boots up the FPGA and then starts transmitting using the I/Q radio. This current is less than the maximum current supported by a typical LiPo 1200 mAh battery [31].

Cost. We also analyze the cost which is an important practical consideration for real world deployment at scale. Table 5.5 shows a detailed breakdown of cost including each component as well as PCB fabrication and assembly based on quotes for 1000 units [48], where the overall cost is around \$55 each.




Figure 5.10: **LoRa Modulator Evaluation.** We evaluate our LoRa modulator in comparison with Semtech LoRa chip.

5.5.2 Evaluating the Case Studies

LoRa using tinySDR. We evaluate various different components of tinySDR using LoRa as a case study.

LoRa modulator. To evaluate this, we use our LoRa modulator to generate packets with three byte payloads using a spreading factor of $SF = 8$ and bandwidths of 250 kHz and 125 kHz which we transmit at -13 dBm. We receive the output of tinySDR on a Semtech SX1276 LoRa transceiver [75] which we use to measure the packet error rate (PER) versus RSSI and plot the results in Fig. 5.10. We compare our LoRa modulator to transmissions from an SX1276 LoRa transceiver. The plots show that we can achieve a comparable sensitivity of -126 dBm which is the LoRa sensitivity for $SF = 8$ and $BW = 125kHz$ configuration. This is true for both configurations, which shows that our low-power SDR can meet the sensitivity requirement of LPWAN IoT protocols.

LoRa demodulator. Next we evaluate our LoRa demodulator on tinySDR. To test this, we use transmissions from a Semtech SX1276 LoRa transceiver and use tinySDR to receive these transmissions. The LoRa transceiver transmits packets with two configurations using a spreading factor of 8 and bandwidths of 250 kHz and 125 kHz. We record the received RF signals in the FPGA memory and run them through our demodulator to compute a chirp symbol error rate. Note that the Semtech LoRa transceiver does not give access to its symbol error rate, but since we have access to I/Q samples, we can compute it on our platform. We plot the results in Fig. 5.11 as a function of




Figure 5.11: **LoRa Demodulator Evaluation.** We evaluate our LoRa demodulator by demodulating chirp symbols at different RSSI.

the LoRa RSSI values. Our LoRa demodulator can demodulate chirp symbols down to -126 dBm which is LoRa protocol sensitivity at $SF = 8$ and $BW = 125kHz$. Both the LoRa modulator and demodulator run in real-time.

Resource allocation. Next, we evaluate the resource utilization of our LoRa PHY implementation on the FPGA. Table 5.6 shows the size for implementing the modulator and demodulator on our FPGA using different SFs. Our LoRa modulator supports all LoRa configurations with different SF with no additional cost. However, in the LoRa demodulator, we need FFT blocks with different sizes to support different SF configurations. This table shows that our FPGA has sufficient resources to support multiple configurations of LoRa and still leave space for other custom operations.

LoRa MAC. We implement the LoRa MAC based on TTN’s Arduino libraries [79]. TTN protocol together with control for the I/Q radio, backbone radio, FPGA, PMU and decompression algorithm for OTA take only 18% of MCU resources. Also, as shown in Table 5.4, our timings are well within the requirements for LoRaWAN specifications [60].

We also measure the power consumption of our platform for LoRa packet transmission and reception. LoRa packet transmission with $SF = 9$ and $BW = 500 kHz$ and radio output power of 14 dBm consumes a total power of 287 mW from which 179 mW is for the radio and the rest is from the FPGA and MCU. LoRa packet reception consumes 186 mW with radio taking 59 mW.

BLE using tinySDR. Next, we evaluate tinySDR using BLE beacons as a case study. First, we

Table 5.6: FPGA Utilization for LoRa Protocol.

SF	LoRa TX (LUT)	LoRa RX (LUT)
6	976 (4%)	2656 (10%)
7	976 (4%)	2670 (10%)
8	976 (4%)	2700 (11%)
9	976 (4%)	2742 (11%)
10	976 (4%)	2786 (11%)
11	976 (4%)	2794 (11%)
12	976 (4%)	2818 (11%)

measure the impact of our BLE beacons transmitted from tinySDR using the TI CC2650 [44] BLE chip as a receiver. We do this by configuring tinySDR to transmit BLE beacons at a rate of 1 packet per second. We transmit 100 packets and set the CC2650 BLE chip to report bit error rate (BER). Fig. 5.12 shows the BER as a function of the received RSSI as reported by the CC2650 BLE chip. The plot shows that we achieve a sensitivity of -94 dBm. This is within 2 dB of the CC2650 BLE chipset's sensitivity, defined by a BER threshold of 10^{-3} .

Next we evaluate the latency of our BLE implementation as BLE beacons are typically transmitted in sequence by hopping between three different advertising channels. We measure the minimum time tinySDR takes to switch between these frequencies by connecting its output to a 2.4 GHz envelope detector and using an MDO4104B-6 oscilloscope to measure the time delay between transmissions. Fig. 5.13 plots the envelope of three BLE beacons in the time-domain transmitted on the different advertising channels and shows that our system can transmit packets with as little as 220 us delay between beacons. The corresponding result when a iPhone 8 transmits beacons is 350 us. Finally, generating BLE beacons requires only 3% of the FPGA resources on the tinySDR and it could run for over 2 years on a 1000 mAh battery when transmitting once per second.

5.5.3 Over-the-Air Programming

An effective OTA programming system should both minimize use of system resources such as power as well as network downtime. Considering the time to reprogram the FPGA and microcontroller from flash is fixed, the downtime for programming a node depends on the amount of data sent and the throughput which varies with SNR.

Raw programming files for our FPGA are 579 kB, however we compress our data using miniLZO. While the exact compression ratio depends on FPGA utilization, our LoRa program compresses to 99 kB and BLE to 40 kB. Our microcontroller programs for both LoRa and BLE are approximately 78 kB and are both compressed to 24 kB. When dividing the files into packets, we would ideally minimize the preamble length and maximize packet length to reduce overhead, however long packets with short preambles lead to higher PER. We choose a preamble of 8 chirps and packets of 60 B which we find balances the trade-off of protocol overhead versus range in our experiments.

To see the impact on a real deployment, we evaluate the time required to program tinySDR nodes in our 20 device testbed shown in Fig. 5.7. We set up a LoRa transceiver configured with $SF = 8$, $BW = 500 \text{ kHz}$ and $CodingRate = 6$ connected to a patch antenna transmitting at 14 dBm as an AP and measure the time it takes to program the tinySDR devices at each location, according to our protocol. We transmit the compressed FPGA and MCU programming data for LoRa and BLE and plot the results as a CDF in Fig. 5.14. The plots show that the LoRa FPGA requires an average programming time of 150 s while BLE, FPGA, and MCU require 59 s and 39 s respectively due to their smaller file size. Decompressing these received files only takes a maximum of 450 ms. The variation of the programming time between different nodes is caused by the variation in the wireless channel and hence the number of re-transmissions for each tinySDR node is different.

Our OTA programming system components, backbone radio and MCU, consume an average energy of 6144 mJ for receiving a LoRa FPGA update and 2342 mJ for a BLE FPGA update when using 14 dBm output power. Using a 1000 mAh LiPo battery, we could OTA program each tinySDR node with LoRa 2100 times and BLE 5600 times. Assuming OTA programming of once




Figure 5.12: **BLE evaluation.** BLE beacons at different power levels.




Figure 5.13: **BLE Beacons Signal.** We show BLE beacon transmissions on three advertising channels from tinySDR using an envelope detector.

per day, the average power consumption would be 71 μW and 27 μW respectively for LoRa and BLE.

5.6 Research Study: Concurrent Reception

An SDR designed for IoT endpoints that can provide I/Q transmission and reception capability opens up opportunities for addressing multiple research questions in IoT networks.

In this section, we focus on the following question: *Can a low-power IoT endpoint device decode multiple concurrent LoRa transmissions at the same time?* LoRa supports long range communication for IoT devices and is gaining popularity as a low-power wide area networking (LPWAN) standard. Supporting long ranges introduces new challenges since it increases the probability of




Figure 5.14: **OTA Programming Time.** We show CDF of OTA programming time for programming LoRa and BLE implementations on tinySDR.

collisions in large scale city-wide deployments. While recent works [96, 106] have explored the feasibility of enabling concurrent LoRa transmissions, they have been designed for decoding on a gateway-style USRP device. In fact, most concurrent transmission techniques in our community [100, 118, 96] have been prototyped on USRPs and it is unclear if a low-power IoT endpoint device can decode concurrent transmissions in real-time within its stringent power and resource constraints. TinySDR enables us to explore such questions and design MAC protocols for decoding concurrent transmissions on IoT endpoints.

Using orthogonal LoRa codes. Here we explore a specific way of enabling concurrent transmissions in LoRa: using orthogonal codes. Specifically, to allow multiple LoRa nodes to communicate at the same time, we exploit LoRa’s support for orthogonal transmissions [36] which can occupy the same frequency channel without interfering with each other. Two chirp symbols are orthogonal when they have a different chirp slope. For a chirp with a spreading factor of SF and bandwidth of BW , the chirp slope is given by: $\frac{BW^2}{2SF}$ [106].

Decoding concurrent transmissions on tinySDR. In order to receive concurrent LoRa transmissions, tinySDR must be able to demodulate LoRa upchirp symbols with different slopes. Suppose we have two LoRa transmissions that use different spreading factor and bandwidth configurations: SF_1, BW_1 and SF_2, BW_2 . To decode them concurrently, we implement decoders similar to Fig. 5.6(b) for each chirp configuration in parallel on our FPGA. Specifically, we first generate a correspond-


ing downchirp symbol for each configuration in real-time using our chirp generator. Note that we generate each chirp with its corresponding configuration on the FPGA and we do not use pre-generated chirps on the FPGA. We then correlate the received signals with their corresponding downchirp symbols using time domain multiplication. After correlation, we take the appropriate length FFT of the result.

Evaluation. We evaluate three key aspects of our design: 1) the platform's effectiveness in decoding concurrent transmissions across a range of RSSI values, 2) the power consumption at the endpoint device while decoding concurrent transmissions and 3) the computational resources required.


We use two SX1276 LoRa transceivers as our transmitters and set them to transmit continuously at two different settings: they both use a spreading factor of $SF = 8$ but have two different bandwidth setups, $BW_1 = 125kHz$ and $BW_2 = 250kHz$. We set the two to send random chirp symbols. The tinySDR platform decodes these two concurrent transmissions and computes the chirp symbol error rate for each transmission. We evaluate two scenarios: 1) when the two transmitters have a similar power level at the receiver, 2) fix the power of one of the transmitters and increase the power of the other one.

Fig. 5.15(a) shows the results when the two transmissions have similar power at the receiver. We lose around 2 dB and 0.5 dB sensitivity for concurrent demodulation of LoRa configurations with $BW_1 = 125kHz$ and $BW_2 = 250kHz$. This is because while in theory the two chirps are orthogonal, in practice, the chirps are created in the digital domain with discrete frequency steps which introduces some non-orthogonality.

Fig. 5.15(b) shows the results when the first LoRa transmitter $BW_1 = 125kHz$ is received near its sensitivity of -123 dBm and the second LoRa transmitter changes its power. Here, the chirp symbol error rate is affected when the other transmission's power is higher than -116 dBm. When two concurrent transmissions are present, one acts as an interferer when decoding the other. The combined power of noise and the interferer, $P_{I,N}$, determines the error rate. When sweeping the power of interferer, at first the $P_{I,N}$ is dominated by noise and we should not see much effect on error rate. Then at some point their power would be equal which results in a 3 dB increase of $P_{I,N}$



((a)) Orthogonal Transmissions with Same Received Signal Power.



((b)) Orthogonal Transmissions with Different Received Signal Power. We set the power of LoRa transmission with $BW_1 = 125kHz$ to -123 dBm and increase the power of the other one.

Figure 5.15: Orthogonal LoRa Demodulation Evaluation

and hence 3 dB sensitivity loss after which the error rate is determined by the interferer power. This demonstrates the need for power control for concurrent transmissions to be received on IoT endpoints.

Our parallel demodulation implementation, uses only 17% of the FPGA resources. This concurrent demodulation implementation consumes 207 mW. Note that Semtech gateway solutions such as the SX1308 [66] can receive multiple transmissions. But, to the best of our knowledge we are the first to show that concurrent LoRa transmissions can be decoded on an IoT endpoint while meeting its power and computational requirements. This would have been difficult to do without tinySDR.

5.7 Conclusion and Research Opportunities

This paper presents the first SDR platform specifically tailored to the needs of IoT endpoints that can be used for large scale IoT network deployments. The goal of tinySDR is to provide a platform that can catalyze research in IoT networks.

Research on PHY/MAC protocols. TinySDR presents an opportunity for researchers to avoid the time consuming endeavor of building their own custom hardware and instead focus on PHY/MAC protocol innovations across the stack: What is the trade-off between packet length and overall throughput? Are there benefits of rate adaptation? What about concurrent transmissions from IoT devices? One could also create multi-hop IoT PHY/MAC innovations, which have not been explored well given the lack of a flexible platform.

Research on IoT localization. TinySDR could also be used to build localization systems as it gives access to I/Q signals and therefore phase across the 2.4 GHz and 900 MHz bands, which forms the basis for many localization algorithms[149]. One could also explore distributed localization solutions that combine the phase information across a distributed set of sensors to create a large MIMO sensing system.

Machine learning on IoT devices. The FPGA on tinySDR opens up exciting opportunities [88] for exploring machine learning algorithms on-board. This would allow researchers to explore trade-offs between the power overhead of running an on-board classifier versus sending data to the cloud. This could also enable use of high bandwidth sensors such as cameras and microphones where the power bottleneck may be communication rather than sensing.

Low power backscatter readers. Recent work on ambient backscatter [132, 119, 120, 114, 190] aims to achieve ultra-low power communication for IoT devices. Many of these proposals require either a single-tone generator [120] or a custom receiver to decode the backscatter transmissions [164, 112, 111, 141]. TinySDR can be used as a building block to achieve a battery-operated backscatter signal generation and receiver.

Better programming interface and protocols. In addition to IoT research opportunities, we can also improve our platform in multiple ways. TinySDR currently requires users to write Verilog or

VHDL to program the FPGA and C code for programming the microcontroller. Future versions can incorporate a pipeline to use high level synthesis tools or integrate with GNUradio for easy prototyping. Further, tinySDR uses a simple MAC protocol for programming with a focus on using minimal system resources to allow for other custom software; however we could explore modified MAC protocols that simultaneously broadcast the updates across the network to reduce programming time.

Chapter 6

CONCLUSION

In conclusion, this dissertation presents wireless technologies and platforms that address challenges in achieving ubiquitous connectivity. We designed a wireless technology that takes a significant step towards enabling battery-free HD video streaming. We implemented a low resolution version of this technology that performs battery-free video streaming to a reader in real-time. We developed wireless technologies to increase the range of backscatter communication to enable wide-area backscatter communication. Building on this work, we presented a new wireless protocol for backscatter networks that scales to hundreds of concurrent transmissions by introducing a distributed coding mechanism. Finally, we presented the first SDR platform specifically tailored to the needs of IoT endpoints that can be used for large scale IoT network deployments. We evaluated our design in various scenarios.

The work presented in this dissertation can be followed up in future along these directions.

- **End-to-end battery-free HD video streaming system.** Existing works on backscatter cameras focus on using discrete components and use cameras and backscatter independently and just interface them. Our key contributions in chapter 2 is to make a case for camera and backscatter architecture. However, we believe that building the ASIC version of this holistic architecture can show the full capacity of the system and enables other research directions. This technology could be used for power-constrained technologies like insect-scale robots [113].
- **Improved data rates in LoRa backscatter.** Our current LoRa backscatter design achieves all the bit rates supported by LoRa protocol. We believe that one could achieve higher data rates at desired ranges by using a combination of multiple frequency bands to transmit at the same time or use multiple antennas. In addition, designing an efficient MAC layer for LoRa backscatter

that adapts the data rate based on signal quality would decrease the number of re-transmissions and improve overall network throughput.

- **Co-located RF carrier and receiver backscatter deployment.** The current design of most novel backscatter systems use a half-duplex approach where the RF carrier source and receiver are separated due to the affect of carrier signal on the receiver's sensitivity. Designing a full-duplex reader [117] for NetScatter system, explained in chapter 4, where both RF carrier and receiver are co-located can simplify its deployment.
- **TinySDR programming interface.** TinySDR currently requires users to write Verilog or VHDL to program the FPGA and C code for programming the microcontroller. Future versions can incorporate a pipeline to use high level synthesis tools or integrate with GNUradio for easy prototyping.
- **Research opportunities using TinySDR.** TinySDR provides flexibility to explore various research directions. It could be used to build a localization system as it gives access to I/Q signals and therefore phase information across both sub-GHz and 2.4 GHz frequency bands. In addition, the FPGA on TinySDR opens up exciting opportunities for exploring machine learning algorithm on-board. Moreover, TinySDR could be used as a battery-powered backscatter signal generator and receiver for backscatter systems.

BIBLIOGRAPHY

- [1] 6-dBi gain patch antenna. https://www.arcantenna.com/index.php/product_documents/get/document/id/492/. Accessed: March 2017.
- [2] Alien swiggle rfid tag. <http://www.rfidtags.com/squiggle-h3-rfid-tag>.
- [3] Ambarella a9 ultra hd 4k camera soc product brief. <http://www.ambarella.com/uploads/docs/A9-product-brief.pdf>.
- [4] Barker code. <http://mathworld.wolfram.com/BarkerCode.html>. Accessed: September 2017.
- [5] bladerf 2.0 micro. https://www.nuand.com/bladeRF_2_micro-brief.pdf.
- [6] CentEye Stonyman image sensor datasheet. <http://www.centeye.com/products/current-centeye-vision-chips/>. Accessed: March 2017.
- [7] The great seal bug. <http://counterespionage.com/the-great-seal-bug-part-1.html>. Accessed: March 2017.
- [8] Head jack pluggable rfid reader for smartphones. http://www.rfidtagworld.com/products/Ear-Jack-reader-UHF_1056.html.
- [9] IGLOO Nano FPGA datasheet. https://www.microsemi.com/document-portal/doc_view/130695-ds0110-igloo-nano-low-power-flash-fpgas-datasheet. Accessed: March 2017.
- [10] Limesdr. <https://myriadrif.org/projects/limesdr/>.
- [11] Limesdr-mini. <https://wiki.myriadrif.org/LimeSDR-Mini>.
- [12] Lmv7219 comparator. <http://www.ti.com/lit/ds/symlink/lmv7219.pdf>. Accessed: September 2017.
- [13] Moving camera on a motorcycle video clip. https://www.youtube.com/watch?v=sHj3xSG-R_E&t=376s.

- [14] NCX2200I low power comparator datasheet. http://www.nxp.com/documents/data_sheet/NCX2200.pdf. Accessed: March 2017.
- [15] An overview of lvds technology. <http://www.ti.com/lit/an/snla165/snla165.pdf>.
- [16] RF5110 amplifier. http://www.rfmd.com/store/downloads/dl/file/id/30508/5110g_product_data_sheet.pdf. Accessed: March 2017.
- [17] Ring video doorbell. <https://ring.com/>. Accessed: March 2017.
- [18] Rtl2832u dvb-t tuner dongles. <https://www.rtl-sdr.com/buy-rtl-sdr-dvb-t-dongles/>.
- [19] Spectacles battery life. <http://www.barrons.com/articles/snapchat-spectacles-review-the-good-the-bad-the-revolutionary-1487846715>.
- [20] Synopsis design compiler. <http://www.synopsys.com/Tools/Implementation/RTLSynthesis/DesignCompiler/Pages/default.aspx>.
- [21] Usrp b200mini. https://www.ettus.com/content/files/USRP_B200mini_Data_Sheet.pdf.
- [22] Usrp e310. https://www.ettus.com/content/files/USRP_E310_Datasheet.pdf.
- [23] Wearable spectacles. <http://www.techradar.com/reviews/snap-spectacles>.
- [24] Youtube website. <http://www.youtube.com>.
- [25] 2.45 ghz balun, filter combination, 2003. https://www.mouser.com/datasheet/2/611/JTI_Balun-Filter-2450FB15A050_2006-09-242325.pdf.
- [26] Max5189 datasheet, 2003. <https://datasheets.maximintegrated.com/en/ds/MAX5186-MAX5189.pdf>.
- [27] Max2831 datasheet, 2011. <https://datasheets.maximintegrated.com/en/ds/MAX2831-MAX2832.pdf>.
- [28] Mdo4000b series datasheet, 2013. <http://www.testequipmenthq.com/datasheets/TEKTRONIX-MD04104B-6-Datasheet.pdf>.
- [29] Smart cities and the internet of everything: The foundation for delivering next-generation citizen services, 2013. http://www.cisco.com/c/dam/en_us/solutions/industries/docs/scc/ioe_citizen_svcs_white_paper_idc_2013.pdf.

- [30] Ad9364 transceiver datasheet, 2014. <https://www.analog.com/media/en/technical-documentation/data-sheets/AD9364.pdf>.
- [31] Li-polymer battery technology specification, 2014. https://cdn-shop.adafruit.com/product-files/258/C101-Li-Polymer_503562_1200mAh_3.7V_with_PCM_APPROVED_8.18.pdf.
- [32] Smartsense temp/humidity manual, 2014. <https://support.smartthings.com/hc/en-us/articles/203040294-SmartSense-Temperature-Humidity-Sensor>.
- [33] 3.5mhz, 500ma synchronous step down dc-dc regulator, 2015. <https://www.semtech.com/uploads/documents/sc195.pdf>.
- [34] Igloo nano fpga datasheet by, 2015. https://www.microsemi.com/document-portal/doc_download/130695-igloo-nano-low-power-flash-fpgas-datasheet.
- [35] The internet of things and the future of farming, 2015. http://bits.blogs.nytimes.com/2015/08/03/the-internet-of-things-and-the-future-of-farming/?_r=0.
- [36] Lora modulation basics, 2015. <https://www.semtech.com/uploads/documents/an1200.22.pdf>.
- [37] Xrcha-f-a series, 2015. <https://www.murata.com/~media/webrenewal/products/timingdevice/crystalu/flyers/vppt-hcrj078-d.ashx?la=en-us>.
- [38] 65 nm cmos process by tsmc, 2016. <http://www.tsmc.com/english/dedicatedFoundry/technology/65nm.htm>.
- [39] 915mhz whip, straight rf antenna by nearson. 2016. <ftp://ftp2.nearson.com/Drawings/Antenna/S463XX-915.pdf>.
- [40] Ad9361 transceiver datasheet, 2016. <https://www.analog.com/media/en/technical-documentation/data-sheets/AD9361.pdf>.
- [41] Ad9363 transceiver datasheet, 2016. <https://www.analog.com/media/en/technical-documentation/data-sheets/AD9363.pdf>.
- [42] Adg904 datasheet by analog devices, 2016. <http://www.analog.com/media/en/technical-documentation/data-sheets/ADG904.pdf>.

- [43] At86rf215 datasheet, 2016. http://ww1.microchip.com/downloads/en/devicedoc/atmel-42415-wireless-at86rf215_datasheet.pdf.
- [44] Cc2650 simplelink datasheet by ti, 2016. <http://www.ti.com/lit/ds/symlink/cc2650.pdf>.
- [45] How much does an RFID tag cost today?, 2016. <https://www.rfidjournal.com/faq/show?85>.
- [46] Lora alliance, 2016. <https://www.lora-alliance.org/>.
- [47] New part 15 rules by fcc, 2016. https://transition.fcc.gov/oet/ea/presentations/files/may05/New_Policies_Pt._15_SD.pdf.
- [48] Pcbminions inc., 2016. <https://pcbminions.com/>.
- [49] Se2435l power amplifier datasheet, 2016. http://www.skyworksinc.com/uploads/documents/SE2435L_202412I.pdf.
- [50] Semtech sx1276 transceiver, 2016. <http://www.semtech.com/wireless-rf/rf-transceivers/sx1276/>.
- [51] Sx1276 datasheet by semtech, 2016. <https://www.semtech.com/uploads/documents/sx1276.pdf>.
- [52] Altera's cyclone v fpgas. 2017. <https://www.altera.com/products/fpga/cyclone-series/cyclone-v/overview.html>.
- [53] Brightvolt product matrix, 2017. <http://www.brightvolt.com/products/product-matrix-2/>.
- [54] Cadence rfspectre. 2017. http://www.cadence.com/products/rf/spectre_rf_simulation/pages/default.aspx.
- [55] CC3200 SimpleLink Wi-Fi and Internet-of-Things Solution, 2017. <http://www.ti.com/product/CC3200/samplebuy>.
- [56] De0-cv development kit. 2017. <http://www.terasic.com.tw/cgi-bin/page/archive.pl?Language=English&CategoryNo=163&No=921&PartNo=2>.

- [57] EFL700A39- rechargeable thin film batteries by stmicroelectronics, 2017. <http://www.st.com/content/ccc/resource/technical/document/datasheet/cd/ac/89/0b/b4/8e/43/0b/CD00270103.pdf/files/CD00270103.pdf/jcr:content/translations/en.CD00270103.pdf>.
- [58] Flexible, printed and thin film batteries 2016-2026: Technologies, markets, players, 2017. <http://www.idtechex.com/research/reports/flexible-printed-and-thin-film-batteries-2016-2026-technologies-markets-players-asp>.
- [59] Lms7002m datasheet, 2017. <https://limemicro.com/app/uploads/2017/07/LMS7002M-Data-Sheet-v3.1r00.pdf>.
- [60] Lora alliance, 2017. https://loro-alliance.org/sites/default/files/2018-04/lorawantm_specification_v1.1.pdf.
- [61] Lte for iot, 2017. <http://resources.alcatel-lucent.com/asset/200178>.
- [62] Msp430fr5969 datasheet by ti, 2017. <http://www.ti.com/lit/ds/symlink/msp430fr5969.pdf>.
- [63] Msp432p401r, msp432p401m simplelink mixed-signal microcontrollers datasheet, 2017. <http://www.ti.com/lit/ds/symlink/msp432p401r.pdf>.
- [64] Sigfox products, 2017.
- [65] Sky66112 power amplifier datasheet, 2017. http://www.skyworksinc.com/uploads/documents/SKY66112_11_203225L.pdf.
- [66] Sx1308 datasheet by semtech, 2017. <https://www.semtech.com/uploads/documents/sx1308.pdf>.
- [67] Ti cc2650. 2017. <http://www.digikey.com/product-detail/en/CC2650F128RHBR/CC2650F128RHBR-ND/5189550>.
- [68] Ad9228 datasheet, 2018. <https://www.analog.com/media/en/technical-documentation/data-sheets/ad9228.pdf>.
- [69] Adalm-pluto overview, 2018. <https://wiki.analog.com/university/tools/pluto>.
- [70] Atmel at86rf215 868/915/928 mhz impedance matched balun + lpf, 2018. <https://www.mouser.com/datasheet/2/611/0896BM15E0025-1518705.pdf>.

- [71] Graphene temporary tattoo tracks vital signs. 2018. <http://spectrum.ieee.org/nanoclast/semiconductors/nanotechnology/graphene-temporary-tattoo>.
- [72] Lfe5u fpga family datasheet, 2018. http://www.latticesemi.com/view_document?document_id=50461.
- [73] minilzo implementation, 2018. <http://www.oberhumer.com/opensource/lzo/#abstract>.
- [74] Rf5110 rf power amplifier, 2018. http://www.rfmd.com/store/downloads/dl/file/id/30508/5110g_product_data_sheet.pdf.
- [75] Sx1257 datasheet by semtech, 2018. https://www.semtech.com/uploads/documents/DS_SX1257_V1.2.pdf.
- [76] Usrc x-300, 2018. <https://www.ettus.com/product/details/X300-KIT>.
- [77] Bluetooth core specification v5.1, January 2019. <https://www.bluetooth.com/specifications/bluetooth-core-specification>.
- [78] The things network, 2019. <https://www.thethingsnetwork.org/>.
- [79] The things network arduino library, 2019. <https://github.com/TheThingsNetwork/arduino-device-lib>.
- [80] Ali Afsahi, Jacob J Rael, Arya Behzad, Hung-Ming Chien, Michael Pan, Stephen Au, Ade-dayo Ojo, C Paul Lee, Seema Butala Anand, Kevin Chien, et al. A low-power single-weight-combiner 802.11 abg soc in 0.13 μm cmos for embedded applications utilizing an area and power efficient cartesian phase shifter and mixer circuit. *IEEE Journal of Solid-State Circuits*, 43(5):1101–1118, 2008.
- [81] Kiarash Amiri, Yang Sun, Patrick Murphy, Chris Hunter, Joseph R Cavallaro, and Ashutosh Sabharwal. Warp, a unified wireless network testbed for education and research. In *2007 IEEE International Conference on Microelectronic Systems Education (MSE'07)*, pages 53–54. IEEE, 2007.
- [82] Narendra Anand, Ehsan Aryafar, and Edward W Knightly. Warplab: a flexible framework for rapid physical layer design. In *Proceedings of the 2010 ACM workshop on Wireless of the students, by the students, for the students*, pages 53–56. ACM, 2010.
- [83] Amay J. Bandodkar and Joseph Wang. Non-invasive wearable electrochemical sensors: a review. *Trends in Biotechnology*, 32(7):363 – 371, 2014.

- [84] H Randall Beard, Alejandro J Marquez-Lara, and Kamran S Hamid. Using wearable video technology to build a point-of-view surgical education library. *JAMA surgery*, 151(8):771–772, 2016.
- [85] A. Berni and W. Gregg. On the utility of chirp modulation for digital signaling. *IEEE Transactions on Communications*, 1973.
- [86] Steven R Best and Bradley C Kaanta. A tutorial on the receiving and scattering properties of antennas. *IEEE Antennas and Propagation Magazine*, 51(5), 2009.
- [87] Ludovic Champion and Nicolas Sornin. Chirp signal processor, 2014. European Patent Application EP2975814A1.
- [88] Justin Chan, Anran Wang, Arvind Krishnamurthy, and Shyamnath Gollakota. Deepsense: Enabling carrier sense in low-power wide area networks using deep learning. *CoRR*, abs/1904.10607, 2019.
- [89] Peter KC Chan, Wei Jin, JM Gong, and NS Demokan. Multiplexing of fiber bragg grating sensors using a fmcw technique. *IEEE Photonics Technology Letters*, 11(11):1470–1472, 1999.
- [90] Vipul Chawla and Dong Sam Ha. An overview of passive rfid. *IEEE Communications Magazine*, 45(9), 2007.
- [91] Long Chen, Xiyuan Tang, Arindam Sanyal, Yeonam Yoon, Jie Cong, and Nan Sun. A 0.7-v 0.6-uw 100 – ks/s low-power sar adc with statistical estimation-based noise reduction. *IEEE Journal of Solid-State Circuits*, 52(5):1388–1398, 2017.
- [92] Tzung-Ming Chen, Yung-Ming Chiu, Chun-Cheng Wang, Ka-Un Chan, Ying-Hsi Lin, Ming-Chong Huang, Chao-Hua Lu, Wen-Shan Wang, Che-Sheng Hu, Chao-Cheng Lee, et al. A low-power fullband 802.11 a/b/g wlan transceiver with on-chip pa. *IEEE journal of solid-state circuits*, 42(5):983–991, 2007.
- [93] Namjun Cho, Seong-Jun Song, Sunyoung Kim, Shiho Kim, and Hoi-Jun Yoo. A 5.1-/spl μ w uhf rfid tag chip integrated with sensors for wireless environmental monitoring. In *Solid-State Circuits Conference, 2005. ESSCIRC 2005. Proceedings of the 31st European*, pages 279–282. IEEE, 2005.
- [94] Daniel MJ Devasirvatham. Time delay spread measurements of wideband radio signals within a building. *Electronics Letters*, 20(23):950–951, 1984.

- [95] Prabal Dutta, Ye-Sheng Kuo, Akos Ledeczi, Thomas Schmid, and Peter Volgyesi. Putting the software radio on a low-calorie diet. In *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks*, page 20. ACM, 2010.
- [96] Rashad Eletreby, Diana Zhang, Swarun Kumar, and Osman Yağın. Empowering low-power wide area networks in urban settings. SIGCOMM '17.
- [97] 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*.
- [98] Lingzhi Fu, Lirui Liu, Min Li, and Junyu Wang. Collision recovery receiver for epc gen2 rfid systems. In *Internet of Things (IOT), 2012 3rd International Conference on the*, pages 114–118, Oct 2012.
- [99] Chuhan Gao, Mehrdad Hesar, Krishna Chintalapudi, and Bodhi Priyantha. Blind distributed mu-mimo for iot networking over vhf narrowband spectrum. In *The 25th Annual International Conference on Mobile Computing and Networking*, pages 1–17, 2019.
- [100] Shyamnath Gollakota and Dina Katabi. Zigzag decoding: Combating hidden terminals in wireless networks. In *Proceedings of the ACM SIGCOMM 2008 Conference on Data Communication*, SIGCOMM '08, 2008.
- [101] Rafael C Gonzalez and Richard E Woods. *Image processing*, volume 2. 2007.
- [102] Yeswanth Guddeti, Raghav Subbaraman, Moein Khazraee, Aaron Schulman, and Dinesh Bharadia. Sweepsense: Sensing 5 ghz in 5 milliseconds with low-cost radios. In *16th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 19)*, pages 317–330, 2019.
- [103] Christian Hambeck, Stefan Mahlknecht, and Thomas Herndl. A $2.4\mu\text{w}$ wake-up receiver for wireless sensor nodes with- 71dbm sensitivity. In *2011 IEEE International Symposium of Circuits and Systems (ISCAS)*, pages 534–537. IEEE, 2011.
- [104] Scott Hanson, ZhiYoong Foo, David Blaauw, and Dennis Sylvester. A 0.5 v sub-microwatt cmos image sensor with pulse-width modulation read-out. *IEEE Journal of Solid-State Circuits*, 45(4):759–767, 2010.
- [105] Mehrdad Hesar, Vikram Iyer, and Shyamnath Gollakota. Enabling on-body transmissions with commodity devices. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 1100–1111, 2016.

- [106] Mehrdad Hesar, Ali Najafi, and Shyamnath Gollakota. Netscatter: Enabling large-scale backscatter networks. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*, 2019.
- [107] Mehrdad Hesar, Ali Najafi, Vikram Iyer, and Shyamnath Gollakota. Tinysdr: Low-power {SDR} platform for over-the-air programmable iot testbeds. In *17th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 20)*, pages 1031–1046, 2020.
- [108] Ralf Hildebrandt. The pseudo dual-edge d-flip-flop, 2011. http://www.ralf-hildebrandt.de/publication/pdf_dff/pde_dff.pdf.
- [109] Pan Hu, Pengyu Zhang, and Deepak Ganesan. Laissez-faire: Fully asymmetric backscatter communication. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication, SIGCOMM '15*.
- [110] Pan Hu, Pengyu Zhang, and Deepak Ganesan. Leveraging interleaved signal edges for concurrent backscatter. In *hotWireless*, 2014.
- [111] Vikram Iyer, Justin Chan, Ian Culhane, Jennifer Mankoff, and Shyamnath Gollakota. Wireless analytics for 3d printed objects. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology, UIST '18*, pages 141–152, New York, NY, USA, 2018. ACM.
- [112] Vikram Iyer, Justin Chan, and Shyamnath Gollakota. 3d printing wireless connected objects. *ACM Trans. Graph.*, 36(6), November 2017.
- [113] Vikram Iyer, Ali Najafi, Johannes James, Sawyer Fuller, and Shyamnath Gollakota. Wireless steerable vision for live insects and insect-scale robots. *Science Robotics*, 5(44), 2020.
- [114] Vikram Iyer, Vamsi Talla, Bryce Kellogg, Shyamnath Gollakota, and Joshua Smith. Inter-technology backscatter: Towards internet connectivity for implanted devices. In *Proceedings of the 2016 Conference on ACM SIGCOMM 2016 Conference, SIGCOMM '16*, pages 356–369, New York, NY, USA, 2016. ACM.
- [115] Meng Jin, Yuan He, Xin Meng, Yilun Zheng, Dingyi Fang, and Xiaojiang Chen. Fliptracer: Practical parallel decoding for backscatter communication. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking, MobiCom '17*.
- [116] Kerem Kapucu and Catherine Dehollain. A passive uhf rfid system with a low-power capacitive sensor interface. In *RFID Technology and Applications Conference (RFID-TA), 2014 IEEE*, pages 301–305. IEEE, 2014.

- [117] Mohamad Katanbaf, Anthony Weinand, and Vamsi Talla. Simplifying backscatter deployment: Full-duplex lora backscatter. *arXiv preprint arXiv:2011.03952*, 2020.
- [118] Sachin Katti, Shyamnath Gollakota, and Dina Katabi. Embracing wireless interference: Analog network coding. *SIGCOMM Comput. Commun. Rev.*
- [119] 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.
- [120] Bryce Kellogg, Vamsi Talla, Shyamnath Gollakota, and Joshua R Smith. Passive wi-fi: Bringing low power to wi-fi transmissions. In *13th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 16)*, pages 151–164, 2016.
- [121] Walt Kester. Understand sinad, enob, snr, thd, thd+ n, and sfdr so you don't get lost in the noise floor. *MT-003 Tutorial*, www.analog.com/static/importedfiles/tutorials/MT-003.pdf, 2009.
- [122] Ahmed Khattab, Joseph Camp, Chris Hunter, Patrick Murphy, Ashutosh Sabharwal, and Edward W Knightly. Warp: a flexible platform for clean-slate wireless medium access protocol design. *ACM SIGMOBILE Mobile Computing and Communications Review*, 12(1):56–58, 2008.
- [123] Manikanta Kotaru, Pengyu Zhang, and Sachin Katti. Localizing low-power backscatter tags using commodity wifi. In *CoNext'17*.
- [124] Ye-Sheng Kuo, Pat Pannuto, Thomas Schmid, and Prabal Dutta. Reconfiguring the software radio to improve power, price, and portability. In *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems*, pages 267–280. ACM, 2012.
- [125] Ye-Sheng Kuo, Thomas Schmid, and Prabal Dutta. A compact, inexpensive, and battery-powered software-defined radio platform. In *Proceedings of the 11th international conference on Information Processing in Sensor Networks*, pages 137–138. ACM, 2012.
- [126] Hasnain Lakdawala, Mark Schaecher, Chang-Tsung Fu, Rahul Limaye, Jon Duster, Yulin Tan, Ajay Balankutty, Erkan Alpman, Chun C Lee, Khoa Minh Nguyen, et al. A 32 nm soc with dual core atom processor and rf wifi transceiver. *IEEE Journal of Solid-State Circuits*, 48(1):91–103, 2013.
- [127] Chungyeol Paul Lee, Arya Behzad, Bojko Marholev, Vikram Magoon, Iqbal Bhatti, Dandan Li, Subhas Bothra, Ali Afsahi, Dayo Ojo, Rozi Roufoogaran, et al. A multistandard, multi-band soc with integrated bt, fm, wlan radios and integrated power amplifier. In *Solid-State*

- Circuits Conference Digest of Technical Papers (ISSCC), 2010 IEEE International*, pages 454–455. IEEE, 2010.
- [128] Juan Antonio Leñero-Bardallo, Teresa Serrano-Gotarredona, and Bernabé Linares-Barranco. A 3.6 μ s latency asynchronous frame-free event-driven dynamic-vision-sensor. *IEEE Journal of Solid-State Circuits*, 46(6):1443–1455, 2011.
- [129] Yu-Te Liao, Huanfen Yao, Andrew Lingley, Babak Parviz, and Brian P Otis. A 3-cmos glucose sensor for wireless contact-lens tear glucose monitoring. *Solid-State Circuits, IEEE Journal of*, 47(1):335–344, 2012.
- [130] Robert LiKamWa, Bodhi Priyantha, Matthai Philipose, Lin Zhong, and Paramvir Bahl. Energy characterization and optimization of image sensing toward continuous mobile vision. In *Proceeding of the 11th annual international conference on Mobile systems, applications, and services*, pages 69–82. ACM, 2013.
- [131] Yuan Lin, Hyunseok Lee, Mark Woh, Yoav Harel, Scott Mahlke, Trevor Mudge, Chaitali Chakrabarti, and Krisztian Flautner. Soda: A low-power architecture for software radio. *ACM SIGARCH Computer Architecture News*, 34(2):89–101, 2006.
- [132] Vincent Liu, Aaron Parks, Vamsi Talla, Shyamnath Gollakota, David Wetherall, and Joshua R. Smith. Ambient backscatter: Wireless communication out of thin air. SIGCOMM '13.
- [133] Vincent Liu, Vamsi Talla, and Shyamnath Gollakota. Enabling instantaneous feedback with full-duplex backscatter. In *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking*, MobiCom '14.
- [134] H. MahdaviFar and A. Vardy. Coding for tag collision recovery. In *2015 IEEE International Conference on RFID (RFID)*, pages 9–16, April 2015.
- [135] George Malim. Hot iot is expanding into cosmetics and medical industries. 2017. <http://www.iotglobalnetwork.com/iotdir/2016/03/22/how-iot-is-expanding-into-cosmetics-and-medical-industries-1203/>.
- [136] U Mallik, M Clapp, Edward Choi, G Cauwenberghs, and R Etienne-Cummings. Temporal change threshold detection imager. In *Solid-State Circuits Conference, 2005. Digest of Technical Papers. ISSCC. 2005 IEEE International*, pages 362–603. IEEE, 2005.
- [137] Fridolin Michel and Michiel Steyaert. A 250mv 7.5 μ w 61db snr cmos sc $\delta\sigma$ modulator using a near-threshold-voltage-biased cmos inverter technique. In *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2011 IEEE International*, pages 476–478. IEEE, 2011.

- [138] Gary J Minden, Joseph B Evans, Leon Searl, Daniel DePardo, Victor R Petty, Rakesh Rajbanshi, T Newman, Qi Chen, F Weidling, J Guffey, et al. Kuar: A flexible software-defined radio development platform. In *2007 2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, pages 428–439. IEEE, 2007.
- [139] James Morra. Iot devices and wearables push development of thin, flexible batteries. 2017. <http://electronicdesign.com/power/iot-devices-and-wearables-push-development-thin-flexible-batteries>.
- [140] B. Murmann. Adc performance survey 1997-2017. <http://web.stanford.edu/~murmam/adcsurvey.html>.
- [141] Saman Naderiparizi, Mehrdad Hesar, Vamsi Talla, Shyamnath Gollakota, and Joshua R Smith. Towards battery-free hd video streaming. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*, 2018.
- [142] Saman Naderiparizi, Zerina Kapetanovic, and Joshua R Smith. Battery-free connected machine vision with wispcam. *GetMobile: Mobile Computing and Communications*, 20(1):10–13, 2016.
- [143] Saman Naderiparizi, Zerina Kapetanovic, and Joshua R Smith. Wispcam: An rf-powered smart camera for machine vision applications. In *Proceedings of the 4th International Workshop on Energy Harvesting and Energy-Neutral Sensing Systems*, pages 19–22. ACM, 2016.
- [144] Saman Naderiparizi, Zerina Kapetanovic, and Joshua R Smith. Rf-powered, backscatter-based cameras. In *Antennas and Propagation (EUCAP), 2017 11th European Conference on*, pages 346–349. IEEE, 2017.
- [145] Saman Naderiparizi, Aaron N Parks, Zerina Kapetanovic, Benjamin Ransford, and Joshua R Smith. Wispcam: A battery-free rfid camera. In *RFID (RFID), 2015 IEEE International Conference on*, pages 166–173. IEEE, 2015.
- [146] Saman Naderiparizi, Pengyu Zhang, Matthai Philipose, Bodhi Priyantha, Jie Liu, and Deepak Ganesan. Glimpse: A programmable early-discard camera architecture for continuous mobile vision. In *Proceeding of the 15th annual international conference on Mobile systems, applications, and services*. ACM, 2017.
- [147] Saman Naderiparizi, Yi Zhao, James Youngquist, Alanson P Sample, and Joshua R Smith. Self-localizing battery-free cameras. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 445–449. ACM, 2015.

- [148] Hiroyuki Nakamoto, Daisuke Yamazaki, Takuji Yamamoto, Hajime Kurata, Satoshi Yamada, Kenji Mukaida, Tsuzumi Ninomiya, Takashi Ohkawa, Shoichi Masui, and Kunihiko Gotoh. A passive uhf rf identification cmos tag ic using ferroelectric ram in 0.35 μm technology. *IEEE Journal of Solid-State Circuits*, 42(1):101–110, 2007.
- [149] Rajalakshmi Nandakumar, Vikram Iyer, and Shyamnath Gollakota. 3d localization for sub-centimeter sized devices. In *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems, SenSys '18*, 2018.
- [150] Revathy Narayanan and Swarun Kumar. Revisiting software defined radios in the iot era. In *Proceedings of the 17th ACM Workshop on Hot Topics in Networks*. ACM, 2018.
- [151] Lalitkumar Nathawad, David Weber, Shahram Abdollahi, Phoebe Chen, Syed Enam, Brian Kaczynski, Alireza Kheirkhahi, MeeLan Lee, Sotirios Limotyrakis, Keith Onodera, et al. An ieee 802.11 a/b/g soc for embedded wlan applications. In *Solid-State Circuits Conference, 2006. ISSCC 2006. Digest of Technical Papers. IEEE International*, pages 1430–1439. IEEE, 2006.
- [152] Shree K Nayar, Daniel C Sims, and Mikhail Fridberg. Towards self-powered cameras. In *Computational Photography (ICCP), 2015 IEEE International Conference on*, pages 1–10. IEEE, 2015.
- [153] Man Cheuk Ng, Kermin Elliott Fleming, Mythili Vutukuru, Samuel Gross, Hari Balakrishnan, et al. Airblue: A system for cross-layer wireless protocol development. In *Proceedings of the 6th ACM/IEEE Symposium on Architectures for Networking and Communications Systems*, page 4. ACM, 2010.
- [154] Jiajue Ou, Mo Li, and Yuanqing Zheng. Come and be served: Parallel decoding for cots rfid tags. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking, MobiCom '15*, pages 500–511, New York, NY, USA, 2015. ACM.
- [155] J. Pandey, Yu-Te Liao, A. Lingley, R. Mirjalili, B. Parviz, and B. Otis. A fully integrated rf-powered contact lens with a single element display. *Biomedical Circuits and Systems, IEEE Transactions on*, 2010.
- [156] Alexandros Pantelopoulos and Nikolaos G Bourbakis. A survey on wearable sensor-based systems for health monitoring and prognosis. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 40(1):1–12, 2010.
- [157] 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.

- [158] 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.
- [159] Yao Peng, Longfei Shangguan, Yue Hu, Yujie Qian, Xianshang Lin, Xiaojiang Chen, Dingyi Fang, and Kyle Jamieson. Flora: a passive long-range data network from ambient lora transmissions. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication, SIGCOMM '18*, pages 147–160. ACM, 2018.
- [160] Carlos Perez-Penichet, Fredrick Hermans, Ambuj Varshney, and Thiemo Voigt. Augmenting iot networks with backscatter-enabled passive sensor tags. In *HotWireless, 2016*, 2016.
- [161] 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.
- [162] Nathan E Roberts and David D Wentzloff. A 98nw wake-up radio for wireless body area networks. In *2012 IEEE Radio Frequency Integrated Circuits Symposium*, pages 373–376. IEEE, 2012.
- [163] Zhiyu Ru, Eric AM Klumperink, and Bram Nauta. A discrete-time mixing receiver architecture with wideband harmonic rejection. In *Solid-State Circuits Conference, 2008. ISSCC 2008. Digest of Technical Papers. IEEE International*, pages 322–616. IEEE, 2008.
- [164] Ali Saffari, Mehrdad Hesar, Saman Naderiparizi, and Joshua R Smith. Battery-free wireless video streaming camera system. In *2019 IEEE International Conference on RFID (RFID)*, pages 1–8. IEEE, 2019.
- [165] Hiroki Sakurai, Yuka Kobayashi, Toshiya Mitomo, Osamu Watanabe, and Shoji Otaka. A 1.5 ghz-modulation-range 10ms-modulation-period 180khz rms-frequency-error 26mhz-reference mixed-mode fmcw synthesizer for mm-wave radar application. In *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2011 IEEE International*, pages 292–294. IEEE, 2011.
- [166] Adel AM Saleh and Reinaldo Valenzuela. A statistical model for indoor multipath propagation. *IEEE Journal on selected areas in communications*, 5(2):128–137, 1987.
- [167] Alanson P Sample, Daniel J Yeager, Pauline S Powledge, Alexander V Mamishev, and Joshua R Smith. Design of an rfid-based battery-free programmable sensing platform. *IEEE Transactions on Instrumentation and Measurement*, 57(11):2608–2615, 2008.

- [168] Rahul Sarpeshkar. Analog versus digital: extrapolating from electronics to neurobiology. *Neural computation*, 10(7):1601–1638, 1998.
- [169] Michael D Scott, Bernhard E Boser, and Kristofer SJ Pister. An ultra–low power adc for distributed sensor networks. In *Solid-State Circuits Conference, 2002. ESSCIRC 2002. Proceedings of the 28th European*, pages 255–258. IEEE, 2002.
- [170] Olivier BA SELLER and Nicolas Sornin. Low power long range transmitter, February 2 2016. US Patent 9,252,834.
- [171] Clayton Shepard, Abeer Javed, and Lin Zhong. Control channel design for many-antenna mu-mimo. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*, pages 578–591. ACM, 2015.
- [172] Clayton Shepard, Hang Yu, Narendra Anand, Erran Li, Thomas Marzetta, Richard Yang, and Lin Zhong. Argos: Practical many-antenna base stations. In *Proceedings of the 18th annual international conference on Mobile computing and networking*, pages 53–64. ACM, 2012.
- [173] Suzanne Smiley. Active RFID vs. Passive RFID: What’s the Difference?, 2016.
- [174] Joshua R Smith, Alanson P Sample, Pauline S Powledge, Sumit Roy, and Alexander Mami-shev. A wirelessly-powered platform for sensing and computation. In *International Conference on Ubiquitous Computing*, pages 495–506. Springer, 2006.
- [175] Nicolas Sornin and Ludovic Champion. Signal concentrator device, October 17 2017. US Patent 9,794,095.
- [176] Jian Sun. Pulse-width modulation. In *Dynamics and control of switched electronic systems*, pages 25–61. Springer, 2012.
- [177] Paul D Sutton, Jorg Lotze, Hicham Lahlou, Suhaib A Fahmy, Keith E Nolan, Baris Ozigul, Thomas W Rondeau, Juanjo Noguera, and Linda E Doyle. Iris: an architecture for cognitive radio networking testbeds. *IEEE communications magazine*, 48(9):114–122, 2010.
- [178] Vamsi Talla, Mehrdad Hesar, Bryce Kellogg, Ali Najafi, Joshua R Smith, and Shyamnath Gollakota. Lora backscatter: Enabling the vision of ubiquitous connectivity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 1(3):105, 2017.
- [179] Vamsi Talla, Mehrdad Hesar, Joshua R Smith, Shyamnath Gollakota, Ali Najafi, and Bryce Kellogg. Backscatter systems, devices, and techniques utilizing css modulation and/or higher order harmonic cancellation, October 20 2020. US Patent 10,812,130.

- [180] Vamsi Talla, Bryce Kellogg, Shyamnath Gollakota, and Joshua R Smith. Battery-free cellphone. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 1(2):25, 2017.
- [181] Vamsi Talla, Bryce Kellogg, Benjamin Ransford, Saman Naderiparizi, Shyamnath Gollakota, and Joshua R Smith. Powering the next billion devices with wi-fi. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*, page 4. ACM, 2015.
- [182] Vamsi Talla and Joshua R Smith. Hybrid analog-digital backscatter: A new approach for battery-free sensing. In *RFID (RFID), 2013 IEEE International Conference on*, pages 74–81. IEEE, 2013.
- [183] Kun Tan, He Liu, Jiansong Zhang, Yongguang Zhang, Ji Fang, and Geoffrey M Voelker. Sora: high-performance software radio using general-purpose multi-core processors. In *6th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 09)*, 2009.
- [184] Fang Tang and Amine Bermak. An 84 pw/frame per pixel current-mode cmos image sensor with energy harvesting capability. *IEEE Sensors Journal*, 12(4):720–726, 2012.
- [185] 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.
- [186] David Tse and Pramod Viswanath. *Fundamentals of wireless communication*. Cambridge university press, 2005.
- [187] Ambuj Varshney, Oliver Harms, Carlos Pérez-Penichet, Christian Rohner, Frederik Hermans, and Thiemo Voigt. Lorea: A backscatter architecture that achieves a long communication range. In *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems*, page 18. ACM, 2017.
- [188] Arseny Vasilyev. *The optoelectronic swept-frequency laser and its applications in ranging, three-dimensional imaging, and coherent beam combining of chirped-seed amplifiers*. PhD thesis, California Institute of Technology, 2013.
- [189] G. Vougioukas, S. N. Daskalakis, and A. Bletsas. Could battery-less scatter radio tags achieve 270-meter range? In *2016 IEEE Wireless Power Transfer Conference (WPTC)*, pages 1–3, May 2016.

- [190] Anran Wang, Vikram Iyer, Vamsi Talla, Joshua R Smith, and Shyamnath Gollakota. {FM} backscatter: Enabling connected cities and smart fabrics. In *14th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 17)*, pages 243–258, 2017.
- [191] Jue Wang, Haitham Hassanieh, Dina Katabi, and Piotr Indyk. Efficient and reliable low-power backscatter networks. In *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*, pages 61–72. ACM, 2012.
- [192] Peng Wei, Wenyi Che, Zhongyu Bi, Chen Wei, Yan Na, Li Qiang, and Min Hao. High-efficiency differential rf front-end for a gen2 rfid tag. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 58(4):189–194, 2011.
- [193] Jeffrey A Weldon, R Sekhar Narayanaswami, Jacques C Rudell, Li Lin, Masanori Otsuka, Sebastien Dedieu, Luns Tee, King-Chun Tsai, Cheol-Woong Lee, and Paul R Gray. A 1.75-ghz highly integrated narrow-band cmos transmitter with harmonic-rejection mixers. *IEEE Journal of Solid-State Circuits*, 36(12):2003–2015, 2001.
- [194] Haoyang Wu, Tao Wang, Zengwen Yuan, Chunyi Peng, Zhiwei Li, Zhaowei Tan, Boyan Ding, Xiaoguang Li, Yuanjie Li, Jun Liu, et al. The tick programmable low-latency sdr system. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*, pages 101–113. ACM, 2017.
- [195] Wanghua Wu, Robert Bogdan Staszewski, and John R Long. A 56.4-to-63.4 ghz multi-rate all-digital fractional-n pll for fmcw radar applications in 65 nm cmos. *IEEE Journal of solid-state circuits*, 49(5):1081–1096, 2014.
- [196] H Yao, Y Liao, AR Lingley, A Afanasiev, I Lähdesmäki, BP Otis, and BA Parviz. A contact lens with integrated telecommunication circuit and sensors for wireless and continuous tear glucose monitoring. *Journal of Micromechanics and Microengineering*, 22(7):075007, 2012.
- [197] Daniel Yeager, Fan Zhang, Azin Zarrasvand, Nicole T George, Thomas Daniel, and Brian P Otis. A 9 μ a, addressable gen2 sensor tag for biosignal acquisition. *IEEE Journal of Solid-State Circuits*, 45(10):2198–2209, 2010.
- [198] Jun Yin, Jun Yi, Man Kay Law, Yunxiao Ling, Man Chiu Lee, Kwok Ping Ng, Bo Gao, Howard C Luong, Amine Bermak, Mansun Chan, et al. A system-on-chip epc gen-2 passive uhf rfid tag with embedded temperature sensor. *IEEE Journal of Solid-State Circuits*, 45(11):2404–2420, 2010.

- [199] Jialiang Zhang, Xinyu Zhang, Pushkar Kulkarni, and Parameswaran Ramanathan. Openmili: a 60 ghz software radio platform with a reconfigurable phased-array antenna. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking*, pages 162–175. ACM, 2016.
- [200] Pengyu Zhang, Dinesh Bharadia, Kiran Joshi, and Sachin Katti. Hitchhike: Practical backscatter using commodity wifi. In *Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems CD-ROM, SenSys '16*, pages 259–271, New York, NY, USA, 2016. ACM.
- [201] Pengyu Zhang, Pan Hu, Vijay Pasikanti, and Deepak Ganesan. Ekhonet: High speed ultra low-power backscatter for next generation sensors. In *Proceedings of the 20th annual international conference on Mobile computing and networking*, pages 557–568. ACM, 2014.
- [202] Pengyu Zhang, Mohammad Rostami, Pan Hu, and Deepak Ganesan. Enabling practical backscatter communication for on-body sensors. In *Proceedings of the ACM SIGCOMM 2016 Conference on SIGCOMM*, 2016.