Low-Power Communication for Environmental Sensing Systems

Zerina Kapetanovic

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

Doctor of Philosophy

University of Washington

2022

Reading Committee:

Joshua R. Smith, Chair

Ranveer Chandra

Payman Arabshahi

Miguel Morales

Adrienne Fairhall

Program Authorized to Offer Degree:
Electrical and Computer Engineering

University of Washington

Abstract

Low-Power Communication for Environmental Sensing Systems

Zerina Kapetanovic

Chair of the Supervisory Committee:
Prof. Joshua R. Smith
Computer Science and Engineering,
Electrical and Computer Engineering

Over the last decade, the Internet of Things (IoT) has been changing the world, from enabling connected electronics, smart homes, to smart agriculture. Today, IoT systems have the potential to make a significant impact when it comes to environmental monitoring, which has become increasingly relevant in the times of the climate change crisis and the need to achieve biodiversity conservation. Imagine being able to use passive wireless communication and deploying battery-less sensors for remote environmental monitoring. This dissertation aims to advance and empower these efforts and presents new methods of low-power wireless communication and sensing systems. First, I introduce FarmBeats, an IoT system for data driven agriculture that solves key challenges related to power, connectivity, and cost. Next, I discuss low-power downlink solutions for ambient backscatter systems. In particular, Wireless Quantization Index Modulation and 'Glaze', which build upon data-hiding techniques to enable downlink communication using existing infrastructure and occupied RF spectrum. Lastly, I present Modulated Johnson noise, a new wireless communication system that uses Johnson noise to enable very low-power wireless communication.

ACKNOWLEDGMENTS

First I wish to express my deep appreciation and gratitude to Joshua R. Smith, my advisor, who has played a key role in shaping who I am as a researcher. Josh is an incredible advisor, who is passionate about research and committed to the success of his students. His curiosity, deep technical knowledge, and the support he provides for all of his students to explore their research interests inspires me. Josh is an excellent example of what to strive for as a researcher, mentor, and an advisor. I am extremely thankful and fortunate to have worked with Josh and learn from him. I hope to have such an impact on my own students one day.

I am also deeply grateful for my mentor, Ranveer Chandra. I first started working with Ranveer as an intern at Microsoft Research and those experiences have significantly shaped my research interests and who I am as a person today. His passion for solving tough problems and having real-world impact is infectious and working with him is both inspiring and extremely fulfilling. I am extremely grateful for his support, invaluable advice, and for always believing in me.

I want to thank Miguel Morales, Adrienne Fairhall, and Payman Arabshahi for providing great discussion and feedback on my research. It was an honor to have them on my thesis committee. It has also been wonderful working with everyone in the Sensor Systems Lab, past and present. It has been a great source of support, mentorship, and collaboration over the last several years.

At Microsoft Research I had a wonderful time working with so many people on the the FarmBeats team, including Deepak Vasisht, Tusher Chakraborty, Bodhi Priyantha, Jonhgo Won, Xinxin Jin, Vasuki Swamy, Amy Kumar, Lopa Kundu, and Talal Ahmad. Some of

my best memories are working on FarmBeats with these amazing people and I will always cherish this.

Finally, and most importantly, I am extremely grateful for the support of my family. They have taught me the value of hard work, perseverance, and sacrifice. They have helped shape me into the person I am today and this thesis would not be possible without them.

PREVIOUSLY PUBLISHED MATERIAL

- Chapter 3 revises a previous publication [87]: Zerina Kapetanovic, Vamsi Talla, Aaron Parks, Jing Qian, Joshua R. Smith. *Wireless Quantization Index Modulation: Enabling Communication Through Existing Signals.* IEEE RFID, 2018.

- Chapter 4 revises a previous publication [88]: Zerina Kapetanovic, Ali Saffari, Ranveer Chandra, Joshua R. Smith. *Glaze: Overlaying Occupied Spectrum with Downlink IoT Transmission.* ACM IMWUT, 2019.

- Chapter 2 revises three previous publication [199, 90, 86]: **(1)** Deepak Vasisht, Zerina Kapetanovic, Jongho Won, Xinxin Jin, Ranveer Chandra, Sudipta Sinha, Ashish Kapoor, Madhusudhan Sudarshan, Sean Stratman. *FarmBeats: An IoT Platform for Data-Driven Agriculture.* USENIX NSDI, 2017. **(2)** Zerina Kapetanovic, Deepak Vasisht, Jongho Won, Ranveer Chandra, Mark Kimball. *Experiences Deploying an Always-On Farm Network.* ACM GetMobile, 2017. **(3)** Zerina Kapetanovic, Ranveer Chandra, Tusher Chakraborty, Andrew Nelson. *FarmBeats: Improving Farm Productivity Using Data-Driven Agriculture.* SIAM News, 2019.

- Chapter 5 revises a previous publication [89]: Zerina Kapetanovic, Miguel Morales, Joshua R. Smith. *Communication by means of Modulated Johnson Noise.* Under Review, 2022.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

Chapter 1
INTRODUCTION

The Internet of Things (IoT) plays a critical role in connecting society to the digital world. Today there are billions of connected devices spanning all industries, from transportation and manufacturing to healthcare. These IoT devices have been shown to provide significant benefits, including increased productivity, reduced costs, and even improving our overall health and wellness. Now, imagine being able to further increase the application domain of IoT systems and use them to solve pressing issues in today's world, particularly when it comes to the climate change crisis.

Sensing devices can be deployed in remote forests to track deforestation and poaching or used to identify detrimental issues causing the population decline of insect pollinators. Even consider the agriculture industry, which is one of the biggest contributors to greenhouse gas emissions. Instead of continuing with traditional farming practices, IoT systems can be used to enable data-driven agriculture, which has shown to increase yield, reduce cost, and ensure sustainability. However, to use IoT for such scenarios requires us to deal with major challenges focused on resource constraints and scale. In this work, we present new approaches to wireless communication and sensing to solve the aforementioned challenges and help achieve the vision of ubiquitous computing.

1.1 Overview of Challenges

Enabling wireless sensing systems for resource-constrained applications requires us to address major challenges focused on power, connectivity, and cost, which are all interconnected. Devices often need to be deployed in the wild for long periods of time and in locations that lack power and Internet connectivity. This becomes even more challenging when considering

the vast size of forests, farms, and oceans, all requiring large-scale deployments in order to be monitored.

- **Power:** In a wireless sensor network, the endpoint devices are often power constrained. Typically the device is powered by a battery that requires routine replacement. In many IoT applications (e.g., homes) this is not a bottleneck. However, it does become an issue when deploying such devices in remote locations at large-scale. In these scenarios, devices are not easily accessible and need to be deployed long-term (e.g., years) to collect various forms of data ranging from temperature and CO_2 measurements to capturing images. It is critical to minimize the overall power consumption of the endpoint devices in order to increase the lifespan of the sensing system and enable long-term data collection.

- **Connectivity:** In many sensing scenarios, there is little to no Internet connectivity. For example, monitoring a farm field in a rural area or wildlife monitoring in remote forests. There are many connectivity solutions that are used to enable IoT systems, such as LoRa, Wi-Fi, Cellular, or BLE. Still, they are often unsuitable for large-scale and remote deployments. These systems require long-range connectivity, minimal infrastructure, and low power consumption. Consider, LoRa, which is popular IoT connectivity solution. With LoRa, we can achieve roughly 3 miles of communication range between the base station and endpoint devices [171]. However, in real-world deployments this is not the case due to many obstructions (e.g., building, tree or crop canopy). Moreover, imagine deploying devices to monitor an area that spans 10s of thousands of acres. This would require 100s of sensing devices and 10s of base station, which quickly scales cost and maintenance.

- **Cost:** Both power and connectivity play a critical role in the overall cost of a system. Long-rage connectivity is needed to minimize the number of base stations required to support a sensor network and low power consumption is needed to minimize overall

maintenance. Moreover, designing cost-effective wireless sensing systems is needed in order to enable widespread adoption of the technology.

Over the last decade, there have been several advances in very low-power wireless communication and sensing techniques that take important strides to realizing the vision of ubiquitous computing. For example, backscatter communication is ultra-low-power and enables devices to be battery-free. Devices can leverage existing RF signals to for both power and communication. This has led to several promising applications of battery-free sensing, including battery-free cameras, phones, and even underwater sensing [110, 131, 82, 189]. However, a major practical limitation with such solution is dependence on a generated or ambient RF source. Commercial RFID tags are one example, which relies on an RFID reader for power and communication at a relatively short range. Similarly, ambient backscatter uses existing RF signals which limits the use cases of the solution. For instance, in remote locations where there are no TV or FM broadcast signals that that devices could utilize for backscatter communication. Lastly, the majority of the proposed ambient backscatter systems only enable peer-to-peer or uplink communication. They lack solution to enable downlink communication as well. In this work, I propose new approaches to wireless communication and sensing that tackle the aforementioned challenges and new designs for end-to-end IoT systems targeted for large-scale environmental monitoring.

1.2 Designing an end-to-end IoT system for resource constrained environments

A major step forward would be to enable an end-to-end IoT system in remote and resource constrained environments that addresses the aforementioned challenges. In Chapter 2 I present our first step towards this by developing FarmBeats, an IoT platform for data-driven agriculture. FarmBeats enables seamless data collection from various sensor types (e.g soil sensors, cameras, drones) and is able to ensure system availability even in harsh weather conditions. For instance, when faced with power or Internet outages, scenarios that are common in farming areas. FarmBeats solves tough challenges such as access to power

and connectivity, automating long-term data collection, and demonstrating the feasibility of long-term deployments in the wild.

1.3 Downlink Communication for Ambient Backscatter Systems

Environmental monitoring extends to urban areas as well, such as monitoring pollution in cities. In these scenarios, it can be challenging to deploy large-scale IoT systems simply due to maintenance. Most, if not all, mainstream IoT solutions require endpoint devices to rely on batteries. In these scenarios, 1000s of devices would require routine maintenance for battery replacement, becoming costly and unsustainable over time. Ambient backscatter communication (ABC) is a promising solution because it operates at ultra-low power allowing devices to be battery-less [4]. ABC uses existing RF signals to encode bits of information by reflecting or absorbing the ambient signal. This enables peer-to-peer and uplink communication at ultra low-power. However, to have an end-to-end backscatter IoT system, downlink communication is necessary.

First, in Chapter 3, I present Wireless Quantization Index Modulation (QIM), a technique that uses existing infrastructure to embed information into existing wireless signals to communicate with IoT devices with negligible impact on the original signal and zero spectrum overhead. We explore the design space for wireless QIM and evaluate the performance of embedding information in TV, FM and AM radio broadcast signals under different conditions. We demonstrate that we can embed messages at up to 8–200 kbps with negligible impact on the audio and video quality of the original FM, AM and TV signals, respectively.

In Chapter 4, we build upon Wireless QIM and present Glaze, a new system that utilizes occupied spectrum to create a new channel for downlink communication. In particular, Glaze uses a new technique that introduces small perturbations in the form of attenuation to existing signals to convey data. We design transmit and receive hardware to evaluate the performance of Glaze and show how it can be used across wireless standards such as FM, TV, or Wi-Fi to communicate with devices with minimal impact on existing data transmissions.

1.4 Wireless Communication using Thermal Noise

Each of the aforementioned solution take important strides in enabling IoT systems for environmental monitoring, but there are still several challenges that remain. For example, how do we maintain large-scale deployments in forested areas? Tree canopy makes solar power limited, routine battery replacement at large-scale is both costly and unmanageable, and enabling connectivity is difficult because of natural obstructions. Battery-free sensing solutions are promising for these scenarios, but backscatter communication has a practical limitation due to its dependence on existing or generated RF signal sources. In Chapter 5, we present the design of a new passive wireless communication method that does not rely on ambient or generated RF sources. Instead, we exploit the Johnson (thermal) noise of a resistor to transmit information bits wirelessly. By switching the load connected to an antenna between a resistor and an open or short circuit, we can achieve data rates of up to 26bps and distances of up to 7.3 meters. This communication method is orders of magnitude less power consuming than conventional communication schemes and presents the opportunity to enable wireless communication for power and connectivity-constrained applications.

1.5 Contributions and Organization of Thesis

In this thesis, we present new low-power wireless communication and sensing solutions for energy-constrained platforms. Additionally, we design and implement end-to-end IoT systems that tackle key challenges related to resource-constrained environments. The outline of the thesis is as follows:

- Chapter 2 presents FarmBeats, an end-to-end IoT system for data-driven agriculture. FarmBeats enables seamless data collection from various ground sensors, cameras, and drones. The system solves key challenges related to power, connectivity, and cost in order to enable data-driven agriculture techniques. FarmBeats has been shown to boost

agricultural productivity and has been deployed on numerous farms across the United States.

- Chapter 3 introduces Wireless Quantization Index Modulation, a technique that uses existing infrastructure to embed information into existing wireless signals to communicate with IoT devices, while having negligible impact on the original signal and zero spectrum overhead. We explore the design space for Wireless QIM and evaluate the performance of embedding information in TV, FM, and AM radio broadcast signals under different conditions. We demonstrate that we can embed messages at up to 8-200 kbps with minimal impact on the audio and video quality of the original FM, AM, and TV signals.

- Chapter 4 presents Glaze, which builds upon Wireless QIM. Glaze is a system that overlays data on existing wireless signals to create a new channel for downlink communication for IoT backscatter devices. In particular, Glaze uses a new technique that introduces small perturbations to existing signals to convey data. We describe the hardware design, evaluate it's performance, and show how it can be used across wireless standards such as FM, TV, or Wi-Fi to communication devices with minimal impact on existing data transmissions.

- Chapter 5 presents Communication by means of Modulated Johnson noise, the first experimental realization of a new, ultra-low power communication method that works by selectively connecting or disconnecting an impedance matched resistor and an antenna. This modulates microwave frequency Johnson noise emitted by the antenna. The data transmission hardware is similar to that of an RFID tag, which communicates by reflecting RF signals; the crucial advantage of the present system is that it requires no pre-existing RF signal. An interesting feature of the system is that all components of the system are at the same physical temperature, but it functions because they have different noise temperatures. The method may also have practical utility:

the elimination of the RF carrier simplifies the system architecture and the reader hardware.

- Chapter 6 concludes the thesis and discusses some future areas of research.

Chapter 2

FARMBEATS: AN AI AND IOT SYSTEM FOR DATA-DRIVEN AGRICULTURE

2.1 Introduction

The demand for food is expected to double by 2050, primarily fueled by an increase in population and upward social mobility[196]. Achieving this increase in food production is even more challenging because of receding water levels, climate change and shrinking amount of arable land. According to International Food Policy Research Institute, data-driven techniques can help us achieve this goal by increasing farm productivity by as much as 67% by 2050 and cutting down agricultural losses [64].

In fact, field trials have shown that techniques that use sensor measurements to vary water input across the farm at a fine granularity (precision irrigation) can increase farm productivity by as much as 45% while reducing the water intake by 35%[7]. Similar techniques to vary other farm inputs like seeds, soil nutrients, etc. have proven to be beneficial [95, 128]. More recently, the advent of aerial imagery systems, such as drones, has enabled farmers to get richer sensor data from the farms. Drones can help farmers map their fields, monitor crop canopy remotely and check for anomalies. Over time, all this data can indicate useful practices in farms and make suggestions based on previous crop cycles; resulting in higher yields, lower inputs and less environmental impact.

While these techniques for agriculture have shown promising results, their adoption is limited to less than 20 percent farmers owing to the high cost of manual sensor data collection (according to US Department of Agriculture [113]). Automating sensor data collection requires establishing network connection to these sensors. However, existing connectivity solutions [40, 58] require a cellular data logger to be attached to each sensor (see Table 2.1

for a detailed comparison). These loggers cost around $1000 each in equipment cost with over $100 in subscription fee. Further, they are limited in the amount of data that they can send to few kilobytes per day. Clearly, these solutions do not scale up for large farms and cannot support high bandwidth sensors like cameras and drones, which rely on sending all their data to the cloud for processing[173, 38]. This situation is further worsened by the fact that farms typically have limited cellular coverage [92] and are prone to weather-based Internet outages.

In this paper, we present FarmBeats, an end-to-end IoT platform for data-driven agriculture, that enables seamless data collection from various sensor types, i.e., cameras, drones and soil sensors, with very different bandwidth constraints. FarmBeats can ensure system availability even in the face of power and Internet outages caused by bad weather; scenarios that are fairly common for a farm. Further, FarmBeats enables cloud connectivity for the sensor data to enable persistent storage as well as long-term or cross-farm analytics. We have deployed FarmBeats in two farms in the US over a period of six months and used FarmBeats to enable three applications for the farmer: precision agriculture, monitoring temperature and humidity in food storage, and monitoring animal shelters. In designing Proteus, we solve three key challenges.

First, to enable connectivity within the farm, FarmBeats leverages recent work in unlicensed TV White Spaces (TVWS) [56, 13, 154] to setup a high bandwidth link from the farmer's home Internet connection to an IoT base station on the farm. Sensors, cameras and drones can connect to this base station over a Wi-Fi front-end. This ensures high bandwidth connectivity within the farm. However, due to the lack of power on the farm, the base station is powered by battery-backed solar power which suffers from power unreliability depending on weather conditions. As shown in past work [75, 176], cloudy weather can reduce solar power output significantly and drain the batteries of the base station to shut it down. To solve this problem, FarmBeats uses a novel weather-aware IoT base station design. Specifically, it uses weather forecasts to appropriately duty cycle different components of the base station. To the best of our knowledge, this is the first weather-aware IoT base station design.

Second, Internet connection to the farm is typically weak making it challenging to ship high bandwidth drone videos (multiple GBs) to the cloud. Furthermore, farms are prone to weather-related network outages that last weeks. Such system unavailability impedes a farmer's ability to take adequate preventive actions, do UAV inspections and leads to loss of valuable sensor data. Thus, FarmBeats uses a Gateway based design, wherein a PC at the farmer's home serves as a gateway for the farm data. The FarmBeats Gateway serves two purposes: a) it performs significant computation locally on the farm data to consolidate it into summaries that can be shipped to the cloud for long-term and cross-farm analytics, and b) the gateway is capable of independent operation to handle periods of network outage, thus leading to continuous availability for the farmer.

Finally, while drones are one of the most exciting farm sensors today, they suffer from poor battery life. Getting aerial imagery for a farm requires multiple drone flights and a long wait time in between when the batteries are being charged. We use the fact that farms are typically very windy, since they are open spaces. Thus, we incorporate a novel path planning algorithm in the FarmBeats gateway, that leverages wind to help the drone accelerate and decelerate, thereby conserving battery. This algorithm is motivated by how sailors use winds to navigate sailboats.

We use the FarmBeats system to enable precision agriculture applications on two farms: one in Washington state and the other in upstate New York. While traditional farming treats the farm as a homogeneous piece of land, precision agriculture adapts the farm inputs over different parts of the farm depending on the requirement. Precision agriculture techniques require a precision map with information about each location in the farm, for example, the soil temperature, soil moisture, nutrient levels, etc. To construct this precision map, existing solutions for precision agriculture require a dense deployment of in-ground sensors [113]. A dense deployment of sensors becomes expensive (as well as cumbersome to manage) as the size of the farm grows. Unless these sensors are deployed densely within a farm, the estimated precision map can be very inaccurate, as we show in Section 3.5. Since FarmBeats' gateway has access to both the drone videos and sensor data, it enables a novel low-cost mechanism

that uses drone videos in combination with sparse ground sensors to generate precision maps for the farm. To the best of our knowledge, this is the first system that can combine the temporal data from sensors, with the spatial data from drones to construct an instantaneous precision map of the farm, such as the one in Figure 2.4.

Beyond FarmBeats' application in precision agriculture, farmers have so far used FarmBeats for two other applications. First, the farmers have been using FarmBeats to monitor temperature and humidity in storage spaces to ensure that the produce does not go bad. Second, the farmers have plugged in cameras at different locations, to monitor cow sheds, selling stations etc[1].

Contributions: To summarize, FarmBeats makes the following key contributions:

- **Long-term large scale deployment:** Our deployments have run over 6 months in each of the farms and collected over 10 million sensor measurements, 1 million camera images and 100 drone videos
- **Novel Weather-Aware IoT Base Station Design:** Adding weather awareness into the IoT base station reduced the base station down time to zero as opposed to greater than 30% downtime during the same month in the previous year in an earlier version of our deployment
- **Novel Inference Techniques for Compression of Aerial Imagery Data:** FarmBeats' gateway achieved a median compression of 1000 times from an aerial drone video to the sensor summaries sent to the cloud. Further, the gateway remained available even when the Internet connectivity to the farm faced a week-long outage
- **Wind-Assisted Drone Flight Planning Algorithm:** FarmBeats' flight planning algorithm improves the area covered by a single drone flight by 30%

2.2 IoT Platform: Objectives

In building FarmBeats, we target the following goals:

[1]Supplementary Material includes detailed description of FarmBeats applications and usage.

- **Availability:** The platform should have negligible downtime. When there is an outage (for example, due to power or network failure), data collection from the sensors should not stop and the platform should continue to deliver services to the farmers.

- **Capacity:** It should support sensors with widely varying requirements: pH sensors reporting few bytes of data to drones sending gigabytes of video. Similarly, the system should be capable of supporting end-user applications with varying needs: from a precision irrigation application that needs the latest sensor data for the entire farm to a crop suggestion application that needs just high level productivity data but across several growing seasons

- **Cloud Connectivity:** Several farming applications, such as crop cycle prediction, seeding suggestions, farming practice advisory, etc. rely on long term data analytics. Besides, a farmer may want to access some applications even when he is not on the farm. Thus, the IoT platform must enable pushing data to the cloud.

- **Data Freshness:** Stale sensor data from the farm can make applications suggest incorrect courses of action to the farmer. Gaps in historical data can also cause applications to misbehave. Moreover, stale data leads to bad user experience. Thus, the platform must strive to maintain maximum data freshness.

2.3 The FarmBeats IoT Platform

While these objectives have been fairly successfully achieved by home IoT platforms like Amazon Echo, achieving these objectives in an agricultural setting introduces several challenges for two main reasons: access and environmental variability. As discussed before (and as shown in Table 2.1), farms do not have access to power and high-bandwidth Internet connectivity unlike indoor IoT systems. Furthermore, energy harvested from the environment and weak network connectivity to the farm is susceptible to failures due to weather

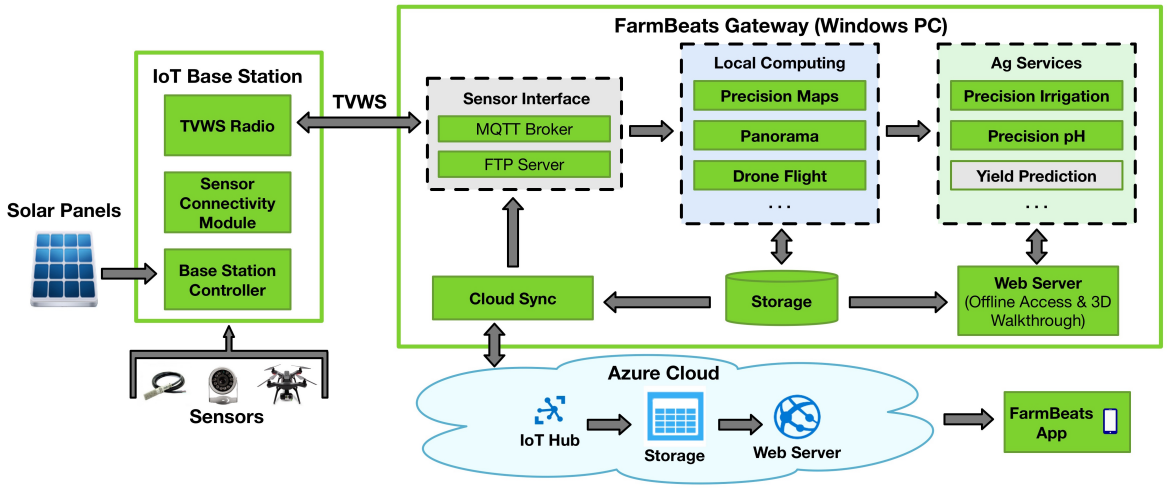

Figure 2.1: **FarmBeats System Overview**

Technology	Cost	Data Restriction
Cellular Connection (Decagon Devices)	Per Sensor: 1000$ + 100$ annual fee	Restricted to sensor data; Uploads every 15 mins at best
Mesh Networks (Ranch Systems)	Base station: 3500$ + 750$ annual fee; Per Sensor: 1100$ + 60$ annual fee	Maximum 25 mesh nodes per base station
Satellite (Iridium)	Per Sensor: 800$ + 100$ monthly fee	Restricted to 2.5 Kbps

Table 2.1: **Cost Comparison of Farm Sensor Networking Solutions**

variability. So, the key question for the design of FarmBeats is: how does one design an IoT platform to meet the objectives in a highly variable, resource constrained environment?

2.3.1 Design Decisions

An overview of the system is given in Figure 2.1. Here, we discuss the main design decisions.

To achieve farm connectivity over long range, we leverage recent work in the TV White Spaces [13, 154, 56] to setup a high-bandwidth connection from the farmer's home to the farm. However, sensors, drones and cameras typically do not support TVWS. Thus, in order to maintain compatibility with sensors along with long-range high bandwidth connectivity, we deploy a two-layer hybrid network. We use a TVWS link to connect the farmer's home Internet connection to a few IoT base stations on the farm. Since it is a high bandwidth backhaul link, each base station can accommodate sensors, as well as cameras and drones.

At the second layer, the IoT base station provides a Wi-Fi interface for connections from sensors and other devices. The Wi-Fi interface ensures that the farmer can not only connect most off-the-shelf farming sensors, cameras and drones; but they can also use their phone to access farming productivity apps. [2]

Variability in harvested solar energy leads to IoT base station downtime in overcast conditions. In fact, in our early deployments, power failures due to environmental factors were the major cause of unavailability. While past work has dealt with this problem in the context of single sensors [75, 200, 176] by duty cycling the sensors, the same approach does not work for a base station. Specifically, the base station has multiple components with different power requirements and duty cycling costs. For example, a farmer is typically inactive at night and is unlikely to check the farm data. So, turning the TVWS device off (which consumes 5x more power than the rest of the base station) can enable the base station to collect data (in a cache) from the sensors more frequently. Further, FarmBeats enables the farmer to turn the base station on to access Wi-Fi for productivity applications, while they are on the farm. This adds another layer of uncertainty in the duty cycling plan. Thus, we propose a novel duty cycle policy (in Section 2.4) wherein the different components of the base station are duty cycled at different rates; while explicitly accommodating these constraints.

Finally, given the weak internet connectivity to the farm, a naive approach of pushing all the data to the cloud does not work. We make the key observation that the data requirements of the farming applications can be broadly classified into two main categories: immediate detailed data and long-term summarized data. Table 2.2 summarizes how the industrial and research applications of farm data can be classified into these two categories. This categorization enables a gateway based IoT design for Proteus. The local gateway sits at the farmer's home at the other end of the White Space link and performs two functions: a) creates summaries for future use and ships them to the cloud and b) delivers applications

[2]Future iterations of the systems would add multiple interfaces to the base station to enable compatibility with more sensor types.

Data Requirement	Applications
Immediate Descriptive Data	Precision irrigation, virtual walkthroughs, productivity apps, farm monitoring, ...
Long-term Summarized Data	Crop suggestions, seed distribution, yield monitoring, financial management, animal health statistics, ...

Table 2.2: **Application classification based on requirements**

that can be provided locally. The summaries are several orders of magnitude lower in size than the raw farm data (3-4 orders of magnitude smaller in case of the precision agriculture application discussed later) and hence, respect the harsh bandwidth constraints.

2.3.2 Architecture

The FarmBeats system has the following components:

Sensors & Drones: FarmBeats uses off-the-shelf sensors for its applications. Each sensor measures specific characteristics of the farm, such as soil moisture and soil pH, and reports this data to the IoT base station over a Wi-Fi connection. In addition to soil sensors, FarmBeats supports cameras for farm monitoring and drones. The cameras are either connected to the IoT base station over Ethernet or report data over Wi-Fi. They take periodic snapshots and transmit this data to the IoT base station. UAV flights are either periodically scheduled or manually initiated using the FarmBeats app on the farmer's phone.

IoT Base Station: The IoT base station on the farm is powered by solar panels, backed by batteries and has three components:

- The TVWS device ensures that the base station on the farm can send the data to the gateway, which then, sends it up to the cloud.

- The sensor connectivity module establishes a connection between the base station and the sensors deployed on the farm. In FarmBeats' current implementation, this module is just a Wi-Fi router.

- Finally, the Base Station Controller is responsible for two functions. First, it serves as a cache for the sensor data collected by the sensor module and syncs this data with the IoT gateway when the TVWS device is switched on. Second, it plans and enforces the duty cycle rates depending on the current battery status and weather conditions.

IoT Gateway: As mentioned before, the goal of the IoT gateway is to enable local services and create summaries from existing data to be sent to the cloud. We use a PC form factor device as the FarmBeats gateway, which is typically placed in the farmer's house or office, whichever has Internet access. The gateway provides an interface for applications to run and create summaries to be sent to the cloud as well as to post data to the local web server. Furthermore, it includes a web service for the farmer to access detailed data when they are on the farm network. This also ensures that FarmBeats remains available even when the cloud connection is not present. Finally, it includes built-in algorithms for drone path planning and for compressing drone data before being sent to the cloud (described in Section 2.5). We illustrate in Section 2.5.3 how applications function on the gateway with the example of precision agriculture applications.

Three aspects of the FarmBeats gateway differentiate it from prior IoT gateways. First, the FarmBeats gateway implements a web service, providing unique services that are different from the FarmBeats web service in the cloud. Second, the gateway can operate offline, and still offer the most important services. Finally, as shown later in the context of precision agriculture, having access to data from multiple types of sensors enables unique feature-based summarization technologies for the drone videos and sensor data.

Services & the Cloud: The Gateway ships data summaries to the cloud, which provides a storage system for long-term data and a web interface for the farmer. The cloud enables three functions: data access outside the farm network (e.g. when traveling), long term applications like crop suggestions, and cross-farm analytics.

2.3.3 Architecture

The FarmBeats system has the following components:

Sensors & Drones: FarmBeats uses off-the-shelf sensors for its applications. Each sensor measures specific characteristics of the farm, such as soil moisture and soil pH, and reports this data to the IoT base station over a Wi-Fi connection. In addition to soil sensors, FarmBeats supports cameras for farm monitoring and drones. The cameras are either connected to the IoT base station over Ethernet or report data over Wi-Fi. They take periodic snapshots and transmit this data to the IoT base station. UAV flights are either periodically scheduled or manually initiated using the FarmBeats app on the farmer's phone.

IoT Base Station: The IoT base station on the farm is powered by solar panels, backed by batteries and has three components:

- The TVWS device ensures that the base station on the farm can send the data to the gateway, which then, sends it up to the cloud.

- The sensor connectivity module establishes a connection between the base station and the sensors deployed on the farm. In FarmBeats' current implementation, this module is just a Wi-Fi router.

- Finally, the Base Station Controller is responsible for two functions. First, it serves as a cache for the sensor data collected by the sensor module and syncs this data with the IoT gateway when the TVWS device is switched on. Second, it plans and enforces the duty cycle rates depending on the current battery status and weather conditions.

IoT Gateway: As mentioned before, the goal of the IoT gateway is to enable local services and create summaries from existing data to be sent to the cloud. We use a PC form factor device as the FarmBeats gateway, which is typically placed in the farmer's house or office, whichever has Internet access. The gateway provides an interface for applications to run and create summaries to be sent to the cloud as well as to post data to the local web server. Furthermore, it includes a web service for the farmer to access detailed data when they are on the farm network. This also ensures that FarmBeats remains available even when the cloud connection is not present. Finally, it includes built-in algorithms for drone path planning

and for compressing drone data before being sent to the cloud (described in Section 2.5). We illustrate in Section 2.5.3 how applications function on the gateway with the example of precision agriculture applications.

Three aspects of the FarmBeats gateway differentiate it from prior IoT gateways. First, the FarmBeats gateway implements a web service, providing unique services that are different from the FarmBeats web service in the cloud. Second, the gateway can operate offline, and still offer the most important services. Finally, as shown later in the context of precision agriculture, having access to data from multiple types of sensors enables unique feature-based summarization technologies for the drone videos and sensor data.

Services & the Cloud: The Gateway ships data summaries to the cloud, which provides a storage system for long-term data and a web interface for the farmer. The cloud enables three functions: data access outside the farm network (e.g. when traveling), long term applications like crop suggestions, and cross-farm analytics.

2.4 Duty Cycling the Base Station

As discussed before, FarmBeats' solar-powered IoT base station on the farm is duty cycled to explicitly account for weather forecasts and current charge state of the batteries. Two aspects of the base station make this problem challenging: a) The sensor connectivity module has significantly lower power requirements than the TVWS device. Thus, we need to intelligently proportion power between these components to achieve optimum performance. b) FarmBeats allows farmers to manually turn the base station on to connect to the Internet to use productivity apps on their phone. This adds a variable component to power consumption.

2.4.1 Duty Cycling Goals

The key goals for the duty cycling algorithm are:

- **Energy Neutrality:** Like past work in the context of duty-cycling sensors backed by

energy harvesting sensor systems [75, 200], FarmBeats aims to achieve the objective of energy neutrality. For a given planning period, the goal is to consume at max as much power as can be harvested from the solar panels.

- **Variable Access:** FarmBeats allows farmers to access Wi-Fi connectivity on-demand. This power consumption is usage-driven and varies across days. FarmBeats must plan ahead for this variable delay.

- **Minimize Data Gaps:** We use the term 'Data Gaps' to denote continuous time-intervals with no sensor measurements available. Such gaps need to be minimized to avoid missing out on interesting data trends. So, FarmBeats' duty cycling algorithm aims to minimize the length of the largest data gaps, under the constraints of energy neutrality and variable access.

2.4.2 Power Budget

The sole power source for the base station is a set of solar panels (backed by a battery). The solar power output varies with the time of day and the weather conditions. We use standard methods [176] to estimate the output of the solar panels, given the weather conditions. Let us say that the energy output from the solar panels over the next planning period is S_I. Because the estimation is not perfect and there is usage variability, there maybe some credit or debit from the previous planning period. Let us denote this credit by C_I. So, the total power budget for the base station over the next planning period is $S_I + C_I$.

2.4.3 Duty Cycling Approach

The duty cycle decisions are made on the order of a planning period, T_p. Since our deployments use solar powered base stations, we set T_p to be one day. We define the average energy loss due to battery leakage and the very low power base station controller during one T_p to be E_D. For the farmer to have on-demand Wi-Fi access, we allocate a fixed time budget of T_v.

If we denote the power consumption of the TVWS device by P_T and the power consumption of the sensor connectivity module by P_S, then, we need to allocate $T_v(P_T + P_S)$ for variable Wi-Fi access. Now the key question is, how do we proportion the remaining power budget?

Duty Cycling the TVWS device: The TVWS module is needed to sync the data in the base station cache with the gateway. Let us assume that we have a schedule, S, the set of sync times advised for the base-station to sync with the FarmBeats gateway. This could depend on the farmer's usage patterns, sensor types and can be either manually programmed or automatically inferred. The sync times in the set S have a corresponding set of weights given by set W. An example of a high-weighted sync time could be sunrise, as that is when the farmer begins their day. Thus, they would like to access the latest sensor data when the activities of the day are planned.

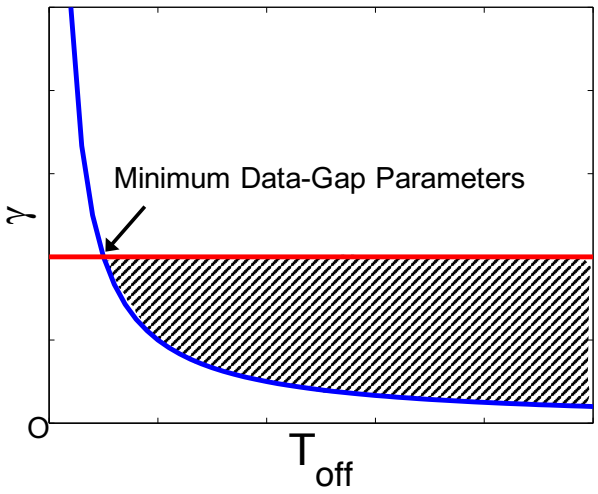

Figure 2.2: **Duty Cycling Approach:** The shaded region shows the feasibility region of Equations 2.2 and 2.4. The latency is minimized when both the equations are satisfied on the boundaries.

To ascertain the subset of syncs that need to be performed, we make a simple observation. If the sensors haven't sent any data to the base station, the base station need not turn on the TVWS device. Specifically, it uses the following greedy algorithm to identify the syncs to be executed. Let us denote by, $S_1 \subset S$, the subset of syncs that are to be executed. This subset is initialized as an empty set. FarmBeats starts by adding the highest priority sync to S_1. After it has done that, it subtracts $|S_1|P_T T_S$ from the power budget, where $|.|$ denotes set cardinality and T_S denotes the time to perform a sync operation. Then, FarmBeats computes the corresponding duty-cycle rate for the sensor connectivity module. If this rate ensures that the second highest weighted

sync in S will have additional data from the sensors to sync with the gateway, it adds this sync operation to the set S_1. It repeats this process in decreasing order of weights until it reaches a state where one of the syncs in S_1 has no new data to share. As we add more sync operations to S_1, the power budget for the sensor connectivity modules decreases. With a lower power budget, the sensor connectivity module can collect data from the sensors less often and hence it becomes less likely for frequent syncs to see new data. Thus, the algorithm implicitly regulates the sync operations between the gateway and the base station.

Duty Cycling the Sensor Connectivity Module: We denote the duty cycling rate for the sensor connectivity module by γ. In particular, it is turned off for a time period, T_{off}, followed by an on period of, T_{on} and $\gamma = T_{on}/T_{off}$. Using the notation we have established so far, the energy expenditure of the system is $E_D + (P_S + P_T)T_v + P_T T_S|S_1| + P_S T_p \gamma$. Since the goal of the planning algorithm is to estimate T_{on} and T_{off} such that the energy expenditure does not exceed the energy budget during the planning period, this imposes the following constraint:

$$S_I + C_I \geq E_D + (P_S + P_T)T_v + P_T T_S|S_1| + P_S T_p \gamma \tag{2.1}$$

$$\implies \gamma \leq \frac{S_I + C_I - E_D - (P_S + P_T)T_v - P_T T_S|S_1|}{P_S T_p} \tag{2.2}$$

Let us denote $T_{connect}$ as the time taken for the sensor connectivity module to turn on and establish a connection to the sensors. Further, let T_{sensor} be the time that it takes for all the sensors to wake up and transmit to the base station. Since the ON time of the module has to be long enough for the sensors to be able to communicate their data to the base station, this imposes a further constraint:

$$T_{ON} \geq T_{connect} + T_{transfer} \tag{2.3}$$

$$\implies \gamma T_{off} \geq T_{connect} + T_{transfer} \tag{2.4}$$

Since our goal is to minimize the data gap under the power constraints imposed by Equations 2.2 and 2.4, we aim to minimize T_{off}. The inequalities from Equations 2.2 and 2.4 define a convex region in the 2-dimensional space of (γ, T_{off}), shown as the shaded region in Figure 2.2. Since the cost function T_{off} is linear, the minimum occurs on a corner of the

intersection region defined by the two inequalities. Specifically, the minimum latency is achieved when the two inequalities are exactly met. The solution is shown graphically in Figure 2.2.

2.4.4 Discussion

At this point, it is worth noting that:

- By explicitly accounting for the credit term, C_I, the formulation absorbs the variability in on-demand Wi-Fi usage patterns. If the on-demand Wi-Fi usage patterns are stable, the term C_I goes down to zero.

- By incorporating flexibility in sync times between the gateway and the base station, FarmBeats can easily adapt to farm applications with different requirements.

- We have not yet discussed the duty cycling of sensor nodes. In our implementation, we set the duty cycle off time for sensors to be less than $T_{transfer}$ to ensure that the sensor can transfer data when the sensor connectivity module is on. An alternative implementation would allow the base station to send wake-up times to sensors. Our design choice was motivated by the availability of very low-power sensors that consume 3-4 orders of magnitude less power than the base station on average.

2.5 The FarmBeats Gateway

In this section, we discuss two key components of the FarmBeats gateway: UAV path planning and stitched imagery (orthomosaic) generation from UAV videos. We also illustrate how the FarmBeats gateway implements a precision agriculture application.

2.5.1 UAV Path Planning

Most UAVs operate in line sweep patterns. Specifically, given a sequence of waypoints defined by their GPS coordinates, they move from one waypoint to the next, in order. However, in

the context of agriculture, our objective is to optimize for the area covered in a single flight. Thus, we aim to minimize the time taken to cover a given area. To that end, we make the observation that increasing the number of waypoints to cover the same area increases the time taken to cover it, even though the total path length may be the same. This is because the quadrotor has to decelerate at each waypoint and come to a halt before it can turn around and accelerate again. We present a novel flight planning algorithm that minimizes the number of waypoints required to cover a given part of the farm.

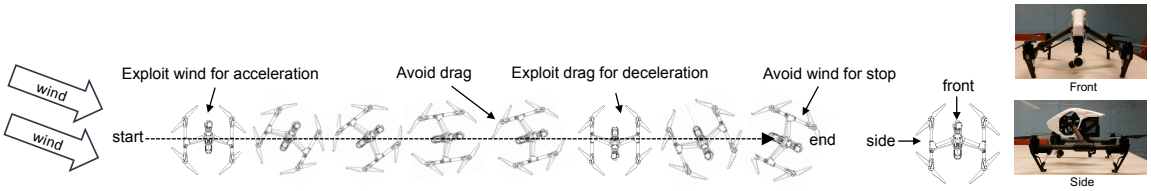

(a) (b)

Figure 2.3: FarmBeats' path planning algorithm uses the asymmetry in front and side profiles of a drone like DJI Inspire 1 (in (b)) to leverage wind to its advantage

Existing commercial systems like Pix4D [141], DroneDeploy [50], etc. offer area coverage services, these systems cover a given area using an east-to-west flight path, without any regards to the number of waypoints required. Recent research proposals like [57] do not guarantee the minimum number of waypoints either. Our area coverage algorithm *Min-waypoint* described below guarantees that the UAV covers an area with the minimum number of waypoints:

- Given an area, construct its convex hull.
- Determine the direction of sweeping lines. For each edge and its antipodal (diametrically opposite) vertex [175], draw two parallel lines and measures the distance between them. The slope of the edge corresponding to the minimum distance between the edge and the antipodal vertex becomes the direction of the sweeping lines.
- Determine the waypoints depending on the flight altitude, the camera's field of view,

and desired image quality.

- Given a start-point and end-point of the flight path, order the waypoints to minimize the total travel distance.

Adaptive Wind-assisted Yaw Control: Since farms are large open spaces and typically very windy, we observed that quadrotors that have an asymmetric physical profile can exploit the wind either for more efficient propulsion or deceleration. Figure 2.3(b) shows an example of a quadrotor (DJI Inspire 1) that has an asymmetrical profile, where its front and the side are considerably different; thus, it can exploit the wind similar to sailboats. Intuitively, when the quadrotor is flying downwind (i.e. wind is helping the quadrotor), the side profile of the quadrotor should face the wind since the side profile has a larger area and hence, will be able to extract the maximum assistance from the wind. In our experiments on the farm, the quadrotor requires significantly more energy (80% higher at 4m/s) to maintain its speed upwind in comparison to the downwind flight.

To leverage this observation, we designed a novel yaw control algorithm to exploit the wind energy on the farm. Specifically, yaw is the angle of the quadrotor with respect to the vertical axis. While we don't describe the algorithm in detail, on a high level, Figure 2.3 describes how the yaw control algorithm would operate for a quadrotor that has a larger area on the sideways profile. For the downwind segment from the start point to the first waypoint, the adaptive control starts by making the yaw perpendicular to the flight path, thereby maximally utilizing the favorable wind as the quadrotor accelerates. However, as the velocity increases, the air drag generated by the quadrotors profile also increases. Consequently, once the quadrotor accelerates the yaw is reduced so as to maximally exploit the wind, while minimizing the parasitic drag due to the side profile. Similarly, the deceleration phase can very effectively exploit the air drag by making its yaw perpendicular to the flight path. This action is analogous to the action that a skier takes to stop.

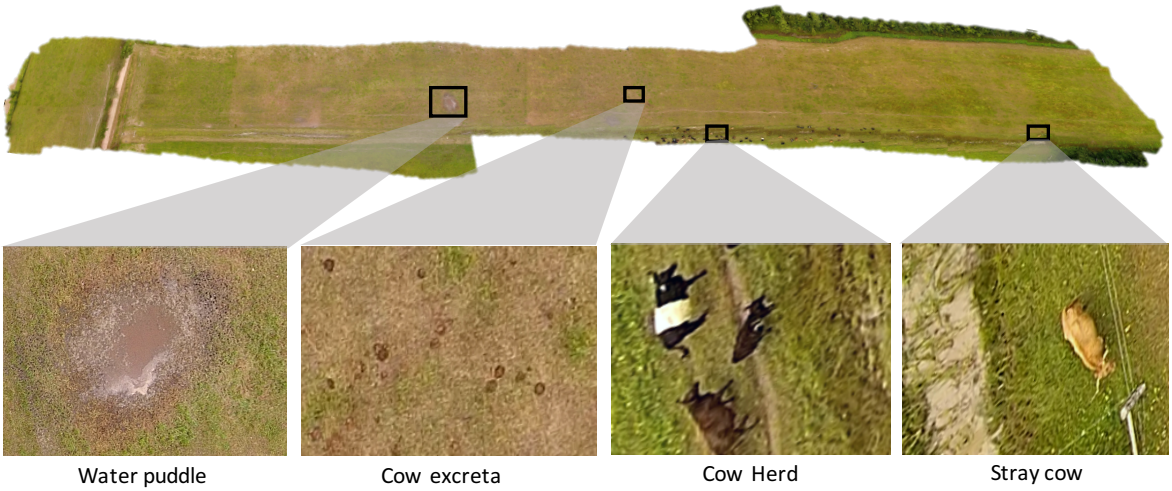

Figure 2.4: **Orthomosaic Generation:** The high resolution orthomosaic generated by FarmBeats for a 5 acre patch in the large farm reveals important visual details to the farmer, such as those shown in the insets – puddles that can make part of the land unavailable for agriculture, cow excreta that becomes manure and enriches the soil, location of individual cows grazing on the farm and their distance from the nearby electric fence.

2.5.2 Generating Orthomosaics from UAV Videos

UAVs generate a prohibitive amount of video that is difficult to transfer to the cloud due to poor network connectivity on farms. For example, a 4 minute flight with a UAV capturing 1080p video at 30 frames per second generates almost a Gigabyte of video data. We make the observation that the unit of interest for the farmer is not the drone video itself, but an overview of the farm that can be provided by a geo-referenced panoramic overview, which is one-two order of magnitude more compact than the full resolution video (see Figure 2.4). The stitched orthomosaic generated from the drone video provides a high resolution visual summary of the farm from a low altitude vantage point, revealing minute details. In fact, existing agricultural drone solutions ([38, 173]) ship the videos to the cloud and convert them into orthomosaics to show to the farmer. Thus, we incorporate the orthomosaics processing pipeline into the FarmBeats Gateway, to process the drone videos locally.

Broadly speaking, the panoramic views can be constructed from the UAV video using two approaches, based on either (i) aerial 3D mapping [141, 172] or (ii) image stitching and mosaicking [124, 11, 23, 185, 178]. While the aerial 3D mapping is a general-purpose method to reconstruct high resolution 3D surface maps of the environment from aerial videos, the image stitching methods treat the world as planar and simply stitch the different images together by finding their relative positions.

Computing high-resolution surface maps is both compute and memory intensive and is not suitable for the resource-constrained farm gateway. On the other hand, while image stitching methods can be incorporated into the gateway, the planar terrain assumption becomes invalid on the farm. Uneven ground geometry, trees, animals or man-made structures observed in the video generates parallax which cannot be handled by the image registration algorithms that assume a planar scene. As we show later in Section 3.5 and as observed in prior work [109], existing image stitchers – Microsoft ICE [124], AutoPano [11] tend to produce distorted orthomosaics in such scenarios. This presents us with an uncomfortable tradeoff: either fly high such that the farm appears planar and sacrifice fine details of the farm, or ship the large aerial videos to the cloud for processing.

Our approach: In order to break this tradeoff, we have developed a hybrid technique which combines key components from both 3D mapping and image stitching methods. On a high level, we use techniques from the aerial 3D mapping systems, just to estimate the relative position of different video frames; without computing the expensive high resolution digital surface maps. Since this process can be performed at a much lower resolution, this allows us to get rid of the harsh compute and memory requirements, while removing the inaccuracies due to non-planar nature of the farm. Once these relative positions have been computed, we can then use standard stitching software (like Microsoft ICE) to stitch together these images. The performance achievements of this hybrid approach are evaluated further in Section 3.5.

2.5.3 Generating Precision Maps

As discussed before, precision agriculture relies on accurate precision maps of the farm that indicate the distribution of a specific characteristic throughout the farm. The FarmBeats-gateway naturally enables a novel approach to precision map generation that can use the aerial imagery from drones to perform spatial inference of sensor values from sparsely deployed sensors.

Specifically, FarmBeatsuses the orthomosaic generated from the drone videos together with the sensor values observed by the sensors planted in the soil, and generates predictions for the entire farm. For example, sensors that observed soil temperature at the discrete locations can inform the machine learning pipeline to make predictions about every location in the farm by considering spatial proximity as well visual similarity of the locations to the sites with the sensors.

Proteus's gateway embeds a machine learning pipeline that draws on probabilistic graphical models that embed Gaussian processes [148]. The key intuition in the proposed model is *spatial and visual smoothness*: areas that are similar should observe similar sensor readings. Specifically, the model relies on two kinds of similarities:

- **Visual Smoothness:** Areas that look similar have similar sensor values. For example, a recently irrigated area would look darker and hence, has more moisture.
- **Spatial Smoothness:** Since we are measuring physical properties of the soil and the environment, the sensor readings for locations that are nearby should be similar.

We encode these two intuitions into a graphical model using standard techniques and formulate it as a Gaussian process regression model [148].

In our current design, FarmBeatsuses the precision maps as units of summarization for the UAV data and ships them to the cloud. This has two advantages over the using orthomosaics as the unit of summary. First, they incorporate sensor data from the farm into drone videos. Second, they can be compressed to two to three orders of magnitude smaller size than a orthomosaic. So, while the orthomosaic is good for giving the farmer a detailed overview of

the farm, precision maps are better for long term storage and shipping. We envision that for other machine learning applications as well, feature maps like the precision maps of the field would be the summaries that get shipped to the cloud, while the descriptive data delivers short-term applications on the gateway.

2.6 Deployment

We deployed FarmBeats in two farms located in Washington (WA) state and in upstate New York (NY), with an area of 5 acres and 100 acres, respectively. The farmer in WA grows vegetables that he sells in the local farmers market. The farm in upstate NY follows the community supported agriculture (CSA) model, and grows vegetables, fruits, grains, as well as dairy, poultry, and meat. Our deployments consist of: sensors, cameras, UAV, the IoT base station, a gateway PC, the cloud service and a dashboard (mobile app and a web page).

Sensors: Each farm was equipped with sensors that measure soil temperature, pH, and moisture. In case of sensors without Wi-Fi support, we interfaced them with Arduinos, Particle Photons or NodeMCUs to add Wi-Fi capability. While the exact number of sensors varied over the deployments and the application of interest, we have deployed over 100 different sensors. Additionally, We deployed Microseven IP [119] cameras in different parts of the field to monitor the farm, as well as to capture IR images of crops. To avoid potential damage from environmental impacts, each sensing platform was encased in a weatherproof box. An example of a sensor deployment can be seen in Figure 2.5(a).

Drones: We used the DJI Phantom 2, Phantom 3 and Inspire 1 for our drone flights.[3] We created an auto-pilot application using the DJI Mobile SDK [45] to interface with Proteus. The user can use the app to first select the flight altitude and determine the area to be covered on an interactive map. FarmBeats' app then plans a flight path using the algorithm proposed in Section 2.5.1. After the drone completes its mission, it automatically returns to its home position and transfers the video recording during the flight to the gateway, through

[3]We received an exemption from the FAA to fly the UAV.

the IoT base station.

IoT Base Station: At each IoT base station deployment, we set up a TVWS network using the FCC certified Adaptrum ACRS 2 radios [4] operating at 20 dBm, and 11 dBi directional antennas with 90 degree sectors. The internet connectivity was provided by the home internet connection of the farmers. To power the base station we setup a solar charging system, which comprised of two 60 Watt solar panels connected to a solar charge controller. The powering system is backed by four 12V-44Ah batteries connected in parallel. The power output goes through an 8-port Digital Logger PoE switch[111]. This provides us the capability to turn on or off individual components of the base station. A Raspberry Pi 3 with 64 GB SD card serves as the base station controller. The sensors interfaced with the base station through a 802.11b router, with a range of over 100 m.

Gateway: The gateway is a Lenovo Thinkpad in the WA farm and a Dell Inspiron laptop in the upstate NY farm.

Cloud: We use the Azure IoT Suite ([120]) for FarmBeats. The sensor readings, camera images, and drone video summaries are populated through the Azure IoT Hub ([121]), to storage. We use blobs for images, and tables for the sensor readings. Although in our current implementation, the different farms share the Azure account, with table-level access control, we plan to have different cloud service accounts for the different farms, as FarmBeats scales up.

2.7 Results

We evaluate the components of FarmBeats below:

2.7.1 Weather Aware Base Station

The FarmBeats base station leverages the algorithm in Section 2.4 to duty cycle different components. It uses the OpenWeather API[137] to get the weather forecasts and plans the duty cycling scheme for the next day. The weather information gives us the cloudiness

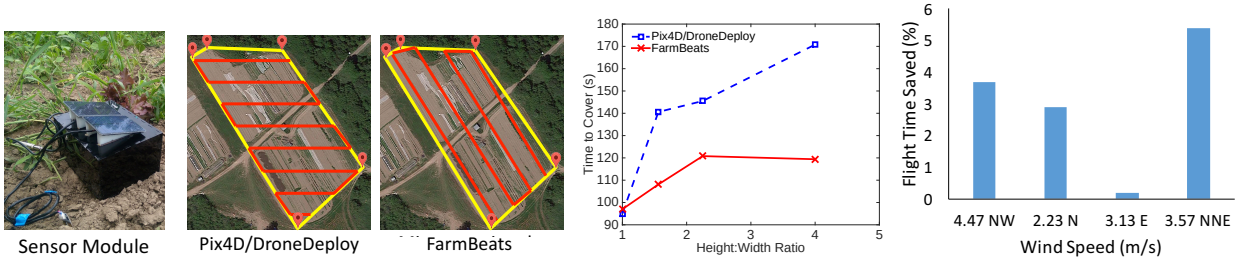

(a) (b) (c) (d)

Figure 2.5: (a) A weather-resistant, solar-powered FarmBeats sensor module.(b,c,d) Drone Flight Planning: (b) FarmBeats' flight planning algorithm minimizes the number of waypoints to cover a region. (c) Depending on the aspect ratio of the field, flights without FarmBeats' algorithm take upto 42% more time. This improves the time by a factor of 1.26 in the average case for our farms. (d) In addition, the yaw control algorithm described in Section 2.5.1 achieves a gain of up to 5% based on the wind velocity.

percentage for each period of three hours. The cloudiness percentage over three days is plotted in Figure 2.6(a).

Over this set of three days, we compare three power-awareness schemes. We define the start of the day as 6AM local time. We periodically record the state-of-charge of our solar power backed batteries. First, we let the base station be always on. As shown in Figure 2.6(b), the battery charge goes up during a sunny day and down during the night. While the base station remains energy neutral during the first day, during subsequent days its battery drains because of cloudy weather, leading to unavailability on the third day. Then, we evaluate the alternate approach. We set the base station to a conservative duty cycling period. While this ensures that the base station is available on cloudy days, the base station battery charges up to 100% during the sunny days thus wasting solar power that could have been utilized. Moreover, its duty cycling interval collects *15 times* less data than the optimal FarmBeats solution, plotted in 2.6(d).

FarmBeats collects data on the first two days more frequently owing to high availability of solar power. However, on the third day, it switches to a conservative duty cycling schedule

to save power. Of the 15x gain in data collection frequency achieved over a fixed duty cycle, a factor of 2 is because of the TVWS client being duty cycled at a different rate than the Wi-Fi router. An earlier version of our deployment which did not duty cycle the base station faced a downtime of 30% in a cloudy month as opposed to zero downtime for our power-aware design in the same month. Thus, FarmBeats' power-aware design achieves its goal of maximizing data-freshness while maintaining energy neutrality.

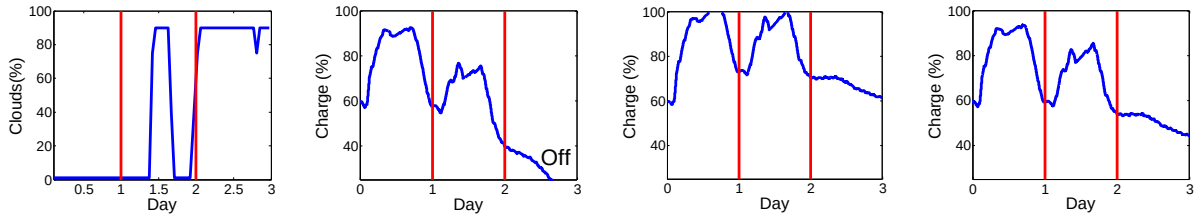

(a) Cloud Cover (b) No Duty Cycle (c) Fixed Duty Cycle (d) Proteus

Figure 2.6: **Power-aware Base Station:** The cloudiness percentage over 3 days. (b) With no duty-cycling, the base station shuts down on a cloudy day. (b) A fixed conservative duty cycle can prevent the base station from going down, but it collects 15 times less sensor data. (d) FarmBeats' Power-aware base station can keep the base station on by reducing the duty-cycling on days are expected to be cloudy.

2.7.2 UAV Flight Planning

As mentioned in Section 2.5.1, we use an efficient area coverage algorithm in addition to leveraging wind assistance to extend drone battery life. To understand the impact of area coverage algorithms on drone flight time, we compare performance of FarmBeats in covering a given area as compared to the state-of-the-art *East-to-west* algorithm (used by Pix4D, DroneMapper, etc.). As shown in Figure 2.5(b), the east-to-west algorithm generates sweeping patterns from the east to the west or vice-versa regardless of the area shape. However, FarmBeats generates a path that minimizes the number of waypoints.

Next, we compare the time taken to complete flights planned by the two algorithms to cover a given area. The maximum speed was set to 10m/s and the altitude was set to

20m. Figure 2.5(c) plots the time taken to complete a flight with the two algorithms in different area geometries defined by their height to width ratio, where height is the distance along the North-South direction and width is measured along East-West. As expected, the gain achieved by FarmBeats increases as the height-width ratio increases. This is because FarmBeats algorithm generates fewer waypoints to cover the same area. In general, for the average case of our deployments, FarmBeats reduced the time taken to cover an area by 26%.

(a) (b) (c)

Figure 2.7: **Orthomosaic Generation:** (a) The Google Earth image for the farm in Figure 2.4. (b) Microsoft ICE image stitching pipeline fails to reconstruct it accurately. (c) Pix4D takes about 2.2x longer on average compared to our approach.

Finally, we evaluate the impact of our yaw control algorithm under different wind conditions. The maximum speed was set to 10m/s and the altitude was set to 30m. For every flight, we fully charged the battery. We measure the percentage of time saved by FarmBeats' yaw control algorithm for each flight and plot it in Figure 2.5(d). As seen in the figure, FarmBeats can save up to 5% time depending on the wind velocity. Moreover, as the north-south component (the principal direction of motion for this set of experiments) of the wind increases, FarmBeats can leverage it better.

2.7.3 Orthomosaic Generation

The novel orthomosaic generation algorithm proposed in this paper advances the state-of-the-art on two fronts. First, our approach of combining sparse 3D reconstruction techniques from video with image stitching techniques is more robust than existing techniques based on either aerial 3D mapping or aerial image stitching. In addition, our approach is computationally

more efficient and runs considerably faster than Pix4D [141], an aerial 3D mapping-based tool catering to Precision Agriculture.

Qualitative Results: We show two representative orthomosaics constructed by FarmBeats and Microsoft ICE in Figure 2.4 and 2.7(b) respectively. Figure 2.7(a) shows what the farm looked like in Google Earth in the past. The orthomosaic generated by Microsoft ICE failed in this case, while our result is consistent and accurate. Our geo-referenced image covers about 5 acres of farmland and provides a detailed visual summary to the farmer. By visually inspecting the high-resolution image, they can discover anomalies such as the water puddle that can render a part of the field unsuitable for agriculture for a couple of seasons. Moreover, the farmer can see where cows are grazing during the day and make a decision about whether they want to move them to another spot for the next day. The decision is based on how much grass they want to leave on the field to be converted into manure.

Processing Time: As shown in Figure 2.7(c), our implementation is 2.2 times faster than Pix4D on average. Specifically, our method took 14 minutes to construct an orthomosaic on average whereas Pix4D took 32 minutes on average on a set of videos captured by our drones at 1080p resolution at 30 frames per second. This demonstrates the improved running time of our method.

Finally, the orthomosaic generated by our system are approximately 5 times smaller than the original video size at full resolution (in .png format) before applying lossy compression. A single pixel in the geo-referenced orthomosaic is about 2 cm in size which is equivalent to a single penny on the ground. The image resolution and compression quality are parameters that can be tuned to meet any target file size.

2.7.4 Generating Precision Maps

As described in Section 2.5.3, FarmBeats uses the visual features from the orthomosaic overview to extrapolate the sensor values and generate precision maps for soil temperature, soil moisture and pH.

Qualitative Evaluation: We show a representative set of these precision maps in Figure

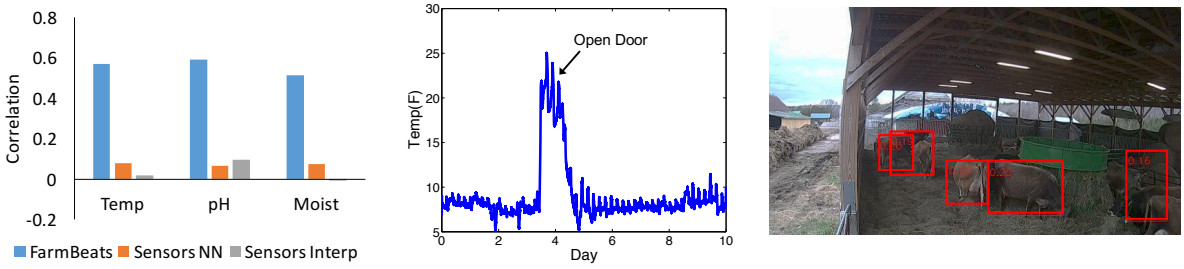

(a) Precision Maps (b) Storage Monitoring (c) Cow Shed Monitoring

Figure 2.8: **FarmBeats Applications** (a) FarmBeats' precision maps are more accurate than standard sensor based interpolation techniques. (b) Temperature (measured in F) in a storage unit can raise an alarm when an employee leaves a door open. (c) Cows being monitored in a cow shed. The red boxes indicate a standard cow detector output.

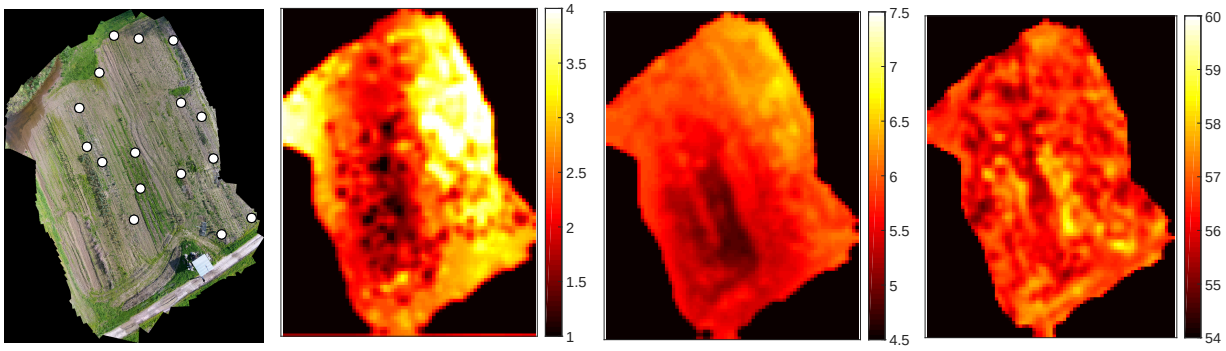

(a) Orthomosaic (b) Moisture Map (c) pH Map (d) Temperature Map

Figure 2.9: **Precision Maps:** (a) A 40 MPixel orthomosaic created from a 3 minute flight over 2 acre area of a farm. Our system infers dense sensor measurements from very few sensors deployed on the farm (indicated by white circles). (b) The predicted soil moisture map (our sensors measures moisture on a scale of 1 to 5). Note that the top left region in the image where the ground appears wet was correctly predicted to have high moisture even though no moisture sensors were present in that part of the farm. (c) The predicted pH map (pH is measured from 0-14, 7 is neutral and 0 is the most acidic). Our system identified that the whole field is slightly acidic, but the bottom left/center is more acidic than the rest. (d) The predicted soil temperature map (in Fahrenheit scale).

2.9. As shown in the figure, based on sensor values in the rest of the farm, the moisture prediction pipeline can estimate that the top left part of the farm has high moisture content even though that part has no sensor there. Similarly, the pH map generates an actionable input in the sense that the bottom left and center of the farm have very low pH and are highly acidic. As a result of this map, the farmer applied lime to enhance the pH and make the soil more neutral.

Note that the pH of the farm varies within the farm at fine granularity. As seen in Figure 2.9, within a couple of acres, the pH can vary from 4 (very acidic) to 7 (neutral). Soil moisture variance is even higher, with variance seen within a few meters. Precision maps generated by FarmBeats capture this variance accurately, by using the drone videos to extrapolate the sensor data.

Quantitative Evaluation: In order to evaluate the accuracy of the precision maps generated by FarmBeats using the approach described in Section 2.5.3, we evaluated our system on 5 datasets constructed from the drone videos and sensor data. Each dataset corresponds to a drone flight over the farm (covering 2 acres) and one set of sensor measurements from the sparse sensor deployment. The hyperparameters are learned by doing 5 fold cross validation. As an accuracy metric, we measure the correlation between the predicted sensor values and the ground truth sensor values to see how well the variations in the field are captured by Proteus. We compare against two techniques, which do not use the drone video based extrapolation of the sensor values:

- **Nearest Neighbor (SensorsNN):** We assign the value from the nearest sensor to each point in the field.
- **Inverse Distance based Interpolation (SensorsInterp):** We linearly interpolate known sensor values in the field, by using inverse distance as a weight. This technique has been previously been proposed in the context of precision agriculture [182, 47].

For all the analysis, we use leave-one-out evaluation, i.e., we generate a precision map after leaving one of the sensors out of the training set and evaluate the map on the left out sensor.

We repeat this process for all the ground sensors and report the averaged results.

The comparison of correlation across the different schemes is shown in Figure 2.8(a). As shown in the figure, FarmBeats outperforms existing sensor based interpolation techniques. In particularly, FarmBeats can accurately estimate the variations of the different sensor values in the field. While sensor based methods do not mirror the variations and hence have nearly zero correlation with the sensor values, FarmBeats' estimates have high positive correlation with the true sensor values, thus indicating the utility of using the drone video in conjunction with the drone estimates. Finally, the precision maps generated by FarmBeats are 3 orders of magnitude smaller in size on average than the video and can be easily shipped to the cloud during periods of connectivity.

2.7.5 Other Applications

Figure 2.8 highlights two other applications that the farmers used FarmBeats for. First, the farmer in NY used FarmBeats sensors to monitor his storage freezers. The temperature in these freezers is carefully regulated below 10° F to prevent produce from going bad. As shown in Figure 2.8(b), an employee leaving the door open could lead to this temperature going up causing loss to the farmer. This problem is solved by FarmBeats by enabling automated notifications based on these sensor readings in the FarmBeats phone application.

Second, the farmers plugged in cameras at different locations like cow sheds and connected them to the nearest FarmBeats base station. One frame of the camera is shown in Figure 2.8(c). While the intent of the current application is to manually monitor the cows, one can potentially build an application that can detect anomalies in cow behavior or use cow motion to track animal health[132]. As a preliminary result, we ran a deep neural network based cow detector on the data. The identification boxes are overlaid on the figure.

2.7.6 End-to-end Deployment Statistics

Data Aggregation: FarmBeats' deployments at both farms have been running for over six months. Over these deployments, FarmBeats interfaced with around 10 different sensor

types, three different camera types, three versions of drones and the farmers' phones. It collected more than 10 million sensor measurements, half million images and 100 drone surveys.

Resilience to Outages: FarmBeats' deployments faced one week-long Internet outage due to a thunderstorm and several smaller term Internet outages. The FarmBeats gateway continued to be available during these times.

Cost: The TVWS client radios cost $200,[4] and there are no additional data charges, than the farmer's existing internet connection. The Particle Photons cost about $20 and can add Wi-Fi support to each sensor. Thus, use of the hybrid networking approach reduces the system cost by an order of magnitude as compared to existing systems which cost over $1000 in equipment cost per sensor and over 100$ annual subscription fee (see Table 2.1).

Applications: Farmers used FarmBeats' precision agriculture system to guide their precision irrigation units. The precision pH maps generated were used by farmers to apply lime in the more acidic regions. As mentioned before, farmers also used FarmBeats for storage monitoring with sensors and animal shelter monitoring, selling station monitoring with cameras. Beyond that, farmers also used FarmBeats base stations to access Wi-Fi while on the farm to run productivity applications like Trello.

2.8 Case Study

Since the initially deployments of FarmBeats starting in 2015, the system has been deployed on farms worldwide and has shown significant benefits. As a case study, we focus on the FarmBeats deployment on Nelson Farm located in Eastern Washington. Nelson farm spans approximately 7500 aces across 45 miles and has a growing season that runs between the months of March and July. The primary crops being grown include dry-land wheat, lentils, peas, and garbanzo beans.

---

[4]With the standardization of IEEE 802.11af [1] standard, we expect the price to the client and base station to be similar to Wi-Fi, of less than 10$. We are testing one such multi-mode TVWS/Wi-Fi chip from a major Wi-Fi vendor.

Figure 2.10: **Nelson Farm Deployment.** A TV whitespace base station and FarmBeats sensor box deployed on Nelson Farm.

In this particular deployment, the farmer has Internet connectivity at their home but it is not capable of providing full coverage on the farm. To address this issue, a broadband TVWS network is deployed to enable backhaul connectivity to the cloud. FarmBeats sensor boxes are deployed across the farm that talk to a base station using either narrowband TVWS or LoRa depending on the application scenario. With narrowband TVWS, the sensor boxes can send data to the base station while being as far as 18 miles away. Every sensor box is powered by a single battery and backed by solar power, which helps minimize overall maintenance of the system. Moreover, the FarmBeats sensor box is designed to be modular and can be integrated with various types of off-the-shelf sensors, both analog or digital. This includes sensors such as soil moisture and temperature, CO_2, wind speed and direction, ambient temperature, and much more. In Figure 2.10 we show a picture of the TVWS broadband deployment and a sensor box.

The FarmBeats system has been operating at Nelson Farm since 2019 and has provided numerous insights to the farmer to improve productivity. Panorama overviews of the farm and precision heatmaps help the farmer make intelligent decisions when it comes to day-to-day tasks. For instance, they enable precision pesticide or herbicide application, where chemicals are only applied where they are needed, which saves cost and improves overall plant

health. Specifically, the application of the most expensive chemicals was reduced by 90% and cost savings increased by 15% [86]. FarmBeats also provides micro-climate predictions of temperature and wind speed and direction, which has immediate impact on the farm. The micro-climate across Nelson farm can have temperature differences of over 10 degrees. When the temperature drops below freezing within 24 hours of some field operations, it can cause a decrease in crop yield by up to 50% and significant revenue loss. Being able to predict such freezing temperatures accurately was crucial in May 2019. The farmer was able to change his plan of action in time when the FarmBeats system predicted below freezing temperatures and his weather apps did not. Moreover, wind microclimate predictions have proven beneficial in running operations, where equipment can be sent to locations that have the least amount of wind and keep conducting field operations rather than completely halting all tasks.

2.9 Conclusion

FarmBeats is a low-cost, highly available IoT platform for agriculture. It supports high bandwidth sensors using TVWS, which is a low-cost, long range technology. FarmBeats uses a weather-aware solar-powered IoT base station, and an intelligent Gateway that ensure that services are available in the Cloud and offline. It also incorporates new path-planning algorithms that extend drone battery life. We have deployed the system in two farms, and the farmers are already using it for three applications: precision agriculture, animal monitoring, and storage monitoring. Moving forward, we are working with the farmers to develop several other applications on top of FarmBeats. Further, we plan to make anonymized data available for researchers to enable more agricultural applications.

Chapter 3

WIRELESS QUANTIZATION INDEX MODULATION

3.1 Introduction

Over the last decade, we have witnessed a rapid growth in the deployment of IoT devices. By some estimates, there will be more than 26 billion connected IoT devices by the year 2020 [62]. However, as more devices connect to wireless networks, *available spectrum is insufficient and existing wireless protocols are ill-equipped to support the growing number of devices.* To understand the challenge, consider a home with wireless cameras, security sensor, smart watches, fitness trackers and wireless speakers. These devices use Wi-Fi/Bluetooth or proprietary wireless in the 2.4 GHz ISM band and operate alongside Wi-Fi routers, smartphones, laptops and tablets. As more devices share the wireless channel, wireless interference and packet collisions increase, negatively impacting the throughput and latency [129] [65].

As wireless spectrum (such as the 2.4 GHz ISM band) becomes crowded, conventional wisdom dictates that we migrate to new protocols in less congested wireless channels. New protocols such as 802.11ah, LoRaWAN [191], SIGFOX [177] operate in 915 MHz ISM band, high speed 802.11 n/ac Wi-Fi is moving to 5.8 GHz ISM band, NB-IoT operates in the licensed cellular bands and TV white space networking [14] operates in unused channels in the TV UHF spectrum. Although these solutions are a step in the right direction, let's discuss these approaches in terms of cost and spectrum utilization.

- *Infrastructure and Maintenance Costs*: Migration to new protocols such as LoRaWAN, SIGFOX or 802.11ah require setup, deployment and maintenance of dedicated expensive gateways and base stations. Instead, if we can reuse existing wireless infrastructure for communication, we could develop a far simpler and cost effective solution.

- *Spectrum Utilization*: Wireless spectrum is an extremely valuable and highly regulated resource. TV white space networking uses allocated but otherwise under-utilized TV spectrum for wireless communication. However, more often than not the availability of unused TV channels in urban areas are scarce and it is very cumbersome and expensive to deploy a TV whitespace network [55]. New protocols such as LoRaWAN, SIGFOX, and 802.11ah are moving to the less crowded 915 MHz ISM band, but over time as the number of devices increase, they are going to run into familiar interference and capacity issues: as the number of devices increase, the spectrum is going to become more crowded and eventually saturate. Unless new spectrum is made available, using traditional methods, it is impossible to scale beyond a certain point.

In this paper, we propose a new cost and spectrally efficient solution for wireless communication. Consider an urban city environment as shown in Fig. 3.1 with existing deployments of AM, FM, TV, and cellular base stations. These base stations have been setup with tremendous infrastructure cost, undergo periodic maintenance and pay licensing fees to transmit at pre-assigned licensed frequencies. The base stations are designed and geographically located for optimal signal coverage. For example, a typical FM tower can be received up to 100 kms.

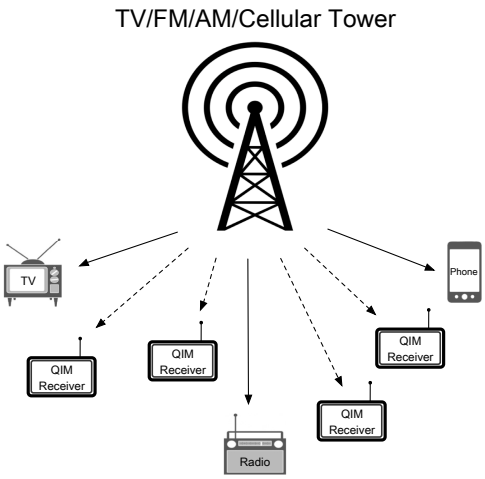

Figure 3.1: **QIM deployment.** QIM uses existing wireless signals to communicate with IoT devices.

We introduce Wireless Quantization Index Modulation (QIM), a communication technique which leverages existing infrastructure and reuses broadcast signals to provide additional communication channels for IoT devices. To understand Wireless QIM, without the loss of generality, let's consider a broadcast TV

station. A TV transmitter can use the QIM technique in its baseband to embed a message into the broadcast TV signal by introducing small perturbations while having a negligible impact on the broadcasted TV signal. Legacy TV receivers in the coverage area decode the broadcasted signal as before while IoT devices with a QIM receiver can decode the embedded message without any prior knowledge of the broadcast signal. So, in summary with a small modification to the baseband of the broadcast station, Wireless QIM reuses infrastructure, spectrum and broadcast signals to simultaneously communicate with QIM enabled IoT devices and legacy AM/FM/TV/cellular devices.

Wireless communication requires both uplink and downlink. However, more often than not, it's asymmetric i.e. depending on the application, either uplink or downlink communication dominates. In this paper, we focus on downlink heavy applications and design a Wireless QIM system for downlink communication. Our target application is a smart city where using Wireless QIM, existing wireless infrastructure provides connectivity for real-time update of electronic bus schedule displays, billboard signs and advertisements, traffic alerts to name a few. With a minimal change in the baseband of existing broadcast towers, we can embed data to wirelessly update devices with a QIM receiver in real-time. These applications would require a minimal uplink channel to send acknowledgement messages, however such a low bandwidth and infrequent task can be accomplished using traditional LoRa, SigFox or cellular radios for the time being. In future work, we will extend the Wireless QIM technique to uplink communication and develop a bi-directional communication system which can leverage existing infrastructure and communicate with smart devices with zero spectrum overhead and minimal additional cost to target a broader set of applications.

To demonstrate the efficacy of Wireless QIM for these applications, we extensively evaluate the design space and explore various tradeoffs between performance of the message and host signal. We implement Wireless QIM on three existing infrastructure broadcast signals: AM, FM, and TV and show that information can be reliably embedded with negligible impact on audio (AM and FM) and video (TV) quality of the host signals. Our results show that using Wireless QIM, we can embed messages for IoT devices at up to 8 kbps in AM

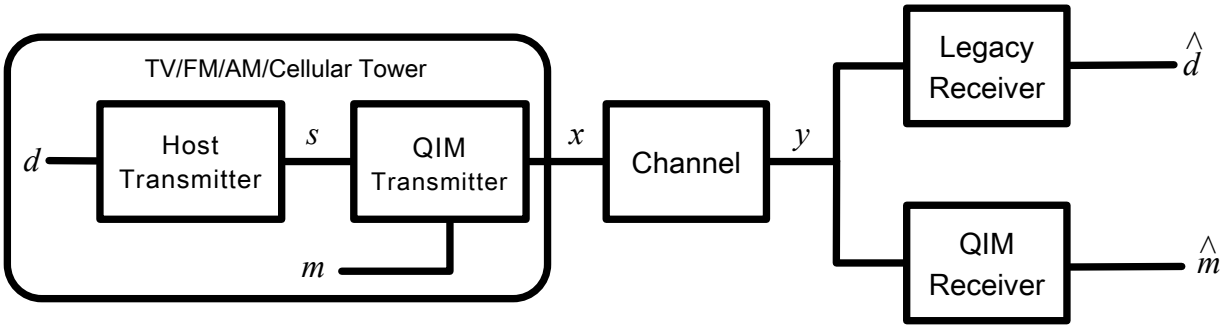

Figure 3.2: QIM embedding model, where a message m is embedded in a host signal s using the QIM embedding function. The signal passes through a noisy channel and the decoder retrieves the estimated message $\hat{m}$ from the received signal, y.

radio signals and 200 kbps in FM signals.

3.2 Quantization Index Modulation

Quantization Index Modulation (QIM) was originally introduced as a scheme for information hiding and digital watermarking [27]. In these applications, a message signal is embedded inside another signal called the host signal such that the embedded message is robust to common degradations, while the host signal suffers minimal degradation. QIM technique achieves an efficient tradeoff between the data rate of the message signal, distortion, and robustness of embedding.

In this work, we use QIM for wireless communication. We show how to embed messages for IoT devices inside existing wireless signals at existing wireless transmitters such that IoT devices with QIM receivers can decode the messages while the broadcasted wireless signal experiences minimal degradation. Fig. 3.2 shows the block diagram for a typical Wireless QIM system. At the AM/FM/TV/cellular broadcast, data d needs to be transmitted and is passed through the host transmitter to generate a baseband host signal s. The host signal can be of any form, for instance, a frequency modulated signal from a FM station. Now, we also want to embed a message m intended for IoT devices. The host signal is passed through

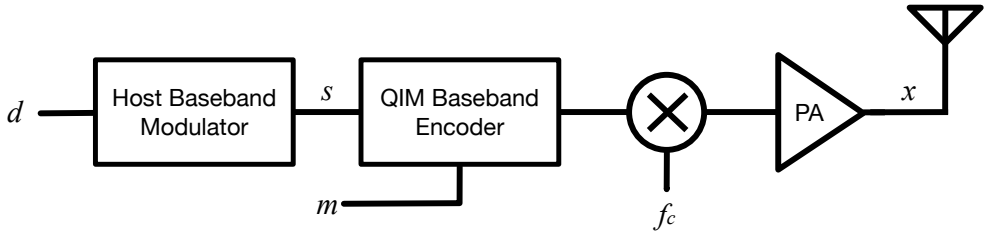

Figure 3.3: Block diagram of a QIM enabled transmitter.

a baseband QIM transmitter that uses a quantizer $Q(s)$ to embed the message inside the host signal. This generates a composite signal x that propagates through the (noisy) channel and signal y is received at the receiver. A legacy device would use standard compliant host receiver to decode the host signal, while an IoT device would use a QIM receiver to decode the embedded message m. QIM is designed to ensure that both the IoT and legacy devices can decode the desired signals with minimal degradation.

A Wireless QIM system consists of a QIM transmitter at the base station and a QIM receiver at the IoT device. To understand how Wireless QIM works, we will first describe the QIM transmitter followed by the QIM receiver.

3.2.1 QIM Transmitter

On a high level, a QIM transmitter embeds information by using a form of quantization to introduce small perturbations in the host signal. Fig. 3.3 shows the block diagram of a QIM transmitter, where data d is passed through the baseband host modulator to generate the host signal s. Then the host signal and message m are passed through the QIM encoder to generate a composite signal x in baseband. Finally, the composite signal is upconverted to carrier frequency f_c and transmitted. The implementation of the baseband QIM encoder depends on the host signal. For analog systems such as AM radio, QIM encoder consists of the digital baseband followed by the ADC, a standard component. Whereas in digital FM radio and TV systems, the QIM encoder can be integrated into pre-existing digital baseband.

In the following section we explain how the QIM embedding process works. First lets define a uniform quantizer $Q(s)$ as

$$Q(s) = \Delta \cdot round(\frac{s}{\Delta}) \tag{3.1}$$

where Δ is the quantization step size and is defined as

$$\Delta = \frac{2 \cdot max(s)}{N} \tag{3.2}$$

Here N is the number of quantization levels. The step size and number of levels determine the embedding resolution for the host signal. Quantizer *Q(s)* can now be used to define the QIM embedding function,

$$Q_m(s) = Q(s - d_m) + d_m \tag{3.3}$$

where d_m is the dither, a function of the message m that is applied to the host signal. d_m can take one of the two following values to represent embedding of either a 0-bit or 1-bit.

$$d_1 = \pm\frac{\Delta}{4} \text{ and } d_0 = \begin{cases} d_1 + \frac{\Delta}{2}, & \text{if } d_1 \leq 0 \\ d_1 - \frac{\Delta}{2}, & \text{otherwise} \end{cases} \tag{3.4}$$

The equation shows that if the 1-bit dither is negative, $d_1 = {}^{-\Delta}/_4$, then 0-bit dither, $d_0 = {}^{\Delta}/_4$, will be a positive value. Similarly, for a positive 1-bit dither, the 0-bit dither would be a negative value. So, in summary we create two dithered quantizers to embed data in the host signal which is dependent on the bit value of the embedded message m.

Fig. 3.4 illustrates this QIM embedding process for a positive 1-bit dither d_1 with $N = 4$ levels, represented by solid horizontal lines. The dashed horizontal lines represent quantization levels for the 1-bit dithered quantizer and dotted lines for the 0-bit quantizer, defined by Eq. 3.3. We can see that because the dither for a 1-bit is defined as ${}^{\Delta}/_4$, the dashed lines are shifted up by ${}^{\Delta}/_4$ from the original set of levels and vise versa for 0-bit quantizer.

Each sample point of the host signal (dashed green line) is perturbed to the appropriate level depending on the message m (shown at the top of the figure). For instance, at the first

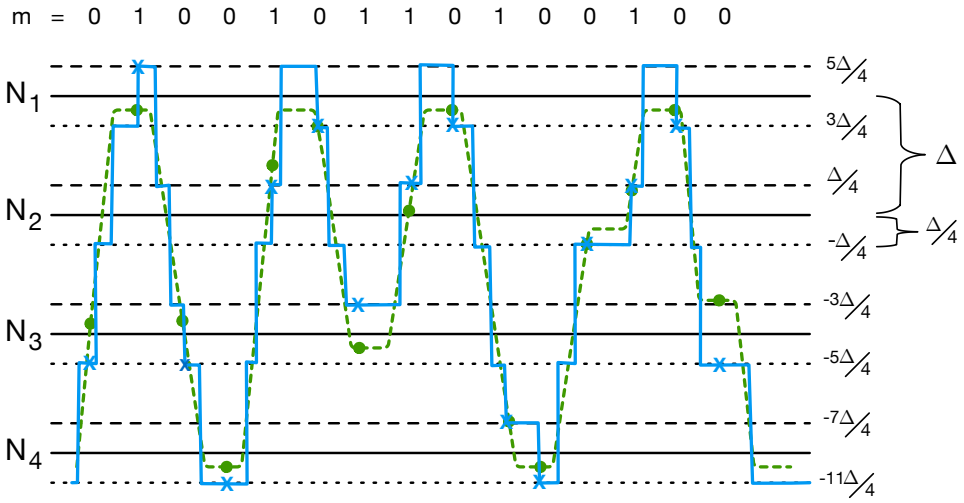

Figure 3.4: Embedding using QIM for $d_1 = \frac{\Delta}{4}$ and $d_0 = \frac{-\Delta}{4}$

highlighted sample point (green dot) we encode a 0-bit and the composite signal (solid blue line) goes down to $^{-5\Delta}/_{4}$, the nearest 0-bit level (blue X). Similarly, at the second highlighted point we embed a 1-bit, which jumps up to $^{5\Delta}/_{4}$, the nearest level for a 1-bit. We can see from the example that the distance between quantization points is uniformly distributed between $[^{\Delta}/_{2}, ^{-\Delta}/_{2}]$. As a result, the mean error due to embedding is equal to $^{\Delta^2}/_{12}$.

Finally, we characterize the impact of the QIM encoder on the host signal. We define distortion in the host signal by comparing the original host signal to the composite signal generated by embedding and can be expressed as,

$$D_s = \frac{1}{K} \sum_{i=1}^{K} |s_i - x_i|^2 \tag{3.5}$$

3.2.2 QIM Receiver

Fig. 3.5 shows the block diagram of a QIM receiver. A QIM receiver is based on standard RF architecture and is similar to a commodity receiver in terms of complexity and power consumption. The QIM decoder is implemented in digital baseband and we simply apply the

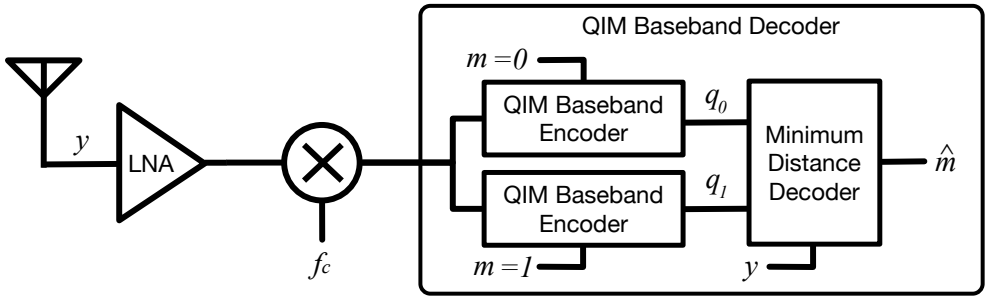

Figure 3.5: Block diagram of a QIM receiver

same QIM embedding function to the received signal for both 0 and 1 bit messages to obtain two quantized signals q_0 and q_1. These are passed through a minimum distance decoder to compare with the original received signal y to obtain estimated/received message,

$$\hat{m} = \arg\min_{m} dist(y, y_m) \tag{3.6}$$

3.3 QIM Techniques

In this section we describe different QIM techniques that can be used to embed messages in a host signal.

3.3.1 Scalar QIM

Scalar QIM is the primary QIM technique which was described in the previous section and can be applied to a real valued or scalar signal.

3.3.2 Distortion Compensated QIM

Distortion compensated QIM (DC-QIM) is an extended version of the QIM method that reduces the distortion in the host signal and significantly improves the distortion to robustness trade-off [27]. The robustness of QIM is a function of the distance between the quantization points. If we increase the distance, robustness of QIM would also increase. This operation

is the same as scaling the QIM embedded signal by a factor, α. For example, if a sample point s_i is shifted to a point, x, then by scaling it by α, the sample point s_i would now be at x/α. However, this increases the distortion of the host signal by a factor $1/\alpha^2$ and results in a $\Delta^2/12\alpha^2$ mean distortion. We compensate for the distortion by adding back a fraction, $1-\alpha$ of the host signal. This operation can be represented by the following QIM embedding function,

$$Q_m(\alpha s) = Q(\alpha s - d_m) + (1-\alpha)s + d_m \tag{3.7}$$

The parameter α is defined as $0 \leq \alpha \leq 1$ where 1 represents the original QIM method. Since, α determines the distortion, a component of noise in the composite signal due to embedding, we can determine the optimal value of α maximizing the signal to noise ratio SNR.

$$\begin{aligned} SNR(\alpha) &= \frac{d_1^2}{(1-\alpha)^2 + \alpha^2\sigma_n^2} \\ \frac{\partial SNR(\alpha)}{\partial \alpha} = 0 \text{ gives } \alpha^* &= \frac{D_s}{D_s + \sigma_n^2}, \end{aligned} \tag{3.8}$$

where, σ_n^2 is the noise power. Hence, the optimal factor α^* is a function of distortion and noise.

3.3.3 Lattice QIM

The QIM technique can be extended to higher dimensional signals by using a lattice form [27]. We can arrange the quantization points to be an integer lattice, Z^N, in an N-dimensional Euclidean space $\mathbf{R}^N$. Let's consider a simplified case of a complex or two-dimensional signal which has both in-phase and quadrature phase components. For N=2, we will have two grids to quantize to either a 0-bit or 1-bit representation. As an extension of scalar QIM, the quantization points for a 0-bit and 1-bit can be represented as co-sets of the lattice,

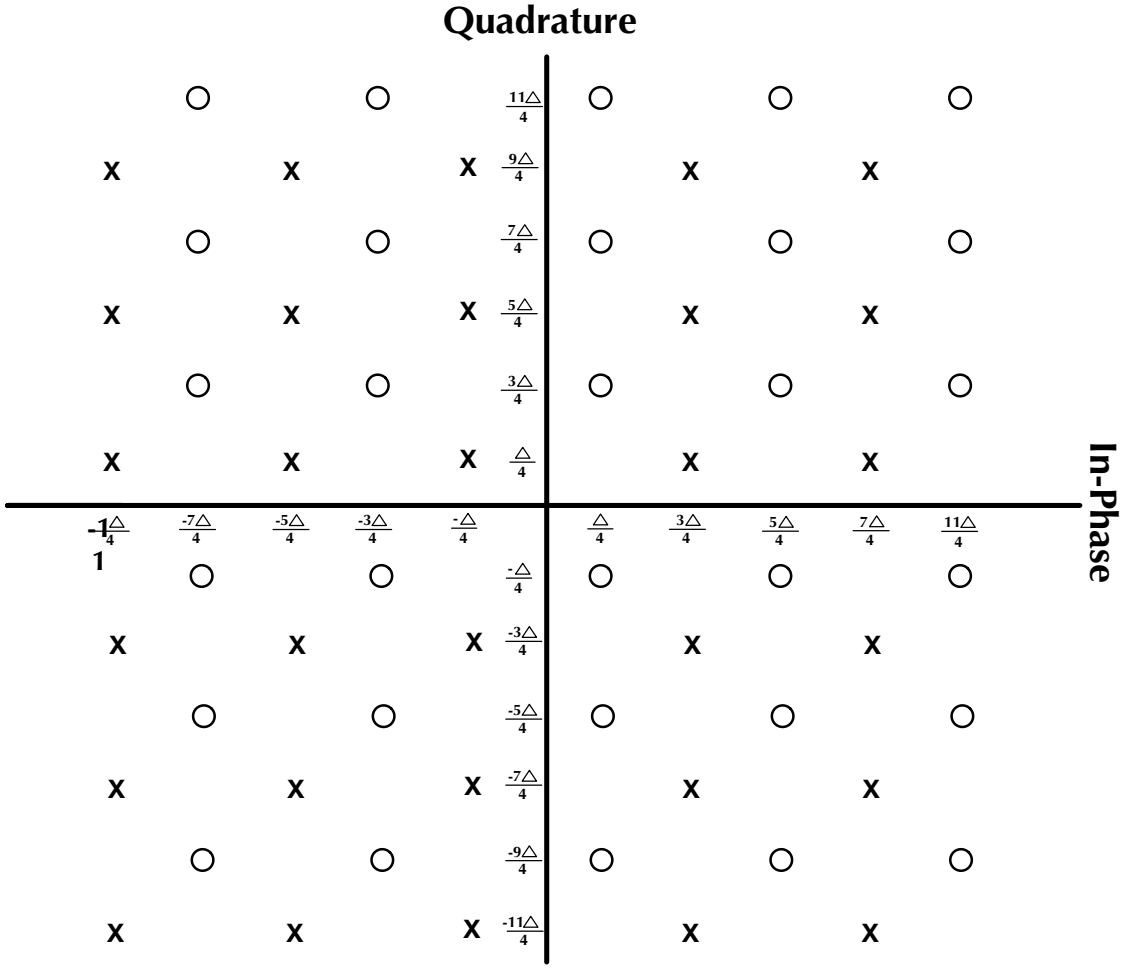

Figure 3.6: Constellation diagram for Lattice QIM with a grid representation. The Xs are quantization points for 1-bit and Os for a 0-bit.

$$\Lambda_0 = \left(\frac{-\Delta}{4} - j\frac{\Delta}{4}\right) + n_i \text{ and } \Lambda_1 = \left(\frac{\Delta}{4} + j\frac{\Delta}{4}\right) + n_i \tag{3.9}$$

where $n_i = 1, 2, ..., l$ and l represents the total number of quantization levels. Fig. 3.6 illustrates the quantization grid for $N = 2$. In lattice QIM, the quantization levels are separated by $\Delta/2$ distance in in-phase and quadrature components resulting in an overall distance of $\Delta/\sqrt{2}$. However, the mean error due to embedding is still $\frac{\Delta^2}{12}$ and because of larger distance between data points, lattice QIM should perform better than scalar QIM. To

implement lattice QIM for complex wireless signal, we quantize both the real and imaginary parts of the signal. We use the same embedding function and modify the dither as

$$d_1 = \frac{\Delta}{4} + j\frac{\Delta}{4} \text{ and } d_0 = \begin{cases} d_1 + (\frac{\Delta}{2} + j\frac{\Delta}{2}), & \text{if } d_1 \leq 0 \\ d_1 - (\frac{\Delta}{2} + j\frac{\Delta}{2}), & \text{otherwise} \end{cases} \tag{3.10}$$

Finally, we note that the distortion compensation technique described above can be applied to the real and complex values of the signal to implement distortion compensated lattice QIM.

3.4 Evaluating the QIM Wireless Link

Evaluation of an embedded communication system (such as QIM) significantly differs from conventional communication systems. Traditional systems deal with only one signal whereas embedded communication systems operate on two signals: the host signal and the embedded signal. We need to evaluate the impact of QIM on the performance of both the host and the embedded message signal: what is the rate and robustness of the embedded signal, and how much did the embedded process distort the host signal and its impact on the output of the host (legacy) receiver.

3.4.1 Performance of Embedded Message

We start by defining the capacity for QIM embedded message in a Gaussian channel. The output of the channel is the sum of the input host signal and Gaussian white noise. Let σ_s^2 be the average power of the host signal and channel noise follows a Gaussian distribution with variance σ_n^2. We can express information capacity, i.e., the supremum of the achievable rate for a real value one dimensional signal as [34]

$$C_{host} = \frac{1}{2}\log_2\left(1 + \frac{\sigma_s^2}{\sigma_n^2}\right) \tag{3.11}$$

where σ_s^2/σ_n^2 denotes the signal-noise ratio (SNR) and assumes that the input also follows Gaussian distribution. Since embedded QIM signal is distortion in the host signal, we can

re-write the capacity expression in terms of distortion by treating the embedded QIM signal as a form of power-limited communication over a Gaussian channel:

$$C_{qim} = \frac{1}{2}\log_2\left(1 + \frac{D_s}{\sigma_n^2}\right) \text{bps/sec/Hz}, \tag{3.12}$$

where the distortion constraint is given by Eq. 3.5. The capacity equation, C_{qim}, shows that the performance of the embedded message is directly proportional to distortion experienced by the host signal. The QIM technique maximizes the capacity of the embedded message for a given distortion of the host signal. The capacity is also a function of the bandwidth of the host signal and a higher bandwidth host signal would enable a higher data rate embedded message. Finally, we note that the capacity in Eq. 3.12 is the maximum achievable data rate per unit bandwidth with arbitrarily small error probability. Practical QIM implementation would require error correction coding mechanisms to achieve performance close to the limits promised by channel capacity.

3.4.2 Impact on the Host Signal

The host signal experiences distortion due to perturbations introduced by the embedded signal which was described in Eq. 3.5. For a fair comparison across host signals, we introduce normalized distortion which is independent of the signal strength of the host signal and is defined as follows,

$$D_s^n = \frac{\sum_{i=1}^{N} |s_i - x_i|^2}{\sum_{i=1}^{N} |s_i|^2} * 100\% \tag{3.13}$$

Finally, in addition to computing the distortion, we will also analyze the quality of the multimedia signal at the output of the host receiver to ensure that QIM embedding operation has minimal impact on the performance of the host signal.

3.5 Simulation Results

The Wireless QIM technique is independent of the host signal and is universally applicable. Here we consider three host signals: AM, FM and broadcast TV which are ubiquitous in cities. We start with a short primer on the host signals.

TV. In the United States, Digital TV (DTV) operates in the UHF band from 470-614 MHz with 6 MHz wide channels and follows the Advanced Television System Committee (ATSC) standard [32]. ATSC uses 8-level vestigial sideband (8-VSB) modulation to transmit data. 8-VSB is a digital modulation technique which uses eight amplitude levels to represent symbols on a 6 MHz channel. Transmissions from a TV tower can be typically received up to 50 miles.

FM Radio. FM radio operates in the 87.8-108 MHz frequency band with 200 kHz wide channels. FM uses analog frequency modulation to encode audio and data i.e. information is transmitted by varying the frequency of the transmitted RF signal. Most FM stations can be heard up to 100 miles from the transmit tower.

AM Radio. In the United States, AM radio operates in the 525-1705 kHz band with 10 kHz channel spacing. AM radio uses amplitude modulation to encode data i.e. information is represented in the amplitude of the signal. AM signals propagate long distances and have been reported to have been received 200 miles away from the station.

We evaluate wireless QIM on recorded AM, FM, and TV signals. The USRP X300 [151] was used to record TV signals centered at 539 MHz (UHF channel 25) with 6.25 MHz sampling rate and FM signal centered at 106.1 MHz with 200kHz sampling rate. We use a WebSDR [204] to record an AM signal centered at 1630 kHz with 8kHz sampling rate. We implement QIM methods described in Section 3.2 and simulate different channel conditions and data rates using MATLAB.

We embed pseudo-random message bits and implement scalar QIM and scalar DC-QIM for real valued AM and FM radio signals. For complex TV signals in addition to scalar QIM, we also evaluate lattice QIM and lattice DC-QIM. In DC-QIM, we set $\alpha = 0.7$, the optimal value as per Eq. 3.8. We introduce additive white Gaussian noise to simulate different

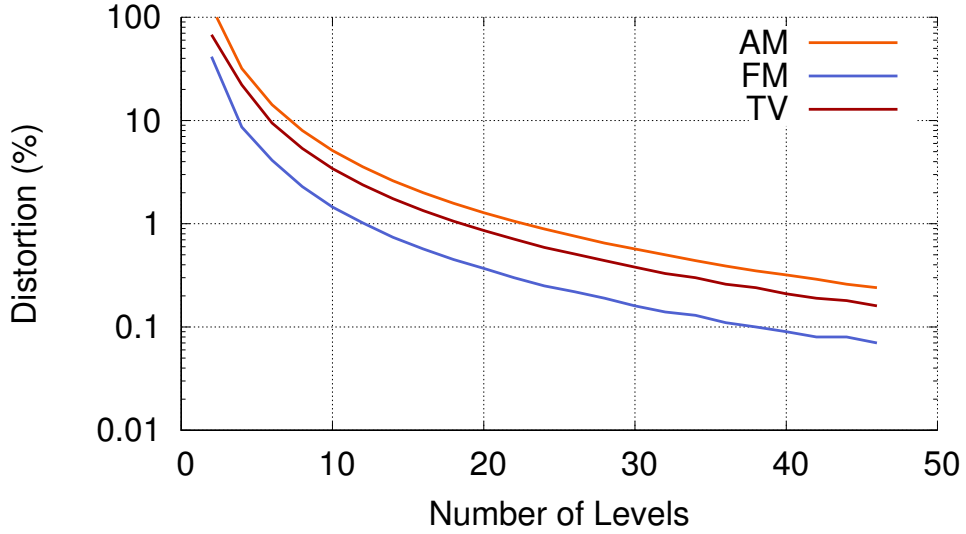

Figure 3.7: Evaluation of distortion for each host signal

channel conditions. Finally, a QIM receiver recovers the transmitted message bits using the algorithm described in Section 3.2.2. The value of Δ, the QIM embedding parameter is known at the receiver. This is a reasonable assumption since it can be either pre-set or periodically updated. We evaluate the system by measuring the impact of QIM on both the host signal and the performance of the embedded QIM signal for different QIM methods at different channel conditions and number of quantization levels (affects distortion).

3.5.1 Impact on the Host Signal

The distortion experienced by the host signal is a function of number of levels N used in the QIM embedding process. We vary the number of quantization levels from 2 to 45 for embedding random messages in TV, FM and AM host signals and measure the normalized distortion and its impact on the performance of the legacy host signal receiver.

Normalized Distortion. Fig. 3.7 shows the percentage of distortion experienced by each host signal as a function of number of levels for scalar DC-QIM technique. The AM signal experiences the most distortion followed by TV and FM. This is expected since AM uses

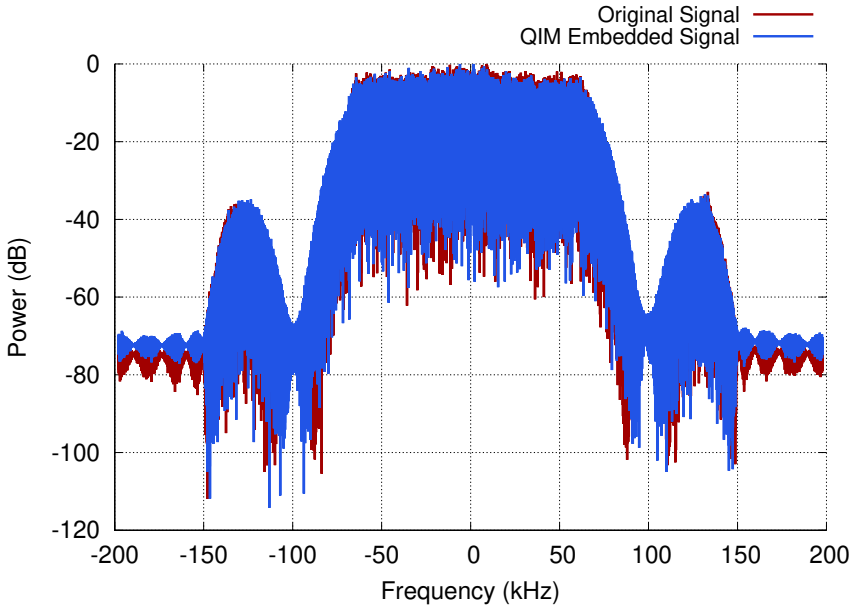

Figure 3.8: Spectrum of the baseband FM signal before and after QIM embedded data.

analog amplitude modulation to encode data and QIM introduces amplitude perturbations, which distorts the information carrying amplitude of the AM signal. Similarly, TV also uses 8 level (digital) amplitude modulation and amplitude perturbations would impact the digital TV signal but since digital amplitude modulation is more robust compared to analog modulation, QIM introduces less distortion in case of TV compared to AM.

The FM signal is the most robust among evaluated signals with less than 8% for four quantization levels since amplitude perturbations introduced by QIM have minimal impact on the frequency modulated FM signal.

Distortion in the Frequency Domain. Next, we evaluate impact of embedding data using QIM in the frequency domain. Fig. 3.8 shows the spectrum of the baseband FM signal before and after the QIM embedding process. The two signals are passed through pulse shaping low pass filters to comply with spectral mask requirements. Our results show that there is small distortion in the in-band spectral characteristics of the baseband signal which corroborate the time domain distortion analysis. The out of band frequency components for

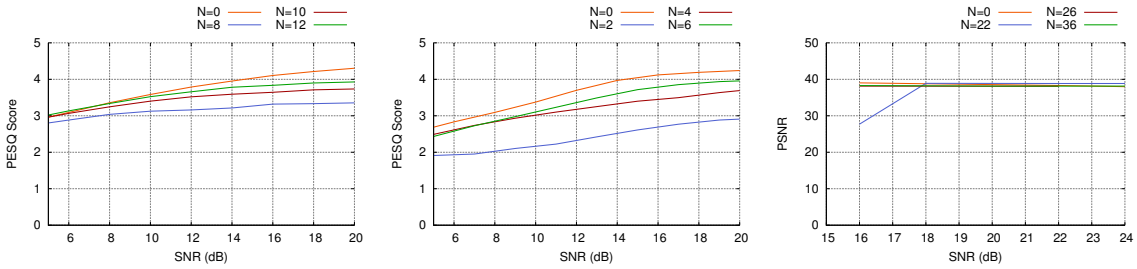

(a) Quality of AM Audio. (b) Quality of FM Audio (c) Quality of TV Video

Figure 3.9: Evaluation of the quality of the host signal multimedia after QIM embedding.

both before and after QIM embedded baseband FM signal are atleast 35 dB below the main lobe thereby having minimal impact on any side channels.

Impact on Host Signal Multimedia. The next step is to translate the distortion to the quality of the multimedia audio and video signal carried by AM, FM and TV signals. This is the key to understanding how QIM impacts the information carried by the host signal. We use the Perceptual Evaluation of Speech Quality (PESQ) metric to quantify the quality of demodulated audio from AM and FM signals. The results model a mean opinion score (MOS) that ranks the quality of speech from 1(bad) to 5(excellent). As a reference, a PESQ ≥ 1 is sufficient for human hearing [76]. We evaluate the PESQ as function of the SNR of the host signal and distortion introduced by scalar DC-QIM technique. Fig. 3.9(a) shows the audio quality of an AM signal as a function of SNR of the host signal and the number of quantization levels. We can see that even at the lowest SNR of 6 dB, PESQ is greater than 2.8 which is sufficient for most applications. As we increase the SNR and number of quantization levels, the audio quality improves. The results confirm that QIM has minimal impact on the audio quality of AM signals. We perform similar analysis for FM signal and Fig. 3.9(b) shows the audio quality of a demodulated QIM embedded FM signal at different host signal SNR and quantization levels. Since the FM signal is more robust to QIM, we evaluate the scalar DC-QIM technique for 2-6 quantization levels (42%-2% distortion) in

the FM baseband. However, due to the robustness of frequency modulation, even at the worst-case SNR of 6 dB and 42% distortion, PESQ is close to 2, which is satisfactory.

Finally, we evaluate the video quality of the QIM embedded TV signal, using the peak SNR (PSNR) metric which is the ratio of maximum possible power of a signal to the power of the distorting noise [195]. The PSNR is computed as follows

$$PSNR = 10log_{10}(\frac{(max(s))^2}{D_s}) \tag{3.14}$$

where D_s is distortion defined in Eq. 3.5. For typical applications, PSNR values between 20-25 dB are acceptable for wireless systems [195].

To evaluate the PSNR of the video signal, we first extract the video by demodulating the QIM embedded TV host signal. We use a software defined radio (USRP X300) to re-transmit the QIM embedded TV signal at different SNR to a TV tuner card by Hauppauge to recover the video. The TV signal was embedded with 20-36 quantization levels which translates to 0.9%–0.3% distortion in the TV baseband signal. In Fig. 3.9(c) we plot the PSNR of the video output of the TV tuner card as a function of SNR of the host TV signal and number of levels. The PSNR of the recovered video was around 34 for majority of the cases expect for the lowest SNR of 16 dB at the highest distortion. However, even the lowest values of 28 dB PSNR is acceptable for most applications. Our analysis considers TV signals above an SNR of 16 dB, a constraint placed by the sensitivity of the TV tuner card. The TV tuner was only able to play video from original distortion free TV signal above an SNR of 16 dB which placed the limit on the SNR of the TV signal evaluated in this work.

To give readers an intuition about the quality metric used in the evaluation, we created a composite video of audio and video clips for AM, FM and TV host signals for different SNR and distortion in the host signal which can be found at the following web link:

https://youtu.be/gKn09ctlFMA.

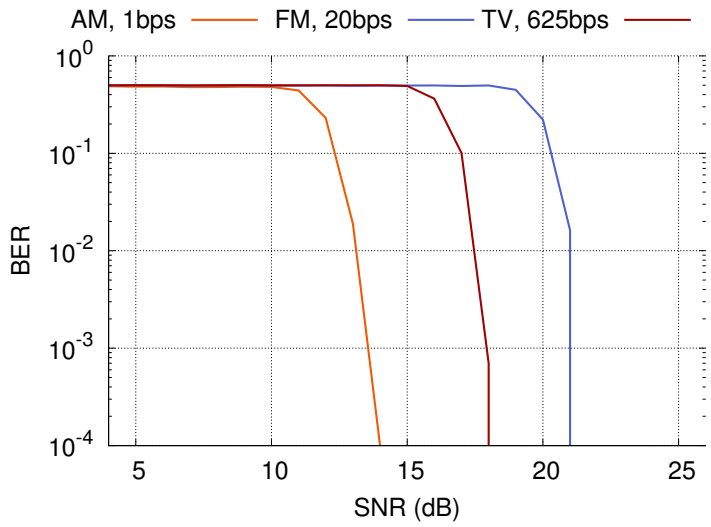

Figure 3.10: A comparison of performance of all three host signals.

3.5.2 Performance of the QIM Embedded Message

The next step is to evaluate the performance of the QIM embedded message signal. We only consider scenarios where the distortion and impact on the host multimedia is within acceptable bounds described in Section 3.5.1.

In Section 3.4 we showed that the information carrying capacity of QIM, like any communication system, is directly proportional to the bandwidth of the host signal. To have a fair comparison across different signals, we normalize the embedded message data rate to the bandwidth of the host signal and evaluate performance for all three host signals. Specifically, we embed at the rate of 1 bps for 10 kHz bandwidth AM signal, 20 bps for 200 kHz bandwidth FM signal and 625 bps for 6.25 MHz bandwidth TV signal. Fig. 3.10 shows the BER of the embedded message using 22 level scalar DC-QIM for the three host signals which translates to 1%, 0.3% and 0.7% distortion respectively in the AM, FM and TV signals. We can see that there is a 4 dB difference between AM and TV and a 7 dB difference between AM and FM. This can be attributed to the fact that for the same number of levels, the AM signal experiences 0.35% more distortion compared to TV and 0.76% more distortion compared to

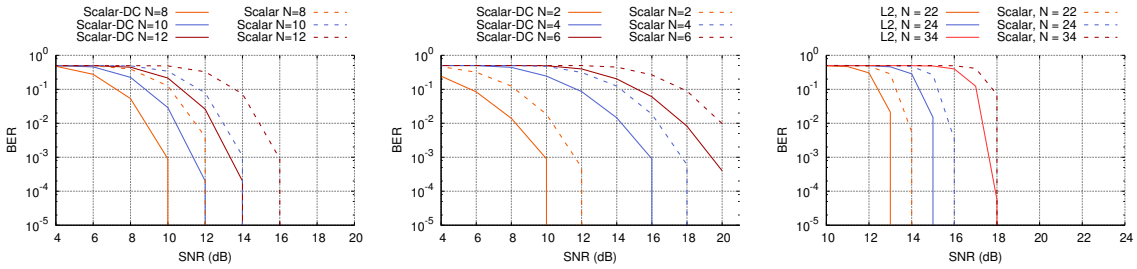

(a) Scalar and Scalar-DC QIM on AM signal.

(b) Scalar and Scalar-DC QIM on FM signal.

(c) Scalar-DC and Lattice-DC QIM on TV signal.

Figure 3.11: We evaluate the performance of embedding message using QIM on each of the three host signals.

FM. Since distortion is the embedded signal, a higher distortion translates to higher signal strength for the embedded message signal and better performance. We empirically note an approximate 1 dB increase in performance for every 0.1% increase in distortion.

Next, we individually analyze each of the three host signals by evaluating the BER of embedded message as a function of the SNR of the host signal and the signal strength of the embedded message. We evaluate the AM signal for 8-16 levels at 200bps embedded message rate for both scalar QIM and scalar DC QIM. Fig. 3.11(a) shows that the BER decreases with decrease in number of levels which translates to an increase in distortion of the host signal or signal strength of the embedded signal. Specifically, for every decrease in two quantization levels, there is a 2 dB increase in the performance. Finally, the distortion compensation technique improves the performance by about 2 dB and this is true for all host signals and both scalar and lattice DC-QIM methods.

For the distortion tolerant FM signals we embed message at 20 kbps and use 2–6 quantization levels which translates to a distortion of 42% to 4%. Fig. 3.11(b) shows the performance of scalar and scalar DC QIM for the host FM signal. We observe a 2 dB increase in performance for every unit decrease in number of quantization levels.

Finally, we evaluate the TV signal for both scalar and lattice DC-QIM techniques. Since

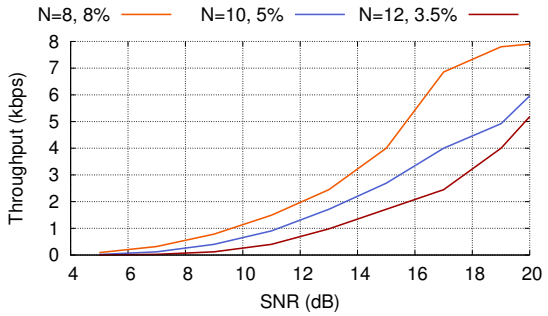

(a) Achievable throughput with low distortion in host signal (AM signal).

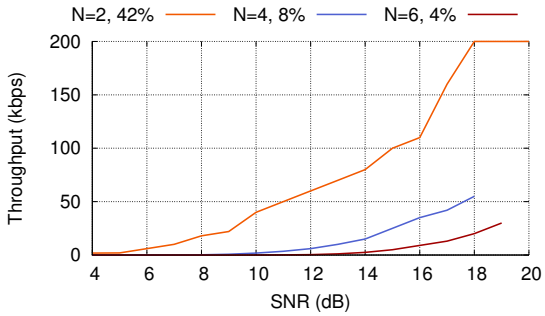

(b) Achievable throughput with high distortion in host signal (FM signal).

Figure 3.12: We evaluate the achievable throughput using Scalar-DC QIM with high and low distortion in host signals.

TV signals are susceptible to distortion, we embed messages at 250 bps and use 22–48 quantization levels which translates to a distortion of 0.7% to 0.1% in the host TV signal. Fig. 3.11(c) plots the performance of an embedded QIM message in a host TV signal and we can see that performance increase with decrease in number of levels or increase in distortion. Additionally, lattice DC-QIM outperforms scalar DC-QIM by about 2 dB.

3.5.3 Achievable Throughput

We analyze achievable throughput for high and low distortion in the host signal. For low distortion in a host signal, we consider the 10 kHz wide AM radio signal and embed messages using scalar DC-QIM method at different data rates using 8–16 quantization levels (8%–3.5% distortion). Fig. 3.12(a) shows the throughput as a function of SNR for different quantization levels. We can see that QIM embedded message achieves the maximum throughput of 8 kbps for SNR of 20 dB for 8 quantization levels (8% distortion). Every decrement in two quantization levels increases the throughput by a factor of 1.5 and every 2 dB increase in SNR increase the throughput by an average factor of 3. We can scale these results for

FM and TV signals by respectively multiplying the AM data rates by 20 and 625. For a high distortion case, we consider the 200 kHz wide FM radio signal and embed messages using scalar DC-QIM method with 2–6 quantization levels (42%–4% distortion). We plot throughput as a function of SNR for different quantization levels in Fig. 3.12(b). We achieve a maximum throughput of 200 kbps.

3.6 Related Work

Our work is related to recent efforts in spread spectrum watermarking of RF signals [208]. Spread spectrum is a promising technique because it is robust against interfering noise. However, it is a linear method and is susceptible to host signal interference. On the other hand, QIM is a non-linear techniques which is aware of the host signal and therefore efficiently manages host signal interference, resulting in better overall performance. Wireless QIM is also related to inter-protocol communication techniques which enable communication between IoT devices using different wireless standards. This is especially beneficial in the crowded 2.4GHz spectrum where devices with a software modification can using existing hardware to communicate between different devices employing different standards. For instance, the WiZip system uses the presence and absence of packets to encode information for transmission from a Wi-Fi device to a ZigBee device [67]. FreeBee uses variance in timing of regular Wi-Fi beacons to transmit information to ZigBee devices [96]. B^2W^2 uses presence and absence of packets to enable BLE to Wi-Fi communication while concurrently supporting existing Wi-Fi and BLE communication [29].

All of these techniques are promising, but they are limited to low data rates, are short range and spectrally inefficient since Wi-Fi and ZigBee use drastically different bandwidths. Instead, the Wireless QIM technique can achieve data rates up to 8 kbps with only a 10 kHz wide host AM signal which is 52, 470, and 5.5 order of magnitude higher when compared to WiZig, FreeBee, and B^2W^2 which use a significantly wider bandwidth (20 MHz) signal.

3.7 Conclusion and Future Work

We have introduced Wireless QIM technique to embed information into existing signals and communicate with smart devices while having negligible impact on the host signal. We have demonstrated communication at up to of 8 kbps for low bandwidth distortion sensitive AM signals and 200 kbps for higher bandwidth, but distortion resilient FM signals. To the best of our knowledge, this is the first work to use QIM technique to embed messages into wireless signals. We believe Wireless QIM presents a new and exciting opportunity for the radio/TV/cellular providers to enable smart cities and IoT applications by reusing their existing infrastructure and deliver additional value with connectivity at zero spectrum overhead.

In this paper, although we have evaluated data rate and reliability of the embedded message, we haven't explored errors correcting codes. Development of error correcting codes for Wireless QIM to achieve data rates closer to the theoretical capacity of the channel is an exciting avenue for future research. Finally, this work was focused on downlink communication and in the future work we will extend the Wireless QIM technique for uplink communication as well.

3.8 Acknowledgements

This work is supported in part by NSF award CNS-1305072 and a Google faculty research award.

Chapter 4

GLAZE: OVERLAYING OCCUPIED SPECTRUM WITH DOWNLINK IOT TRANSMISSIONS

4.1 Introduction

Over recent years there has been a rapid growth in the deployment of IoT devices. These devices are used in home, cities, even large industries such as agriculture. However, deploying current solutions for IoT systems at large scale has two key drawbacks: cost and maintenance.

Mainstream IoT solutions require the setup of additional equipment, such as a base station to transmit data. Setting up the antenna, tower hardware, tower space and management for the base station quickly increase the cost of the IoT deployment. Furthermore, many of these IoT devices have high power consumption and rely on single use batteries, making it difficult to maintain at a large scale. Imagine deploying a large network, tens of hundreds of devices, to collect data. These devices would require routine maintenance just to replace batteries, negatively impacting productivity and increasing overall cost.

Recent breakthroughs in ambient backscatter technology make improvements in reducing the cost and power consumption of IoT systems by enabling both peer-to-peer and uplink communication for IoT devices by utilizing existing wireless signals. For example, devices can communicate peer-to-peer by backscattering TV broadcast signals or Wi-Fi signals can be modulated by a Wi-Fi tag to send uplink data to a base station [110] [93]. Other examples such as FM backscatter, demonstrate that devices can communicate by backscattering FM broadcast signals as well [203]. These techniques have been used to develop devices that can even operate at ultra low power, allowing them to be completely battery-free [160] [131]. While these ambient backscatter solutions are very promising they still lack reliable downlink communication to enable a fully end-to-end IoT system.

To enable downlink communication with existing ambient backscatter solutions, the setup of an additional base station would be required, which could not rely on backscatter techniques. For instance, ambient backscatter tags that rely on TV broadcast signals for peer-to-peer or uplink communication would need a separate base station for downlink communication. This is true for FM or Wi-Fi signals as well, which in most cases would quickly drive up the cost. A more promising solution that has been proposed for Wi-Fi backscatter, is using the presence and absence of Wi-Fi packets to convey downlink data [93]. However, this solution is constrained to Wi-Fi systems and requires modifications to the Wi-Fi access point.

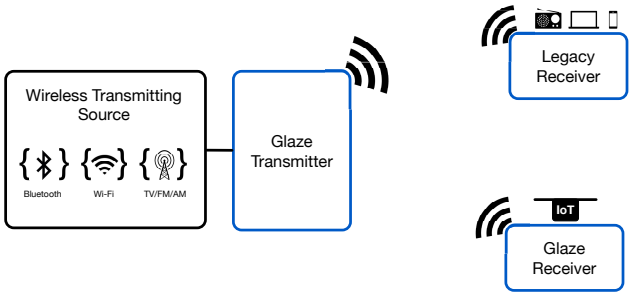

Figure 4.1: **Glaze System.** The Glaze module connects to the antenna port of a preexisting wireless transmitter to overlay data. The transmitted signal is then received by a Glaze receiver and legacy receivers.

In this work we bridge the gap by introducing a new system called Glaze, a general paradigm for downlink communication that can be applied to any preexisting wireless signal to enable downlink transmissions to IoT endpoint devices. Since TV transmissions already exist, they could carry downlink data for IoT endpoint devices. Similarly, FM transmissions, Wi-Fi or Bluetooth signals could carry this information as well. If this data can be decoded at low-power at the endpoint device, it can be used to enable an end-to-end ambient backscatter communication system. However, there are three main challenges in directly reusing existing wireless signals and infrastructure. First, requiring backscatter devices to have TV, FM, Wi-Fi or other receivers would be very power consuming and expensive. Second, it is non-trivial to modify the hardware, firmware, or encoding pipeline of existing transmitters to overlay additional IoT data. Lastly, we must not degrade the performance of the existing communication channels. We solve these challenges as follows.

To transmit additional data using existing wireless signals, for example Wi-Fi, without

requiring a Wi-Fi receiver at the backscatter device to decode the signal, we overlay a low rate data transmission on top of the existing wireless signal. The IoT backscatter device then needs only a very simple low-power radio to decode the low rate Glaze downlink message.

To avoid making changes in the transmitting radio hardware or firmware, we propose connecting a Glaze module to the antenna port of wireless radios to overlay data on top of the RF signal. For instance, a Glaze module can be connected to the antenna port of a Wi-Fi access point to encode downlink data transmissions for ambient backscatter devices, while Wi-Fi packets are still being transmitted and received by legacy receivers, as shown in Figure 4.1.

Lastly, to minimize interference towards existing communication channels our data embedding technique stems from the domain of digital watermarking and information hiding. There have been several proposed techniques to hide messages in media such as images or audio, with minimal distortion impact. One example is additive spread spectrum modulation, where many small increases and decreases in pixel brightness are introduced in an image to embed additional data while being imperceptible to the human eye [180]. We extend some of these concepts to wireless communication to send additional data while introducing minimal interference. We propose an encoding scheme that introduces small amounts of attenuation to the existing wireless signals to overlay additional data. The key insight to this approach is that most wireless receivers are built to tolerate a small amount of noise in the wireless signal, allowing us to convey additional data to the IoT endpoint devices.

In the remainder of this paper we will explain how downlink communication can be enabled for ambient backscatter systems by leveraging existing wireless signals and infrastructure. Specifically, we go over the following key contributions,

- *We present Glaze, a data overlay system for preexisting wireless signals that enables downlink communication for ambient backscatter systems.* Glaze uses a data encoding scheme that can be applied to most wireless signals to convey additional data to IoT endpoint devices, without disrupting existing communication links. The Glaze mod-

ule can be added to an existing wireless transmitter without modifying hardware or firmware components.

- *We show that Glaze can be applied to packetized data networks, such as Wi-Fi.* Such networks do not continuously transmit data and the module would need to predict the Wi-Fi downlink activity to successfully embed data.
- *We design and implement hardware prototypes for the Glaze module and receiver.* The Glaze module and receiver operate at low-power and are compatible with existing backscatter uplink solutions. Specifically, we show that downlink communication can be achieved at rates up to 10 kbps and distances of up to 3.5 meters for Wi-Fi signals. Moreover, we demonstrate that with FM and TV broadcast signals the Glaze system can achieve rates of up to 10 kbps at distances of 9 meters when the transmit power is set to 20dBm.

4.2 Challenges

Ambient backscatter communication is a promising approach to enable communication for IoT systems because it is a low-cost and low-power solution, however current implementations are still limited to uplink or peer-to-peer communication. While most IoT data is uplink, for instance sensor data sent to the edge or base station, all control and signaling data is downlink. This includes information such as the reporting rate (e.g. how frequently the sensor should send data), updates (e.g. a latest version of the firmware), acknowledgements, diagnostic data, among others. In order to enable an end-to-end IoT system using ambient backscatter communication techniques there must be a downlink channel as well. For example, ambient backscatter leverages existing infrastructure to send data from tag to tag. However, sending any downlink data still requires dedicated infrastructure. With these limitations, IoT endpoint devices are unable to take full advantage of the cost benefits offered by backscatter techniques since they would still need an expensive IoT radio receiver for downlink communication.

Glaze solves this problem by removing the need for dedicated infrastructure or high power

receive radios and instead, overlays downlink IoT transmissions on existing RF signals that can be received by backscatter devices. Furthermore, because of its modular design Glaze can be used with any wireless transmitter without requiring modifications to the hardware or firmware. To bring Glaze to fruition there are several challenges that need to be addressed and can be broken down into two categories: broadcast and data networks.

4.2.1 Broadcast Networks

Broadcast networks account for signals such as digital TV or FM and AM radio, that stream data continuously. Applying the Glaze system to such signals is challenging for two reasons. First, it is inevitable with any data overlay technique that the embedded message will cause distortion towards the wireless signal. While small amounts of distortion are acceptable, we do not want to exceed a threshold that will cause channel degradation. The Glaze system must adapt to these constraints depending on the wireless signal and optimize the performance of the embedded message. In particular, for broadcast signals, the modulation scheme used by the broadcast transmitter plays a key role in defining the data overlay technique parameters.

The modulation scheme of the wireless signal used to encode its own data will also determine how resilient the signal will be to distortion caused by the data overlay technique. We mentioned previously that the Glaze system introduces small amounts of attenuation to embed data on top of the wireless signals. This can be viewed as introducing small shifts in amplitude, which means a signal that uses frequency modulation will be more resilient in comparison to a signal that uses amplitude modulation. A specific example would be that a digital TV broadcast signal would be more susceptible because it uses 8-VSB (a form of amplitude modulation) in comparison to analog FM broadcast signals, which would be much more robust.

Lastly, since the major appeal of ambient backscatter communication techniques is the capability of being a low-cost and low-power solution for IoT systems, the Glaze system must also maintain these characteristics. In particular, the Glaze receiver needs to operate at very

low-power and be compatible with existing ambient backscatter uplink solutions. To enable this, we must design hardware that can decode the additional embedded data without the need for power consuming component such as ADCs or RF oscillators. Moreover, we need a low complexity data overlay scheme that is also resilient to channel noise in order for the Glaze receiver to decode the transmitted data without expensive computation.

4.2.2 Data Networks

Integrating the Glaze system with data networks is even more challenging. Data networks account for wireless signals such as Wi-Fi or BLE, which instead of continuously transmitting, send data intermittently. Take Wi-Fi as an example, which uses packetized protocols to transmit data to clients. In order to reliably overlay data on top of Wi-Fi or other signals using packetized protocols, the Glaze system must be aware of incoming packet transmissions and the duration of the packets. These two insights will determine when the Glaze module should begin to overlay data and if an entire Glaze packet can be overlaid during the duration of the packet transmission. Specifically, if the packet transmissions are very short, then the Glaze system might need to overlay data across several packets. Thus, not only do we need to know the aforementioned parameters, we also need an adaptive data overlay scheme to accommodate the varying signal characteristics.

Another challenge with applying Glaze to data networks is that the distance between the wireless transmitter and receiver plays a role in how much distortion can be introduced to a wireless signal. In the case of Wi-Fi, clients can be very close or much further away from the Wi-Fi access point (AP). A client that is further away from the AP cannot tolerate as much distortion introduced to the signal in comparison to a client that is closer in proximity. Thus, the Glaze system must be aware of an APs associated clients along with their proximity to adapt the data overlay technique and in turn avoid channel degradation.

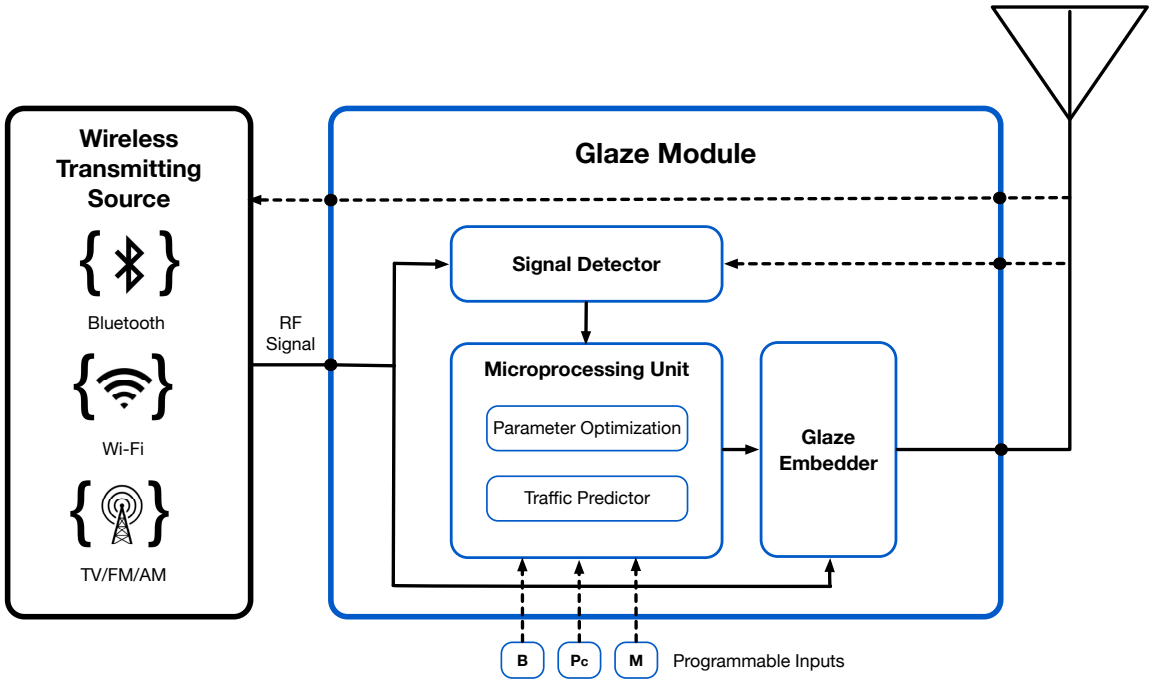

Figure 4.2: **Glaze Module Design**. The Glaze module connects to the antenna port of wireless transmitters to overlay additional data.

4.3 Glaze System Overview

Glaze is a novel communication system that enables a new channel of communication for ambient backscatter devices by overlaying data on existing wireless signals. By utilizing existing infrastructure and spectrum, we envision Glaze to be used as a low-power and low-cost modular solution for downlink backscatter communication. As shown in Figure 4.1, Glaze has two key components: a Glaze module and receiver. The module is connected to the antenna port of a wireless transmitter, which can be anything from a Wi-Fi AP to FM or TV broadcast tower. It overlays data on top of wireless signals and outputs a composite signal to be transmitted and received by legacy receivers and Glaze receivers. In the remainder of this section, we describe the components of the Glaze module and receiver necessary for the system to be compatible with both broadcast and data networks.

4.3.1 Glaze Module

The Glaze module overlays data by introducing small amounts of attenuation to the wireless signals. For example, it can modify a Wi-Fi transmission by changing the signal in a way that Wi-Fi receiver's performance is not affected, yet the signal carries additional data for ambient backscatter devices.

Achieving the above is non-trivial. One approach to add additional data to an existing signal is to modify the firmware of existing radios. For instance, a Wi-Fi transmitter could be modified to carry additional information in the guard bands or pilot tones. Although this can be implemented on software defined radio platforms, implementing such techniques on commercial off-the-shelf Wi-Fi radios would require hardware changes. Furthermore, modifying guard bands and pilot tones might affect the reception of existing signals. Instead, we use a different approach in Glaze, which does not require hardware or firmware changes to the radio platforms. Our system operates as a module that is added to the antenna ports of existing transmitters. It does not need to decode the original signal, yet it is able to overlay data to enable downlink transmissions for IoT endpoint devices. This is accomplished by using the following components in the Glaze module:

- **Signal Detector:** The Glaze module should only overlay data when there is a signal being sent by the wireless transmitter. A FM transmitter always transmits a signal, while a Wi-Fi transmitter might only transmit a packet when it has something to send. Instead of trying to decode the transmissions, the Signal Detector monitors for incoming signal transmissions and activates the rest of the Glaze module pipeline.
- **Parameter Optimizer:** Before overlaying the data, the Glaze system must define the data encoding parameters to avoid degrading the original wireless signal. The Parameter Optimizer takes into account the modulation scheme, channel power and bandwidth of the wireless signal. Moreover, it considers the distance between the wireless transmitter and its associated clients for data network scenarios. All of these inputs are used to optimize the data embedding parameters. For instance, how much

attenuation can be introduced to the wireless signal or the rate of embedding.

- **Traffic Predictor:** Since some wireless transmitters, for instance Wi-Fi APs, use packetized protocols, the Glaze module must be aware of the start of a packet transmission and its duration. The Signal Detector can detect the packet transmissions, however having some knowledge of packet duration is needed because the Glazed data might not always fit in a single transmission. For example, if a low bit rate is being used to overlay the additional data it would need to be embedded across several Wi-Fi packets. The Traffic Predictor is used to estimate the likelihood of signal transmissions for a select unit of time. In turn, the Glaze module can make intelligent decisions when overlaying data, such as optimizing data rate and choosing the ideal time to begin embedding.
- **Glaze Embedder:** Finally, the data is overlaid by the Glaze Embedder. It overlays the Glaze data on the RF signal based on the parameters determined by the Parameter Optimizer. At a high level, the Glaze Embedder introduces small perturbations to the RF signal in the form of attenuation, but in a way that does not degrade the reception of the preexisting wireless signal. All the while, being able to be received by the IoT endpoint devices.

4.3.2 Glaze Receiver

The Glaze receiver needs to be very low-power and yet should be able to receive and decode the Glaze signal transmissions. Typically in wireless systems, data is encoded in the frequency domain and in turn, receivers require power hungry components to perform high complexity tasks needed to decode. One approach to avoid power consuming components would be to overlay data on the baseband of the wireless signal in the time domain. This would enable the Glaze receiver to decode without needing power consuming hardware components to perform tasks such as Fast Fourier Transforms. However, this would instead increase complexity at the Glaze module. That is, the module would need to downconvert the wireless signal to the baseband frequency, overlay the Glaze data, then upconvert back

to the RF frequency. While this approach is possible, it would not be a very practical implementation. Instead, overlaying data directly on the RF signal avoids these changes to the Glaze module and enables the Glaze receiver to have a low-power design.

4.4 Glaze Physical Layer

To realize the Glaze system described in the previous section, we need to implement the network physical layer. This includes two key components. First, defining an encoding and decoding algorithm used to overlay data on the existing wireless signals. Second, defining a scheme to detect Glaze packets, perform synchronization, and decode data at the receiver.

4.4.1 Modulation and Bit Encoding

We define Glazing, a data embedding technique for RF signals. It introduces small amounts of attenuation to RF signals to encode additional data. The technique is built upon the key intuition that most receivers are designed to withstand some amount of noise in the signal. As long as we can convey data in the noise we can send downlink transmissions to ambient backscatter devices, and as long as we add very little noise to the RF signal we can reduce the impact on receiver performance. To understand the Glazing technique, we first start by defining the data embedding function,

$$X(s) = s * \alpha_m, \quad \text{where,} \quad \alpha_m = \begin{cases} \frac{1}{\sqrt{10^{\frac{\Delta}{10}}}}, & \text{if } m = 0 \\ 1, & \text{if } m = 1 \end{cases} \tag{4.1}$$

where s is the RF host signal (eg. Wi-Fi, TV, or FM) and Δ is the amount of change introduced to the signal in the form of attenuation to overlay a data message, m. Specifically, if we want to overlay a 0-bit then Δ will equal the maximum amount of attenuation allowed for any particular host signal. Moreover, if we want to overlay a 1-bit then no change will be made and Δ=0. The maximum amount of attenuation allowed for data embedding can be determined by taking into account the minimum signal strength seen at an associated client device, P_c and the typical client receiver sensitivity, P_{limit}, for any particular wireless

signal and protocol. Thus, $\Delta = P_c - P_{limit}$. Furthermore, to improve the robustness of the overlaid data we implement a Manchester encoding scheme where a 0-bit is represented as a zero to one transition and vice versa for a 1-bit. Lastly, to decode the overlaid data the Glaze receiver compares the received signal with and average of itself to distinguish between two data levels.

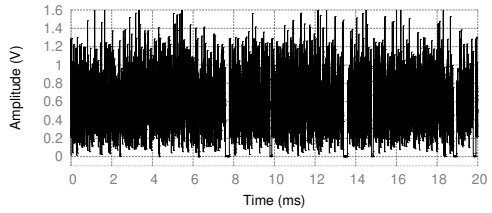

(a) **Wi-Fi Packet Transmissions.** The envelope of Wi-Fi packet transmissions before applying the Glazing technique.

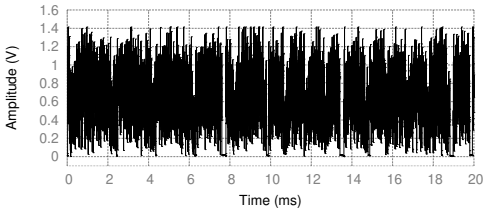

(b) **Glazed Wi-Fi Packet Transmission.** The envelope of Wi-Fi packet transmissions after applying the Glazing technique.

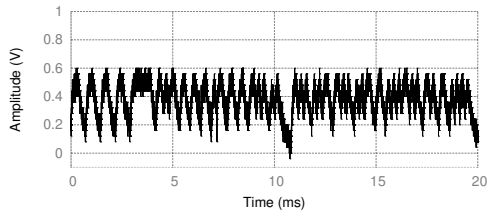

(c) **Signal at the Glaze Receiver.** The signal seen at the Glaze receiver after passing through an envelope detector.

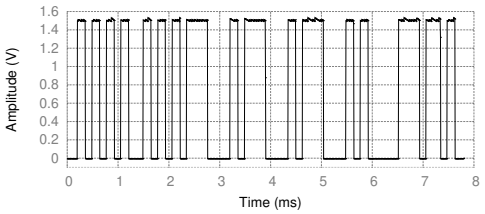

(d) **Decoded Glaze signal.** Data decoded by the Glaze receiver

Figure 4.3: **Glazing technique.** An example of the Glazing technique applied to Wi-Fi packet transmissions at the Glaze module and decoding at the Glaze receiver.

An example of the above data overlay mechanism is shown in Figure 4.3, where the encoding and decoding scheme are applied to Wi-Fi packet transmissions. First, Figure 4.3(a) shows Wi-Fi packet transmissions over 20ms and Figure 4.3(b) shows those same Wi-Fi packets after applying the Glaze embedding technique with Δ=4dB attenuation. The Wi-

Figure 4.4: **Packet Structure for Glaze Embedding**. The Glaze system uses a short and long training field, followed by data, to overlay packets on top of existing signals.

Fi signal is then transmitted and received at the Glaze receiver. Figure 4.3(c) shows the received signal after passing through an envelope detector. The Glaze receiver decodes the the overlaid data and outputs the estimated message shown in Figure 4.3(d).

4.4.2 Packet Detection and Synchronization

To be able to decode a Glaze packet transmission the receiver must perform a series of tasks to identify an incoming packet, which includes synchronization and extracting important parameters needed for decoding. The Glaze system uses the packet structure shown in Figure 4.4 to perform packet detection and synchronization. There are three main fields: short training field (STF), long training field (LTF), and data. The short training field includes a 7-bit wake-up signal that is a bit sequence of redundant 1s and 0s used to trigger the Glaze receiver to wake-up from sleep mode. The last two bits of the STF carry the bit rate used to encode data, which has four possible values. For instance, in a Wi-Fi scenario the bit rate can be 3bps, 200bps, 5kbps, and 10kbps corresponding to the binary representation 00, 01, 10, and 11. The second field is the LTF, which is a 13-bit Barker code preamble. The preamble is used for synchronization and enables the Glaze receiver to identify the start of the data in the received packet and begin decoding. Finally, the last field is data and can carry up to 25 bits. To enable coexistence of multiple receivers associated with one Glaze module, a device ID field can be added to the packet structure as well. In other words, each device will have a unique ID, which would be included as a header in the packet.

4.5 Glaze System Design

In Section 4.3 and 4.4 we described the Glaze system and the components necessary to overlay data on existing wireless signals. Moreover, we detailed challenges with ensuring there is minimal distortion impact on the wireless signals all the while being able to decode the overlaid data at the IoT endpoint devices. In the following subsections we go over the design of the Glaze module and receiver needed to realize the Glaze system.

4.5.1 Glaze Module Design

The Glaze module consists of several components all needed to enable overlaying data with minimal distortion towards the wireless signal. First, the module must be able to detect a signal transmission to know when to start overlaying data. Second, it needs to optimize the data overlay parameters to ensure that the Glaze data does not interfere with the existing communication channel. Lastly, the module needs to adapt the Glazing technique when overlaying data on wireless signals that use packetized protocols (eg. Wi-Fi). We address each requirement as follows.

- **Signal Detector:** The Signal detector is used to to detect incoming signal transmissions and trigger the remainder of the Glaze module pipeline. This can be achieved by using a simple energy detector. The energy detector is always measuring the signal strength of the wireless signal and when the signal strength exceeds a predetermined power threshold the rest of the Glaze module pipeline is enabled. The Signal Detector can be implemented by using an off-the-shelf energy detector.
- **Parameter Optimizer:** The solution to achieve parameter optimization is twofold. Input parameters such as modulation scheme, channel power, and bandwidth can be programmed in software for the Glaze module. However, we also need to know the distance between the wireless transmitter and associated clients to determine the maximum amount of distortion that can be introduced without disrupting any communication link. We address this by using the energy detector to sense both incoming and

outgoing packet transmissions. The Glaze module uses the output of the energy detector to measure the signal power of incoming packets. Each time the module senses a received packet, it immediately begins to sense for an outgoing packet transmission as well. The key intuition with this approach is that every time a Wi-Fi access point receives a packet from an associated client it will always send an acknowledgement within a fixed time frame. We use this Wi-Fi characteristic to determine if the client is associated with the AP. If so, the signal strength of the received packets are used to determine how much distortion can be introduced when overlaying data.

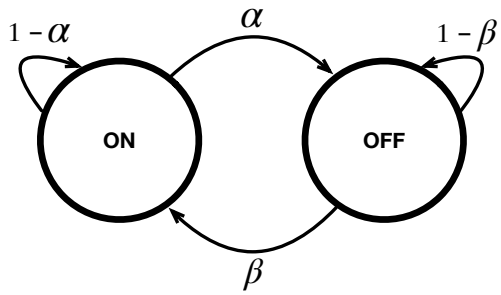

(a) **Markov Model.** A two state Markov model for Wi-Fi traffic prediction. The ON and OFF states represent packets being transmitted or not.

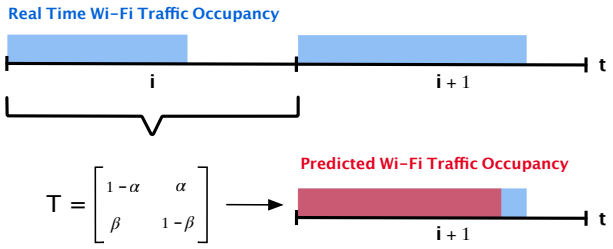

(b) **Occupancy Prediction.** An example of Wi-Fi traffic occupancy prediction using the Markov model.

Figure 4.5: **Wi-Fi Traffic Occupancy Model.** A two state Markov model shown on the left is used to predict short-term Wi-Fi traffic occupancy. On the left, an example is illustrated of how Wi-Fi occupancy prediction is implemented.

- **Traffic Predictor:** As mentioned above the energy detector is used to sense incoming and outgoing signal transmissions. The detection of an outgoing signal transmission is also used to trigger the rest of the Glaze module pipeline when using the Glaze system in data network scenarios. However, challenges with data networks (eg. Wi-Fi) still persist. That is, since Wi-Fi packets are transmitted in bursts, the Glaze module needs to have an idea of when to begin overlaying data and determine the optimal data

embedding rate for any given unit of time. A simple approach would be to embed each bit using a very low data rate (few bps), spanning many Wi-Fi packets, to ensure that there is always enough energy per bit to be decoded at the Glaze receiver. However, this approach would greatly limit the performance of Glaze. If there is a way to know the total Wi-Fi downlink traffic occupancy in a short time frame, then we can optimize the data embedding rate to improve throughput and choose the optimal time to begin overlaying data.

A useful characteristic of Wi-Fi traffic that can help address this challenge is that it is very persistent. That is, typically if Wi-Fi traffic occupancy is high, then it is likely to stay high for a very long time. Similarly, if it is low then it is likely to stay low. This has been modeled in several studies by using the Hurst exponent [16, 37]. At a high level the Hurst exponent is a calculated value used to measure the long-term memory of a time series. The exponent can be used to classify the predictability of a time series. Specifically for Wi-Fi time series data, studies have show that the Hurst exponent ranges from 0.75 to 0.90 and data is considered persistent if it has a Hurst exponent greater than 0.5 [115]. The key takeaway from the mentioned results is that in most cases Wi-Fi traffic has a high Hurst exponent and in turn we should expect very persistent traffic flow. We leverage this trait and develop a promising approach to estimate short-term downlink Wi-Fi traffic occupancy by using a two state Markov model. In fact, using Markov models to forecast time series data has shown to be a promising approach in the wireless networking domain. For instance, improving the estimation of expected transmission count (ETC) in wireless links or a slight variation by using Interrupted Poisson Processes to model Wi-Fi traffic [15] [97].

Figure 4.5 shows the model where the two states represent a Wi-Fi AP transmitting packets (ON) or not transmitting (OFF). Using this method we can predict the likelihood of having a very long OFF period and also the average downlink Wi-Fi traffic per unit of time. With this knowledge, the data embedding rate can be optimized when using the Glazing technique. For example, when the downlink Wi-Fi traffic occupancy

is predicted to be very high we can use a higher bit rate when overlaying data. On the other hand, if the Wi-Fi traffic is expected to be low we wait to overlay data. If the low Wi-Fi traffic occupancy persists for a predetermined time frame (eg. 100ms) then we switch to low bit rate to overlay the next Glaze packet.

To implement the Traffic Predictor, the energy detector is used to determine the total ON and OFF time of the Wi-Fi AP for a select unit of time. The data is then used to generate a transition matrix that gives the probability of staying in each state or transitioning,

$$T = \begin{bmatrix} 1-\alpha & \alpha \\ \beta & 1-\beta \end{bmatrix} \qquad where, \qquad \begin{matrix} \alpha = P_{ON} \rightarrow P_{OFF} \\ \beta = P_{OFF} \rightarrow P_{ON} \end{matrix} \tag{4.2}$$

where α is the probability of transitioning from ON to OFF and β is the probability of transitioning from OFF to ON. With this, the probability of staying in each state is represented by 1 - α or 1 - β. We use the transition matrix, T, to estimate the total amount of Wi-Fi packet transmissions in the next unit of time. To perform the data processing and occupancy prediction a low-power microcontroller (MCU) can be used. Figure 4.5(b) shows an example of this process. The time unit i is used to generate the transition matrix. The MCU keeps track of total ON and OFF times during the select time unit by using a power threshold. If the power of the incoming signal exceeds the threshold, this is considered ON, otherwise it is treated as OFF. After monitoring the incoming signal transmissions, the transition matrix, T, is generated and used to predict the Wi-Fi downlink traffic activity of the next time unit, $i + 1$.

Using the Traffic Predictor, the Glaze module can make intelligent decisions for the data embedding scheme. However, there still may be scenarios where the predicted Wi-Fi traffic occupancy is greater than what is actually seen during the time of data embedding. In these cases, the Glaze receiver might not be able to decode the packet and would result in increased error rates. To supplement the Wi-Fi traffic occupancy model, we also use the energy detector monitor the total amount of energy per unit of

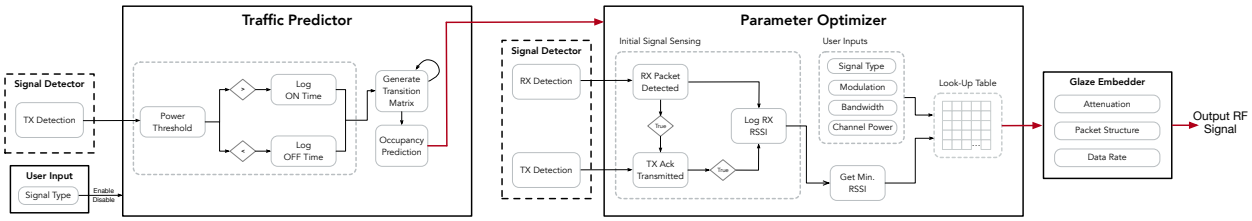

Figure 4.6: **Glaze Module Pipeline.** The Glaze module pipeline goes through a process of performing occupancy prediction and parameter optimization before beginning to embed data.

time during the data overlay process. If it senses that the total occupancy per time unit is less than what was predicted the Glaze module will re-transmit the packet and update the data embedding rate.

- **Glaze Embedder:** The Glaze Embedder is the last component in the Glaze module pipeline, which is used to perform the Glazing technique. The design consists of a microcontroller and digital attenuator. The microcontroller controls the digital attenuator by defining how much attenuation to apply to the wireless signal as well as the data rate. The output of the Parameter Optimizer and Traffic Prediction blocks define how much attenuation is acceptable for the wireless signal at hand and also what data embedding rate should be used for the packet that is to be overlaid.

Figure 4.2 shows the Glaze module implementation. The energy detector is connected to monitor outgoing and incoming signal transmissions. The output of the energy detector goes to the MCU. The MCU uses the power measurements from the energy detector to perform parameter optimization, packet prediction, and adaptive rate embedding. Moreover, the MCU makes decisions using a look up table to pick the best data rate and the attenuation amount that can be used given the input parameters. The look-up table is generated by having prior evaluation of the impact of Glaze embedding for different signal characteristics. In Section 4.7 we demonstrate this by evaluating the impact of the Glaze system on Wi-Fi,

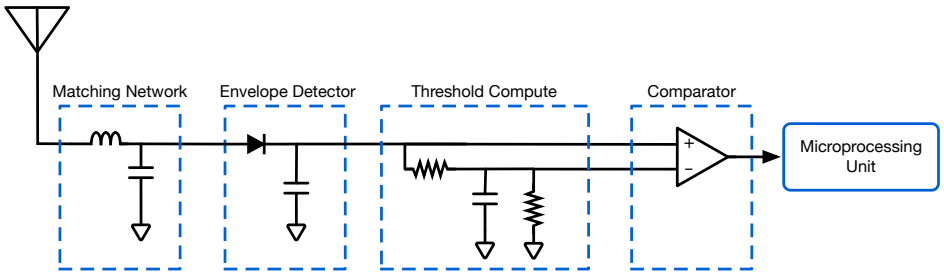

Figure 4.7: **Glaze Receiver Design.** A high level circuit diagram of the Glaze receiver.

TV, FM signals. Lastly, the digital attenuator takes input from the MCU and performs the data embedding. Depending on the wireless signal that the Glaze module will be overlaying data on, certain components of the Glaze module would be disabled. For example, the signal detector, adaptive rate embedding, and occupancy prediction are only necessary for intermittent signals. Figure 4.6 illustrates the the process that each component of the Glaze module performs.

4.5.2 Glaze Receiver Design

Next we describe how the Glaze receiver achieves a low-power design and decodes the Glazed data. Figure 4.7 shows the design of the Glaze receiver, which has four stages: 1) impedance matching 2) envelope detection 3) threshold compute and comparator 4) microprocessing unit.

- **Impedance Matching:** An impedance matching circuit is passive and consists of inductors and capacitors used to tune the receiver to a specific radio frequency. For each wireless system that Glaze is integrated with, the Glaze receiver must be tuned to the appropriate frequency to have optimal performance. For instance, if we are overlaying data on TV broadcast signals then the Glaze receiver should be tuned to one of the channels in the UHF or VHF bands. Similarly, for Wi-Fi or FM the matching network of the receiver would need to be modified.

- **Envelope Detector:** The task of the envelope detector circuit is to remove the high frequency carrier from the received signal (eg. 2.4GHz for Wi-Fi). It is completely passive by using only diodes and capacitors, in turn allowing it to operate at very low-power. We choose specific values for each one of the circuit components to remove the carrier frequency and only the amplitude of the signal remains.
- **Threshold Compute and Comparator:** In principle, an ADC can sample the output of the envelope detector and distinguish between two signal levels. However, to enable a low-power receiver, we should not use the ADC to decode data. Instead, we design a RC circuit to compute a dynamic threshold based on the input signal. Then a comparator compares this threshold with the input signal and outputs the embedded data. The RC circuit is equivalent to a single-bit ADC that has a dynamic threshold based on the input signal.
- **Microprocessing Unit:** Similar to the Glaze module, the Glaze receiver also uses a microcontroller (MCU). In this case, the MCU performs only one task. It takes the output of the comparator and decodes the received data. Then based on the data tasks can be performed. This can be anything from provisioning the IoT endpoint device or triggering a sensor to collect data.

4.6 Implementation

In order to perform an evaluation of the Glaze system we implement prototypes of the Glaze module and receiver described in Section 4.5. The Glaze module is implemented using off-the-shelf components as shown in Figure 4.8(a). Specifically, an LTC5505 Analog Devices energy detector, MSP432 microcontroller (MCU), and an HMC1119 Analog Devices digital attenuator. The energy detector is used to enable the Signal Detector and Traffic Predictor. The microcontroller enables the Parameter Optimizer and coupled with the digital attenuator to enable the Glaze Embedder.

A 4-layer printed circuit board (PCB) was developed to prototype the Glaze receiver

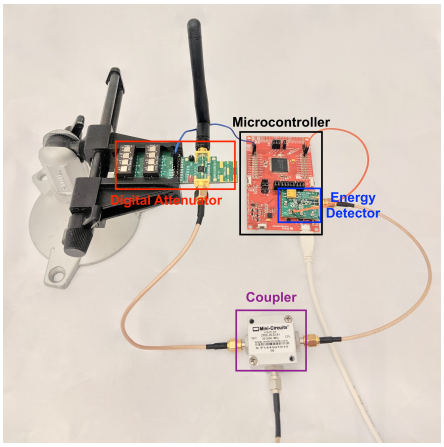

(a) **Glaze Transmitter Prototype.** Off-the-shelf components used to implement the Glaze module.

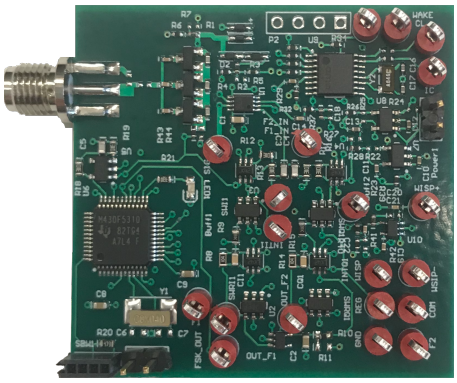

(b) **Glaze Receiver Prototype.** The Glaze receiver printed circuit board.

Figure 4.8: **Prototype Implementations.** The Glaze module and receiver prototype hardware.

design using off-the-shelf components. The receiver is implemented using a MSP430F5310 microcontroller and NCS2200 TI comparator. The prototype receiver board is shown in Figure 4.8(b). To receive data, the MCU needs to sample the output of the comparator, perform cross-correlation to find the preamble, and finally, decode the Manchester encoded data. The maximum achievable throughput with this receiver is limited to 10kbps, a constraint set by the MCU and its power consumption. That is, the MCU needs to sample the output of the comparator, performs an if-else statement and saves the results in the memory. This process can be done for data rates less than 10kbps . However, running the MCU with higher frequency crystals would increase this limit but also the power consumption. By increasing the frequency, the MCU would be able to sample the output of the comparator and save the results in the memory faster and as result the data rate could be higher. The Glaze receiver prototype uses an 8MHz crystal and this limits the data rate to 10kbps.

Fundamentally, there is a trade-off between power consumption and performance. For the

Glaze receiver to operate at low-power, we use less power consuming components such as the completely passive envelope detector and comparator to decode data. However, these design choices also limit the sensitivity of the Glaze receiver to -25dBm. To improve sensitivity, other design approaches can be taken. For instance, we can add a Low Noise Amplifier (LNA) to the Glaze receiver to amplify the input signal. Another approach to improve the sensitivity is by using the ADC of the MCU rather than the comparator and threshold compute components. Since the ADC has better sensitivity or in other words it is able to detect signals with lower signal strength, we would be able to achieve a higher communication range at the receiver. While these approaches improve sensitivity they also increase the power consumption of devices and would not be practical for low-power or battery-free systems. With the current topology of the Glaze receiver, we achieve very low-power consumption. The Glaze receiver consumes only 50μW in listening mode and 800μW when decoding data at 5kbps. This amount of power consumption allows us to have a battery-free design that can harvests the required energy from ambient light or ambient RF signals.

4.7 Results

In this section, we focus on evaluating the Glaze system for three different RF signals: Wi-Fi, TV, and FM. Specifically, we first evaluate the impact that the Glaze system has towards each wireless signal, followed by the performance of the Glaze system itself.

4.7.1 Impact towards the Wireless Signals

We are interested in the impact that the Glaze system has on the original wireless signals to determine the upper bound for introduced distortion and to ensure that we do not degrade any communication links when overlaying Glaze packets. Distortion impact towards Wi-Fi packet transmissions is evaluated by measuring packet error rate for different amounts of attenuation used for Glaze embedding. Moreover, the impact towards FM signals is gathered by measuring the quality of audio. Similarly, for the TV signals the distortion impact is measured by evaluating video quality. For each evaluation we setup bench-top experiments,

where a software defined radio (SDR) is configured as the wireless transmitter, the Glaze module is connected to the antenna port of the SDR to overlay data, and off-the-shelf legacy receivers are used to receive the transmitted signals. The distortion impact for each RF signal was evaluated for Δ = 4, 6, and 8dB attenuation used to embed data. Lastly, the insights gathered from this evaluation are used to formulate the look-up table stored at the Glaze module. The Glaze module uses the look-up table to tell the Glaze embedder how much attenuation can be used to embedded date and a what rate. For instance, the maximum amount of attenuation that can be used per signal type. Moreover, for data networks such as Wi-Fi, how much attenuation can be used depending on the RSSI of associated clients.

Wi-Fi Packet Transmissions

To evaluate distortion impact for Wi-Fi signals we overlay Glaze data on top of Wi-Fi packet transmissions and analyze the PER as a function of RSSI. A USRP X300 is configured as a Wi-Fi AP to send 802.11g Wi-Fi packets with a data rate of 9Mbps and packet length of 346 bytes. We transmit 200 consecutive packets with an interarrival time of 25μs. A Dell laptop PC is used as a Wi-Fi receiver to receive the Wi-Fi transmissions. Finally, the total packets received are used to determine the PER. Figure 4.9(a) shows a plot of PER versus RSSI ranging from -20 to -90dBm with different amounts of attenuation applied. For a fair comparison, we also plot the baseline of the PER for our experimental setup.

Intuitively, we expect that as the RSSI decreases the PER will increase and as more attenuation is introduced to overlay data, we expect the PER to increase further. This trend is evident when the attenuation is Δ=6 dB or greater, however, the distortion impact is still minimal. For instance, when Δ is greater than 4 dB the average deviation from the baseline is approximately 8%. The key takeaway from this experiment is that when the Glaze module uses $\Delta \leq$ 8dB the system does not need to be concerned about degrading the Wi-Fi packet transmissions, even at low RSSI. However, if a larger Δ is to be used, the RSSI of associated clients must be taken into account.

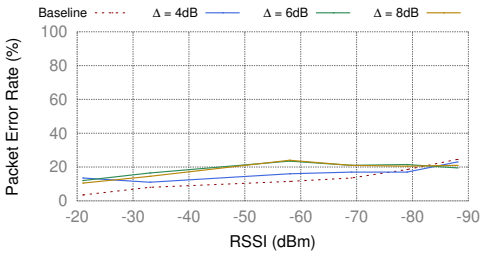

(a) **Impact on Wi-Fi Signals.** The PER as function of RSSI for Wi-Fi transmissions with Glaze embedding.

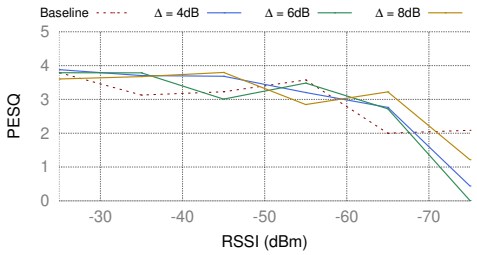

(b) **Impact on FM signals.** The PESQ of demodulated FM audio as function of RSSI with Glaze embedding.

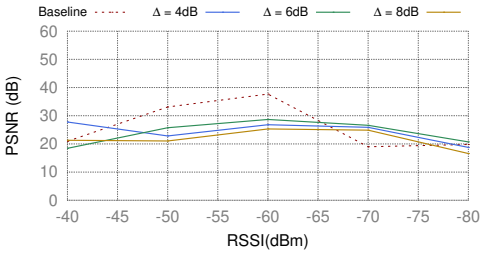

(c) **Impact on TV signals.** The PSNR of video quality from demodulated digital TV with Glaze embedding

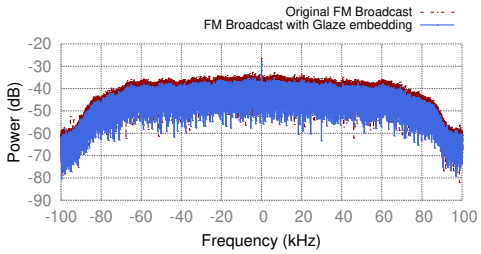

(d) **Impact in the Frequency Domain.** FM broadcast signal spectrum before and after Glaze embedding.

Figure 4.9: **Distortion impact of Glaze on wireless signals.** The distortion impact on Wi-Fi, FM, and TV signals was evaluated for different amounts attenuation used for Glaze embedding.

Analog FM Broadcast Signals

The distortion impact towards FM broadcast signals is evaluated by using the Perceptual Evaluation of Speech Quality (PESQ) metric to measure the quality of audio after data embedding. The PESQ metric models a mean opinion score (MOS) that ranks the quality of speech from 1 being very bad to 5 being excellent. As a reference, a PESQ score greater than 1 is considered sufficient for human hearing [77]. A real FM broadcast signal was recorded using a USRP at 91.7MHz for 15 seconds. The signal recording was retransmitted using

the USRP, with the Glaze module connected to the antenna port to embedded data. The signal is received by a FM radio and the audio was recorded to then calculate PESQ as a function of RSSI. Figure 4.9(b) shows the results for PESQ versus RSSI ranging from -25 to -75dBm. The PESQ results show that FM broadcast signals are resilient towards Glaze embedding for attenuation ranging from Δ=4-8dB, even at very low RSSI. Furthermore, the PESQ maintains a value greater than one up until -80dBm, even using very high amounts of attenuation.

Digital TV Broadcast Signals

We evaluate the impact Glaze has on digital TV signals by measuring the quality of audio using the Peak Signal-to-Noise (PSNR) ratio. PSNR is the ratio of the maximum power of a signal to the power of the distorting noise [195]. Typically in wireless systems a PSNR ranging from 20-25dB is considered sufficient for image quality [195]. Similarly to FM, a real TV signal was recorded at 539MHz for 10 seconds. The signal recording was re-transmitted using a USRP with the Glaze module connected to the antenna port to embedded data. The signal was received using an off-the-shelf TV tuner. We recorded 10 second videos for different Glaze embedding scenarios as function of RSSI. Figure 4.9(c) shows the measured PSNR versus RSSI for a TV broadcast signal. We see the impact that the Glaze embedding has on the video the TV signal was carrying almost immediately. At an RSSI of approximately -45dBm the PSNR begins to deviate from the baseline results. However, when applying attenuation ranging from Δ=4-8 dB the PSNR maintains a value greater than or equal to 20 dB, even at very low RSSI.

Impact in the Frequency Domain

Evaluating the impact that the Glaze system has on an RF signal in the frequency domain in necessary to ensure spectral mask requirements are maintained. Figure 4.9(d) shows an FM signal before and after overlaying data using 4dB attenuation. We can see that the the FM signal with overlaid data is attenuated uniformly across the entire 200kHz bandwidth.

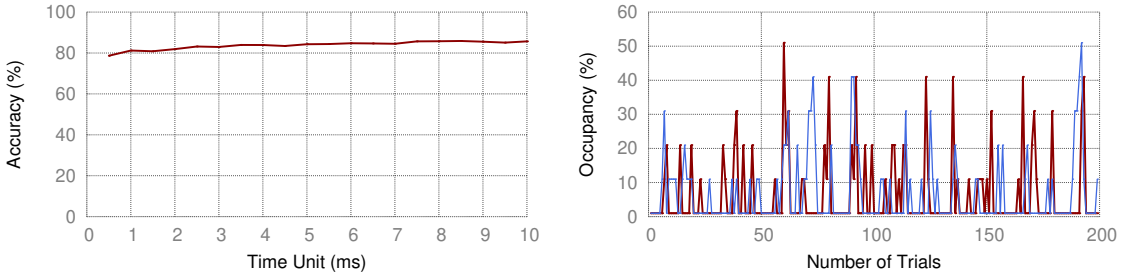

(a) **Wi-Fi Traffic Occupancy Model.** Accuracy of the model for different units of time in milliseconds

(b) **Predicted Occupancy.** An example of predicted occupancy in comparison to the actual Wi-Fi traffic occupancy from the collected data set.

Figure 4.10: **Markov Model Evaluation.** The accuracy of the two state Markov model to predicted Wi-Fi traffic occupancy.

This resulted is expected because the digital attenuator at the Glaze module is designed to apply attenuation across the entire channel bandwidth. Moreover, a minimum of 25dB is maintain between the main and side lobes of the signal.

4.7.2 Glaze System Performance

To evaluate the performance of the Glaze system, we first look at how well the Glaze module can optimize the Glazing technique for Wi-Fi signals. We discussed in previous sections that the module can adapt the bit rate for data embedding and optimize parameters to achieve the best performance. To gather the feasibility of this implementation, we first evaluate the performance of the two state Markov model to predict Wi-Fi traffic occupancy and discuss how the Glaze module can use these inputs, among others, to optimize performance. Lastly, we evaluate the performance of the Glaze receiver. In particular, we are interested in the bit error rate at varying distances from the Glaze module. Moreover, we evaluate the bit error rate for the overlaid Glaze data when using different amounts of attenuation to overlay data. In the following sections, we detail the results of each evaluation.

Modeling Wi-Fi Traffic Occupancy

A two state Markov model is used to predict the total Wi-Fi traffic occupancy per unit of time, as described in Section 4.5.1. These insights can help the Glaze module determine the best bit rate per packet and in turn enable adaptive rate embedding. To evaluate the accuracy of the model, we collect Wi-Fi traffic data in a home setting where there is a single Wi-Fi AP using 802.11b/g/n protocols and four active clients. The clients are all performing different tasks ranging from simple web browsing to video or music streaming. We use the Wi-Fi traffic data to generate transition matrices and make predictions for expected Wi-Fi traffic occupancy in the next time frame. For example, the first 5ms seconds of Wi-Fi traffic data is used to generate a transition matrix, and this set of probabilities are used to predict the total Wi-Fi traffic occupancy in the next 5ms and so on. To evaluate the model we run 7500 trials and compare the predicted occupancy to actual occupancy in the Wi-Fi traffic data set collected and calculate the accuracy percentage.

Figure 4.10(a) shows the accuracy of the model when predicting occupancy for units of time ranging from 0.5-10ms. Furthermore, Figure 4.10(b) shows an example of the actual occupancy and predicted occupancy for a 10 ms unit of time and for 200 trials. The accuracy of the model ranges from 80-85%, where predicting occupancy for very short units of time, for instance 0.5ms, is not as accurate in comparison to longer units of time, 5-10ms. Intuitively, it makes sense that units of time that are 0.5ms or less have lower accuracy. The reasoning is that typically the duration for Wi-Fi packet transmissions are very short. For instance, high throughput protocols (eg. 802.11g/n/ac) have packet durations that are in the μs range. In this case, if we continue to decrease the time unit used for the occupancy model then the accuracy would eventually be equivalent to a coin toss. In other words, decrease towards 50%.

These results show that using the two state Markov model to predict short-term Wi-Fi traffic occupancy can help the Glaze transmitter choose the optimal parameters to begin embedding and also the best data rate to use for any given unit of time. Figure 4.11(a)

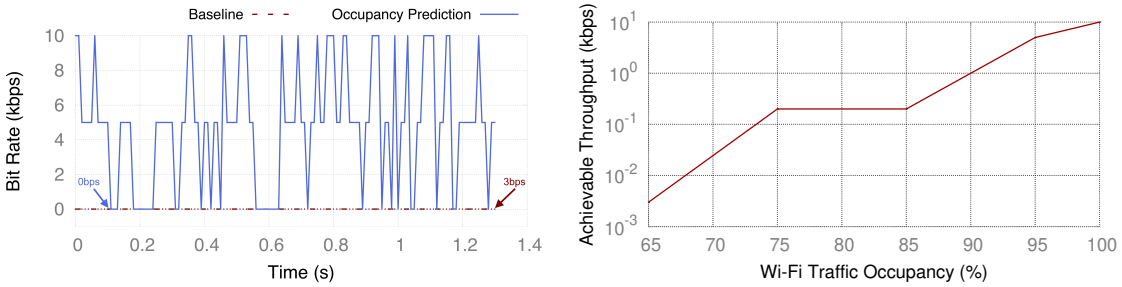

(a) **Adaptive Rate Embedding.** The performance improvement of Glaze when using the Traffic Predictor

(b) **Wi-Fi Achievable Throughput.** The maximum achievable throughput for the Glaze system as a function of downlink Wi-Fi traffic occupancy.

Figure 4.11: **Achievable throughput for data overlay on Wi-Fi transmissions.** The achievable throughput for Glaze embedding for Wi-Fi scenarios using adaptive rate embedding.

shows an example of bit rate as a function of time with and without occupancy prediction for real Wi-Fi traffic. If occupancy prediction is enabled at the Glaze transmitter the bit rate used to overlay data can be adapted to higher rates when there is high occupancy. However, without having any knowledge of the channel occupancy the Glaze transmitter would need to fall back to the baseline data rate of 3bps.

Achievable Throughput

In order for the Glaze module to optimize data overlay parameters, it also needs to know the achievable throughput for different amounts of Wi-Fi traffic occupancy during a select unit of time. To gather this information we applied the Glaze data overlay technique to Wi-Fi transmissions for occupancy varying per 1ms unit of time. For each occupancy we overlaid Glaze data using data rates ranging from 0.03-10kbps and attenuation of Δ=8 dB. Figure 4.11(b) shows a plot of the achievable throughput as a function of Wi-Fi traffic occupancy ranging from 65-100%. We can see that at 100% Wi-Fi traffic occupancy we can achieve the

Table 4.1: **Comparison to other Downlink Backscatter (BS) solutions.** A comparison to other backscatter solutions in enabling downlink communication.

Solution	Downlink	Signal Compatibility	Device Modification	Downlink Throughput	Range
Ambient BS [110]	×	×	×	×	×
FM BS [203]	×	×	×	×	×
Wi-Fi BS [93]	✓	Wi-Fi	✓	20, 5kbps	2.2, 3m
Glaze	✓	Wi-Fi,FM,TV	×	10kbps	3.5m

maximum data rate of 10kbps and as the occupancy decreases the achievable throughput decreases to the baseline data rate of 3bps. It is important to note that the upperbound on achievable throughput of the Glaze system is limited by the prototype implementation. Specifically, the MCU used at the Glaze receiver can only decode data at rates up to 10kbps. In future implementations, this can be improved by using higher frequency crystals at the MCU. Table 4.1 compares the performance of Glaze to other backscatter solutions in terms of downlink capability, signal compatibility, modification to the hardware/firmware of existing wireless transmitters, throughput, and range. Solutions such as Ambient and FM Backscatter do not have downlink capabilities. On the other hand, Wi-Fi backscatter enables downlink communication for IEEE 802.11- protocols, but also requires modification to Wi-Fi APs. Wi-Fi backscatter achieves data rates of up to 20kbps at distances $\leq$2.2m and 5kbps at 3m, whereas, Glaze reaches data rates of up to 10kbps at 3.3m. Given these results, the two systems are not directly comparable in terms of throughput performance. Moreover, referring back to Section 4.6, the achievable throughput of Glaze is limited by the 8MHz crystal in the prototype implementation and can be further improved with minor hardware modification to the Glaze receiver.

Lastly, this throughput analysis is utilized at the Glaze module, along with results from Section 4.7.1 and 4.7.2, to perform parameter optimization and adaptive rate embedding to ensure that we maximize performance while minimizing the interference towards the original wireless signals.

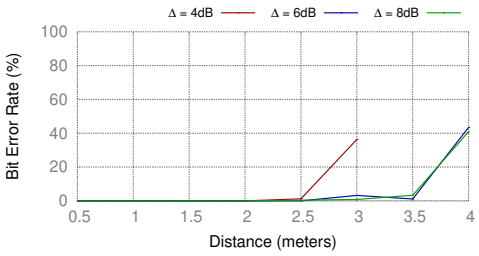

(a) **Wi-Fi Overlay Performance.** The BER as a function of distance for the Glaze system in a Wi-Fi scenario.

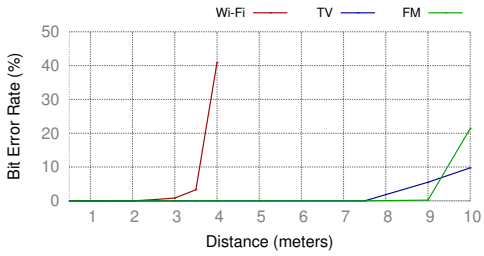

(b) **Performance Comparison.** The BER as a function of distance for the Glaze system in FM, TV and Wi-Fi scenarios.

Figure 4.12: **Performance of the Glaze System.** The achievable throughput and BER performance for the Glaze system when overlaying data on Wi-Fi, TV, and FM signals.

Bit Error Rate versus Distance

Next, we evaluate the bit error rate (BER) as a function of distance between the Glaze module and receiver. We perform our experiments on Wi-Fi packet transmissions in an indoor lab setting and FM and TV broadcast signals outdoors on a university campus. Each wireless transmitter is placed in a fixed position with the Glaze module attached to the antenna port. The Glaze receiver is placed at a varying distance from the transmitter. To evaluate the Glaze system for Wi-Fi signals, we use an 8dBi directional antenna at the transmitter and a 3dBi omnidirectional antenna at the Glaze receiver. The transmit power at the wireless transmitter was set to 20dBm. In the case of evaluating the system for FM and TV broadcast signals, we also used an 8dBi directional antenna at the transmitter and 3dBi omnidirectional antenna at the Glaze receiver. Since it is impractical to emulate FM and TV broadcast tower, we transmit TV and FM broadcast signals at 915MHz with a transmit power of 20dBm.

Figure 4.12(a) shows the BER as a function of distance for Glaze packets overlaid on Wi-Fi packet transmissions. We performed a controlled experiment with a bit rate of 10 kbps for Wi-Fi traffic occupancy approximately 100% to demonstrate the the maximum performance

of the Glaze system. The results show that with an attenuation of Δ=4dB the Glaze system can achieve data rates of up to 10 kbps at a distance of 2.5 meters and when the attenuation is further increased to 6 and 8dB the distance improves to 3.5 meters. Figure 4.12(b) compares the performance of the Glaze system for all three RF signals. Similarly to Wi-Fi, Glaze packets were overlaid on TV and FM broadcast signals using a bit rate of 10 kbps for varying amounts of attenuation. We can see that FM has the best performance with range up to 10 meters, followed by TV at 9 meters, and Wi-Fi at 3.5 meters.

Fundamentally, Wi-Fi has a shorter communication range in comparison to FM and TV just by considering the carrier frequency. In other words, as the carrier frequency increases the path loss increases as well, in turn decreasing the communication range of a wireless communication system. We see this in our results, where the communication range for the Glaze system is approximately half in comparison to the FM and TV broadcast signal results, which where transmitted at 915MHz. This makes sense because 915MHz is approximately half of the carrier frequency for Wi-Fi, 2.4GHz. While the evaluation for FM and TV was performed at 915MHz, we would expect in a real-life scenario the communication range for broadcast signals to be significantly larger. The reasoning behind this is twofold. First, digital TV operates in the UHF and VHF band in the United States, which is much lower than 915MHz. Similarly, FM broadcast signals in the United States are transmitted at carrier frequencies ranging from 87.8-107MHz. Second, the transmit power at TV and FM broadcast towers are drastically higher than 20 dBm. An estimate of communication range for real-life FM and TV broadcast scenarios is 125km and 75km, respectively.

4.7.3 Proof of Concept

For proof of concept, we integrated the Glaze system with a commodity FM transmitter. Figure 4.13 shows the proof of concept implementation. The Glaze module is connected to a commodity FM transmitter and a laptop PC is used to provide an audio input to the FM transmitter. The FM signal is broadcast at 108MHz and received by a commodity FM radio. Meanwhile, the Glaze receiver is able to decode the overlaid packets. To give

readers an intuition of FM audio quality with Glaze embedding and performance of the Glaze receiver we provide a video clip demonstrating the Glaze system in the following web link, https://youtu.be/FIJPrj2Pt04

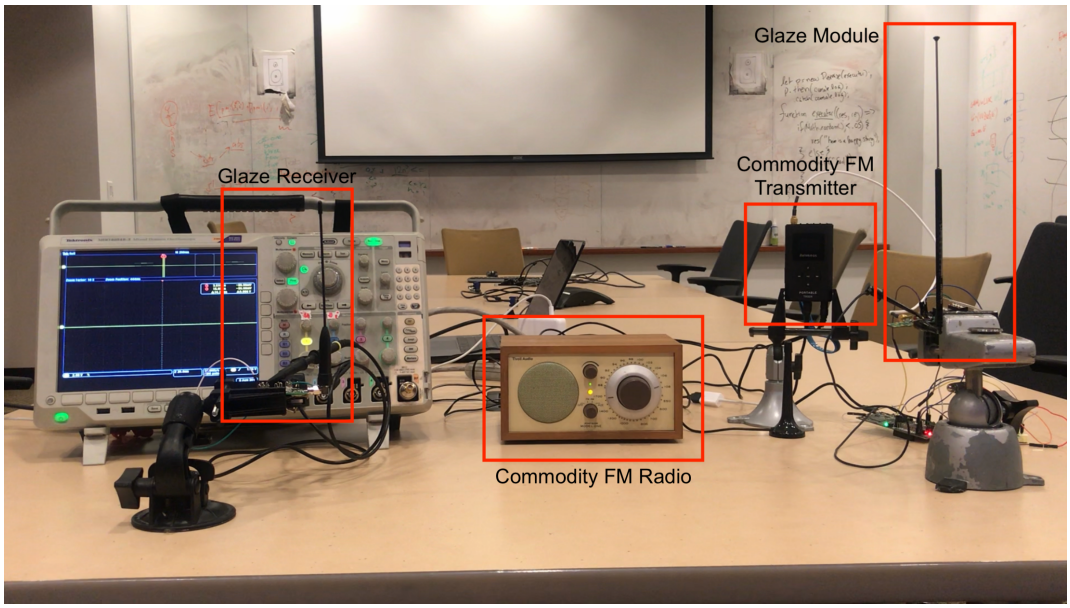

Figure 4.13: **Proof of Concept.** The Glaze system integrated with a commodity FM transmitter.

4.8 Conclusion

We introduced Glaze, a new communication system that overlays data on existing wireless signals to enable downlink communication for ambient backscatter systems while having minimal distortion impact towards preexisting communication channels. In this section, we highlight several future research directions for the Glaze system.

- **Improving Throughput and Sensitivity.** We have shown that Glaze can be integrated with a variety of wireless transmitters to overlay downlink data transmissions. However, the performance for certain scenarios can still be improved. For instance,

the range in Wi-Fi scenarios is currently limited to 3.5m. This can be improved by using better performing hardware, for instance replacing the MCU, but a possibly more power conscious approach would be exploring different coding techniques to increase throughput and sensitivity. In relation to throughput, performing multilevel embedding can be implemented. For instance, multiple states of attenuation can be used to embed multiple bits per symbol. Developing protocols to perform adaptive multi-level embedding based on channel conditions would also help boost performance. This would require the Glaze receiver to have knowledge of the data rate and number of attenuation states used for encoding data, requiring a more complex physical layer implementation.

- **Forecasting Wi-Fi Traffic.** A two-state Markov model was implemented to aid in predicting the short-term downlink Wi-Fi traffic occupancy and in turn optimizing Glaze embedding. While this technique has proven to perform well, there are many other approaches that can be explored to further improve the the performance of Glaze embedding. For instance, recurrent neural networks can prove to provide in-depth insights about Wi-Fi traffic. One example is the long short-term memory model and can be used to forecast Wi-Fi traffic, predict interarrival times, or average packet size with presumably better accuracy.

- **Multi-band Glaze Receiver.** The current topology of the Glaze receiver requires modifying the impedance circuit depending on what wireless signal is being used to overlay data due to the varying carrier frequency. A potential research direction to solve this problem is to explore implementing a multi-band Glaze receiver that could operate at ISM, UHF, and VHF bands.

Chapter 5

COMMUNICATION BY MEANS OF MODULATED JOHNSON NOISE

5.1 Introduction

In passive wireless communication, an energy-constrained data transmitter sends information by modulating an RF signal generated by an RF source that is not power constrained. Because the data transmitter does not have to generate an RF signal, the power necessary to send data is orders of magnitude less than in conventional RF communication. In Modulated Backscatter Communication, a continuous wave RF carrier is generated by a dedicated device on the high-power side of the link, and the low power side encodes data by selectively reflecting this RF signal [184]. Ambient Backscatter is another form of passive communication that makes use of pre-existing, ambient RF signals such as those generated by broadcast TV or radio towers [110]. While the low-power of the data transmission side is attractive, both methods rely on a pre-existing RF signal.

This paper introduces a new form of passive wireless communication, Modulated Johnson Noise, in which the signal to be modulated is the Johnson noise in an unbiased (un-powered) resistor. This scheme retains the benefits of prior passive wireless communication schemes while eliminating the need for an external RF signal. This has the potential to reduce the overall energy consumption of the system, to allow more stealthy and low-interference operation, and to allow operation in areas where no ambient RF signals are available.

Consider the frequency spectrum measurements shown in Fig. 5.1A, which shows the measured signal of a 50Ω terminator and open circuit terminator connected to the input of a receiver. We can see that there is a clear difference between the two measurements, which can be exploited to enable wireless communication. By selectively connecting and discon-

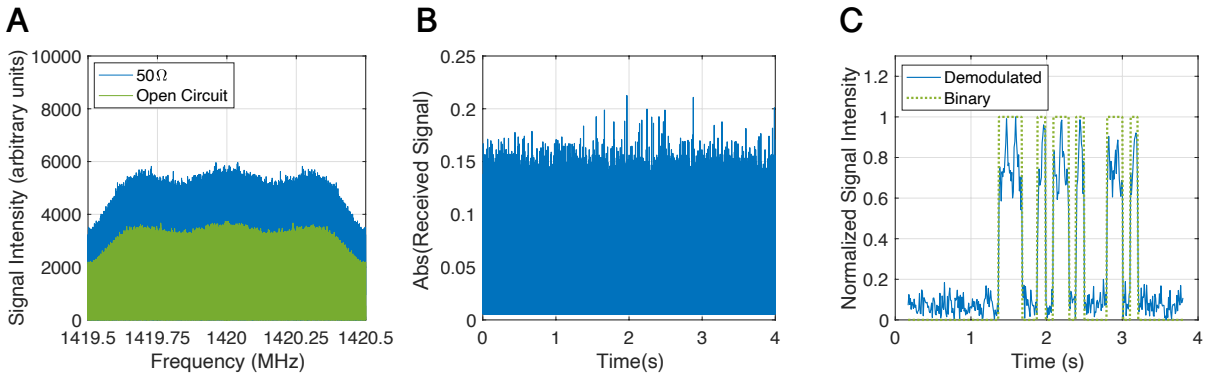

Figure 5.1: **Johnson Noise Communication.** Wireless communication can be enabled by modulated Johnson noise. **(A)** shows the measured frequency spectrum at 1.42GHz, which compares measurements when a 50Ω load is connected to a receiver and an open circuit load. **(B)** shows what a received signal looks like when wirelessly transmitting a data packet by modulated Johnson noise and **(C)** shows the received data packet after demodulation.

necting an impedance matched 50Ω resistor to an antenna, information bits can be wirelessly transmitted. For example, in Fig. 5.1B we show the received signal of a data packet that was wirelessly transmitted by modulating Johnson noise. While this looks like a noisy signal, after performing demodulation, the data packet can be extracted as shown in Fig. 5.1C. While thermal noise communication has been proposed from a theoretical perspective [99], we present the first system that enables wireless communication by modulating Johnson noise. In this paper, we discuss the design and experimental implementation of the system and evaluate the overall performance of the wireless communication scheme that relies on modulated Johnson noise. The contributions of the paper are summarized below:

- We introduce wireless communication by means of modulated Johnson noise, the first wireless system that enables devices to communicate without reliance on generated or ambient RF signals.

- We present the designs and prototype of hardware that enables Johnson noise commu-

nication. Moreover, we demonstrate the transmitter can be designed to be battery-free.

- We evaluate the performance of the wireless system in terms of achievable throughput and communication range. The performance evaluation shows that data rates of up to 26bps can be achieved at distances of up to 7.3 meters.

Overview of Johnson Noise

Johnson noise is caused by the thermal vibrations of charge carriers inside of an electrical conductor (e.g., resistor) and is characterized by its mean-squared voltage,

$$V_n^2 = 4kTBRe(Z) \tag{5.1}$$

where k is Boltzmann's constant, T is temperature, B is bandwidth, and Re(Z) is the the real part of the electrical conductor's impedance [**?**]. A resistor with Johnson noise can be modeled as a Thevenin equivalent circuit that includes a noiseless resistor and a noise voltage generator with voltage given by Eq. 5.1. With a matched load resistor connected, the maximum noise power provided by the noisy resistor is

$$P_n = \frac{V_n^2}{4R} = kTB \tag{5.2}$$

which is independent of resistance [143]. The thermal noise power is a function of temperature and bandwidth. As an example, a resistor at room temperature (296K) with a system bandwidth of 500MHz would result in P_n = -86.9 dBm. We also note that Johnson noise is a white noise noise source and therefore independent of frequency (within the system's bandwidth). Since it is white, Johnson noise has a Gaussian distribution.

5.2 Design and Implementation

We design and implement a transmitter (TX) and receiver (RX) to enable wireless communication by modulating Johnson noise. First, the transmitter requires an RF switch and a

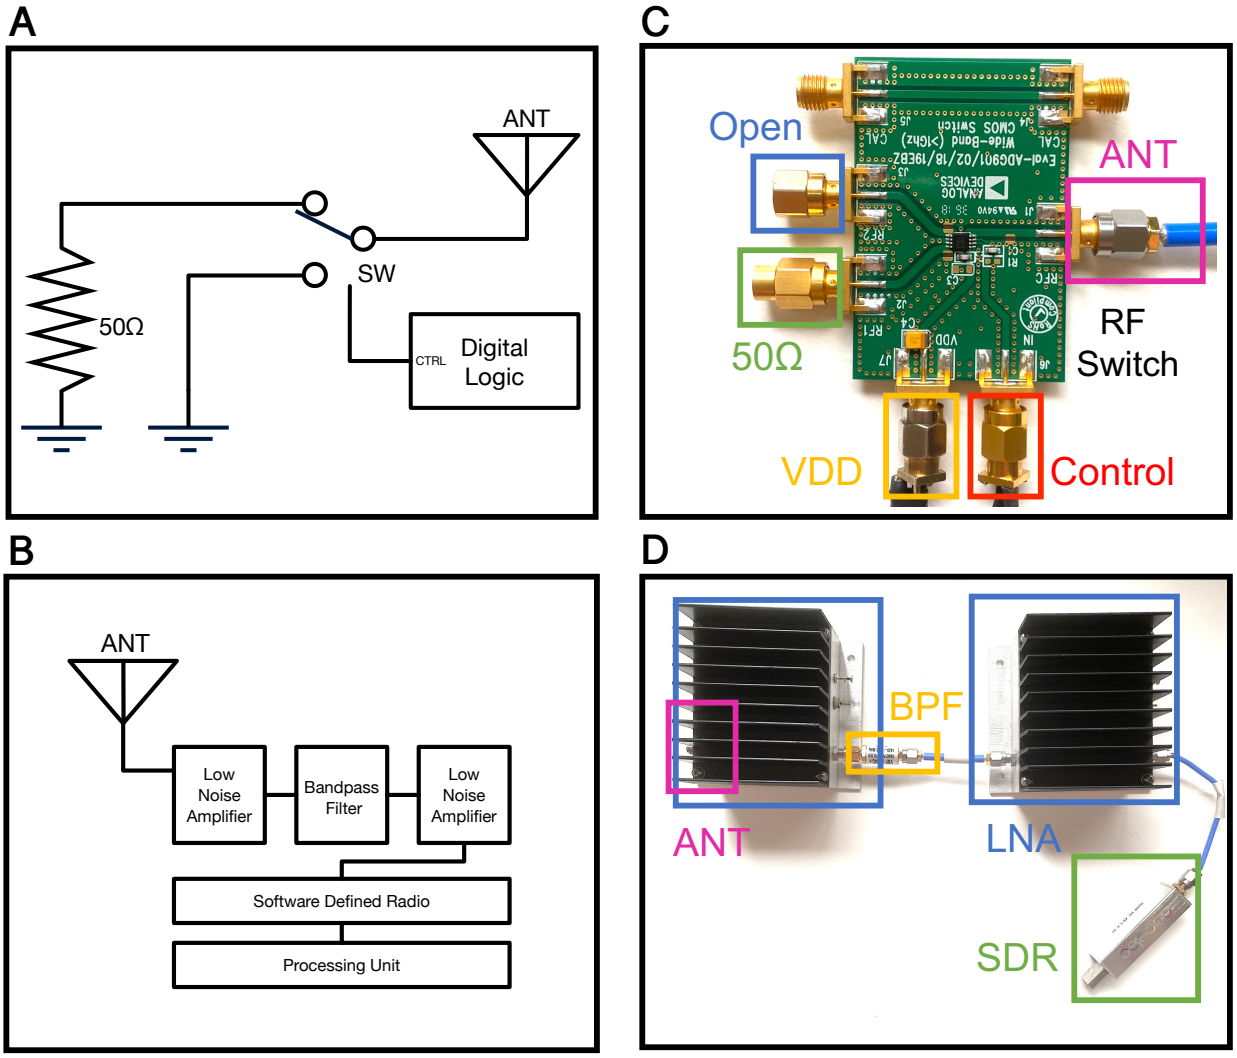

Figure 5.2: **Design and Implementation.** **(A)** shows the transmitter design, which switches between a 50Ω load and short circuit (or open circuit) to modulated information bits. **(B)** shows the receiver, which is composed of two low noise amplifiers (LNA) who's output is fed into a software defined radio. The bandpass filter between the two LNAs is added to prevent feedback oscillation and the data received by the SDR can be processed by laptop PC or small-form factor computer (e.g. Raspberry Pi). **(C)** and **(D)** show the prototype implementation of the transmitter and receiver, respectively.

processing unit to control switching between an open or short circuit, and resistor load connection shown in Fig. 5.2A. On the receive side, shown in Fig.5.2B, two high gain, low noise amplifiers (LNA) amplify the very low-power signals being transmitted while maintaining a good signal-to-noise ratio. Between the two LNAs is a bandpass filter to prevent feedback oscillation. The entire system is designed to operate at 1.42GHz with 50Ω impedance. In Fig. 5.2C and 5.2D we show the prototype implementation of the transmitter and receiver. The transmitter uses an RF switch and switches between a 50Ω RF terminator (with 50Ω impedance) and an open circuit terminator, both of which are very well shielded [43]. On the receive side there are two LNAs with combined gain of approximately 84dB, a 50MHz bandpass filter, and a software defined radio (SDR) [126, 133, 125]. We constructed pyramidal horn antennas for both the TX and RX side. The antennas were characterized to have approximately 13.6dBi of gain.

Data Encoding and Modulation

Since the transmitter is switching between two states, data can be modulated by performing ON-OFF keying. A 0-bit is transmitted by simply staying in the OFF state (open circuit) and a 1-bit is transmitted by switching between both ON and OFF states (open circuit and 50Ω) using a square-wave subcarrier frequency. Mathematically, the transmitted signal is written as,

$$s_{rx} = m \cdot sgn\big(sin(2\pi f_{sc} t)\big) \tag{5.3}$$

where m represents the bit to be encoded and takes on a 0 or 1 value. The subcarrier frequency is defined as f_{sc}.

Packet Detection and Demodulation

Our system assumes that the f_{sc} is known on the RX side, which allows us to perform heterodyne detection. Our demodulation is inspired by techniques used in radio astronomy,

in particular, the Dicke radiometer [**?**]. We can view our entire system as a distributed Dicke radiometer, where the switching between two states occurs on the TX side, while the amplification and integration occurs on the RX side. The receiver performs heterodyne detection which allows the system to reject noise by filtering out any received signal power that is outside the narrow bandwidth of the subcarrier. In other words, the receiver generates the same square-wave subcarrier signal with amplitude values of +1 and -1, multiplies the received signal with this receive-side subcarrier, and accumulates (integrates) the product values for a duration that is less than or equal to that of the duration of one bit. This process results in a demodulated signal intensity and is given by,

$$\langle s_{rx}, s_{sc} \rangle = \int_0^T s_{rx}(t) s_{sc}(t) dt \tag{5.4}$$

where T is the total integration time. The demodulated signal intensity is compared to a threshold value to determine whether a 1 or 0-bit was transmitted. In Fig. 5.1B we show an example of a received data packet before demodulation that was wirelessly transmitted, and in Fig. 5.1C we show the data packet after demodulation and its binary format after performing thresholding.

5.3 System Validation

With this method of wireless communication comes two key questions that must first be answered: (1) Are we indeed modulating Johnson noise and not another source (e.g., ambient RF or control signal feedthrough)? (2) How does modulated Johnson noise work if the entire system is at thermal equilibrium?

5.3.1 Evaluating Feedthrough

One concern that comes up with this system is whether there is feedthrough from the control signal used for the RF switch. We have performed control experiments to evaluate signal isolation. If the system is truly using thermal noise, then if it is configured to switch between

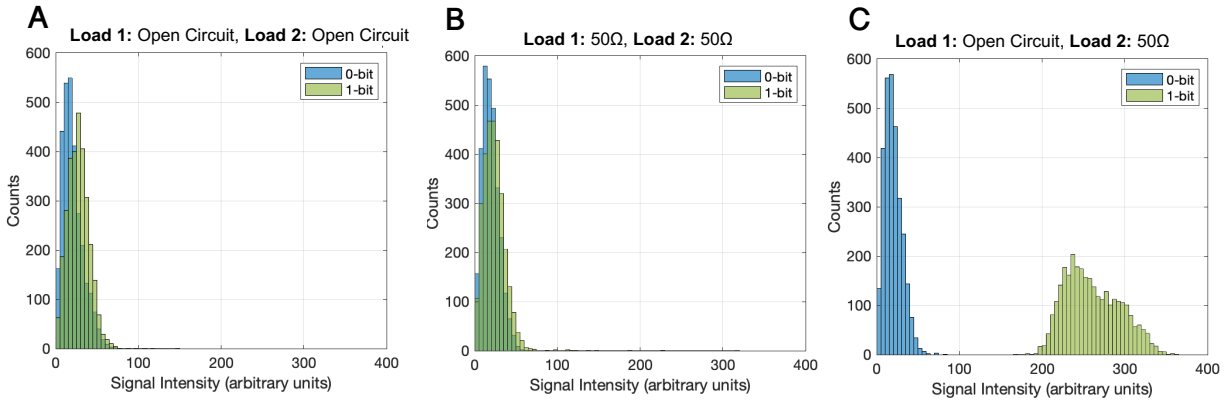

Figure 5.3: **Feedthrough Evaluation.** To ensure the system has proper signal isolation, we evaluate the possibility of feedthrough from the control signal used for the RF switch in the transmitter. In **(A)** the histogram distribution of 0 and 1-bit transmissions after demodulation when switching between two identical open circuit loads is shown. Similarly, **(B)** shows the results when switching between to identical 50Ω loads at room temperature. In **(A)** and **(B)**, the 0 and 1-bits are clearly not distinguishable. However, when switching between an open circuit and 50Ω load, shown in **(C)**, there is a clear difference between 0 and 1-bit transmissions.

two open circuit loads or two 50Ω loads at room temperature, it should not be possible to decode data transmissions. We evaluate these two scenarios and compare the results to that of switching between 50Ω and an open circuit.

An additional potential concern is that we might not truly be using Johnson noise. Perhaps some pre-existing RF signal, such as 60Hz, broadcast TV, or even galactic emissions from neutral hydrogen, are being picked up in our apparatus, and then being modulated by our two different load resistances, resulting in different outputs from our TX antenna. To control for this, we tested our data transmission method by switching between two 50Ω loads at drastically different temperatures inside an anechoic chamber. If we truly are using Johnson noise, and not some other mechanism related to the different impedances, then two

identical impedances at different temperatures should also work. We demonstrate that data can be modulated by switching between a 50Ω load at room temperature (296K) and a 50Ω load submerged in liquid nitrogen (77K).

(1) Control Experiment: Feedthrough

A cabled benchtop experiment is set up to evaluate the three different scenarios. Data bits (1 and 0) are transmitted and then demodulated by the receiver. The demodulated data is presented in the form of histogram distributions which show the signal intensity of demodulated 1 and 0-bits. The open circuit and 50Ω loads used were RF shielded terminators. In Fig. 5.3A we show the system evaluated when switching between two open circuit loads. In Fig. 5.3B we show the system evaluated when switching between two 50Ω loads. We compare this to the results shown in Fig. 5.3C, where the system is switching between open circuit and 50Ω. From these results, we can see that when switching between two open circuit states or two 50Ω states, the distribution of 1 and 0-bits look nearly identical. On the other hand, when switching between open circuit and 50Ω this is a clear difference between 1 and 0-bit transmissions.

(2) Control Experiment: Temperature Modulation

Using the same hardware setup as for the feedthrough experiments, we modulate information bits by switching between two 50Ω loads that are at two different temperatures. A 50Ω load is submerged in liquid nitrogen (77K), while the other 50Ω load is at room temperature (296K). With this setup, we transmit data packets. In Fig. 5.4A we show the demodulated data and in Fig. 5.4B we show the distribution of 1 and 0-bits after demodulation. Here, data is transmitted using a data rate of 5bps and a subcarrier frequency of 100Hz.

5.3.2 Evaluating Noise Temperature

One may worry that if the entire experiment is done at a single temperature (e.g. room temperature) then it should not be possible to observe modulated thermal signals due to

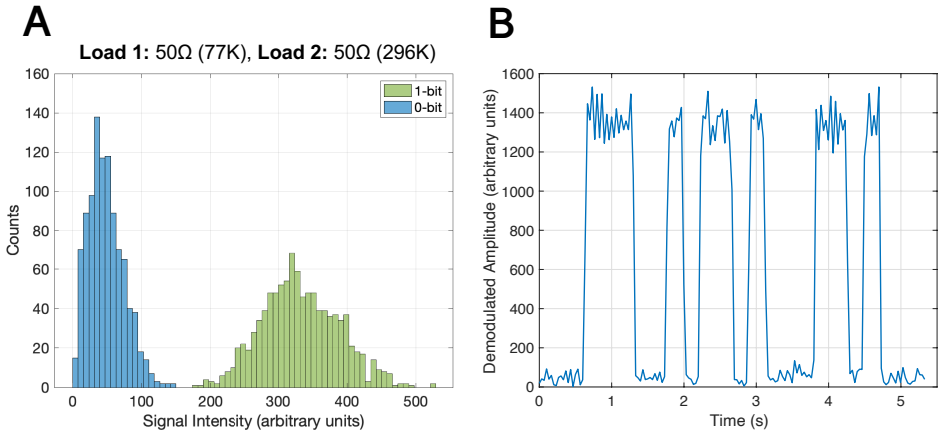

Figure 5.4: **Temperature Modulation.** Since Johnson noise power is a function of temperature, it should be possible to transmit information by switching between two 50Ω loads that are at drastically different temperatures. We modulate information bits by switching between two 50Ω loads, where one load has a temperature of 296K and the other has a temperature of 77K. **(A)** shows the histogram distribution of 1 and 0-bits after demodulation and **(B)** shows a demodulated data packet.

a lack of contrast between the resistor and its surroundings. The question can be refined further and there are two separate potential issues. First, the transmit side and the receive side are at the same temperature and second, the two loads on the transmit side (50Ω and open circuit) are at the same temperature as one another. In the following sections we address these concerns by first discussing the temperature contrast between TX and RX, as well as 50Ω and open circuit loads. We then present experimental measurements that estimate the noise temperature of each load and compare the measured noise of each load to a theoretical analysis.

(1) TX-RX Contrast

If the receive side and the transmit side of the system are at the same temperature, does this mean that our method should not work? If this was the case, then by the same argument, it would not be possible to observe Johnson noise with a room temperature oscilloscope,

which in fact is easily observed. Furthermore, it is straightforward to perform a benchtop experiment that is analogous to our communication scheme: if one connects or disconnects a room temperature resistor (e.g 1MΩ) to an oscilloscope with a 1MΩ input impedance, one sees a much larger signal in the connected case than in the disconnected case. This is analogous to our communication set up, but in the communication scheme the scope lead has been replaced by a pair of antennas. In the disconnected case, the oscilloscope's LNA is poorly matched, and one is left seeing just the amplifier's internal noise. In the connected case, the noise from the resistor is much larger than the amplifier's internal noise.

Active amplifiers, such as those found in an oscilloscope or the low noise amplifiers (LNA) in our receiver, are not in thermal equilibrium and are characterized by a noise temperature (T_N) that is "cold": lower than the physical temperature (T_P) of the apparatus itself. This means that the radiation emitted by the receive antenna coupled to the LNA is equivalent to that of a colder 50Ω resistor.

(2) 50Ω-Open Contrast

The two loads we switch between for signaling are at the same temperature. Furthermore, the Johnson noise *power* is independent of the value of the resistance. Thus one might worry that no useful contrast can be generated. The Johnson noise *mean squared voltage* on the other hand is proportional to the value of the resistance. We analyze this quantitatively below, but in summary this is because the 50Ω load is well matched to the LNA input, and the open and short are poorly matched.

To evaluate this explanation quantitatively, we first measure the impedance of the open, short, and 50Ω loads at our operating frequency of 1.42GHz. Fig. 5.5 shows the impedance of the open, short, and 50Ω loads. As expected the open circuit has a high resistance, approximately 17kΩ, and the short circuit has a low resistance, approximately 0.25Ω. To predict the mean squared Johnson noise voltage across the load resistor, we plug these measured impedance values into the expression for the mean squared Johnson noise voltage defined by Eq. 5.1. Now, we need to model the effect of the LNA, in particular the effect

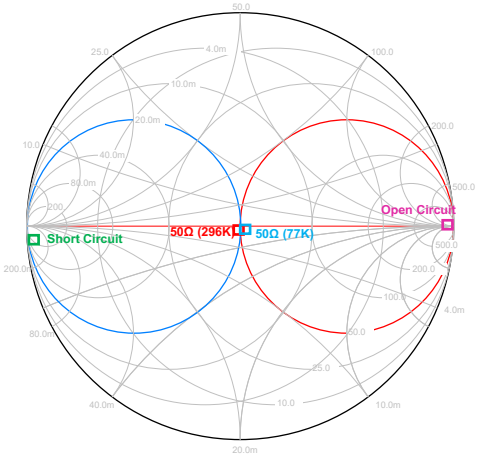

Figure 5.5: **Load Impedance.** A Smith chart showing the impedance of each evaluated load. The open and short circuit are both mismatched and have a resistance of approximately 17kΩ and 0.25Ω, respectively. The 50Ω load is very well matched and has a 50Ω impedance at both 296K and 77K.

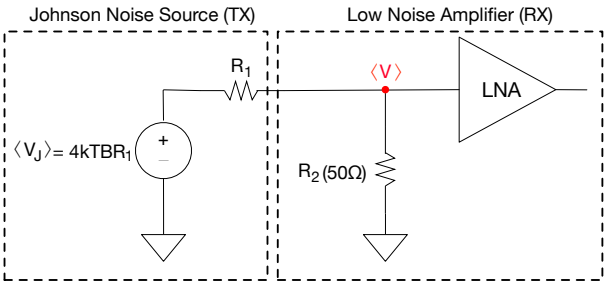

Figure 5.6: **Equivalent Circuit.** The equivalent circuit model for a cabled scenario showing the voltage divider formed by a noisy resistor and LNA input impedance.

of its input impedance on the measured signal. The input impedance of the LNA is 50Ω, meaning that it can be modeled as an equivalent circuit with a 50Ω resistor to ground and an infinite input impedance amplifier observing the voltage at the junction between the LNA's input and this 50Ω shunt resistor. Fig. 5.6 shows the voltage divider formed by the noise resistor and LNA input. The mean squared voltage observed by the amplifier, $\langle V^2 \rangle$, is given by

$$\langle V^2 \rangle = \langle V_J^2 \rangle \cdot g^2 = 4kTBR_1 \cdot g^2 \tag{5.5}$$

where

$$g = \frac{R_2}{R_1 + R_2} \tag{5.6}$$

is the gain of the voltage divider. As shown in Fig. 5.5, R_1 represents the resistance of

Load	T_P **(K)**	T_N **(K)**	G_{RX}	**B (Hz)**	**Offset**	$\widehat{V}^2$ **Predicted from Measured** T_P	$\widehat{V}^2$ **Calculated from Extracted** T_N	$\widehat{V}^2$ **Measured**
50Ω (Room Temperature)	296	296	2.3e11	1e6	0.0212	0.0676	-	0.0676
50Ω (Liquid Nitrogen)	77	77				0.0332	-	0.0333
Open Circuit	296	46				0.0217	-	0.0274
Short Circuit	296	40				0.0221	-	0.0283
LNA Input	-	39.5				-	0.0274	0.0273

Table 5.1: **Computing Theoretical Gaussian Distribution.** This table shows the parameters used to compute the theoretical $\widehat{V}^2$ for each load. The theoretical $\widehat{V}^2$ values can be compared with the measured $\widehat{V}^2$ in the last column of the table. In the theoretical Gaussian distributions later in the paper, the variable $\widehat{V}^2$ from this table is renamed σ^2. The T_P values above are measured physical temperatures. The T_N values for the 50Ω loads are assumed, to produce the linear calibration curve shown in Fig.5.7. The other T_N values in the column are extracted by applying this linear calibration curve to the measured $\widehat{V}^2$ values.

the Johnson noise source and $R_2 = 50\Omega$ is the LNA's input impedance. The infinite input impedance amplifier block in Fig. 5.6 measures the voltage V at the node where R_1 and R_2 connect. The mean squared input voltage to the amplifier is given by Eq. 5.5. Expressed as multiples of kTB, the mean squared noise voltages $\langle V^2 \rangle$ for our open, short, and 50Ω loads are predicted to be

$$\langle V^2_{open} \rangle = 0.58 \cdot kTB$$

$$\langle V^2_{short} \rangle = 1.0 \cdot kTB$$

$$\langle V^2_{50} \rangle = 50 \cdot kTB$$

The varying constants in front of kTB demonstrate that the impedance matching differences lead to small signals from the open and short, and larger signals from the 50Ω load. Later (after addressing calibration) we provide a detailed comparison of predicted mean

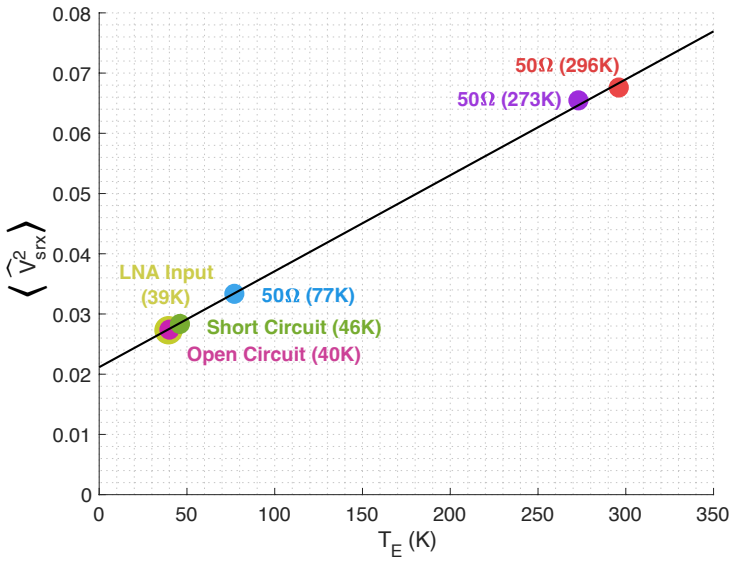

Figure 5.7: **Noise Temperature.** The measured mean square noise voltage and noise temperature for each load. The three 50Ω loads at known temperatures were used to calibrate the noise temperatures for the other loads. Here the linear fit equation is, $\langle \hat{V}^2_{s_{rx}} \rangle = 0.000159 \cdot T + 0.0212$

squared noise voltages to experimentally measured values for each of the loads of interest, shown in Table 5.1. Note that the experimentally observed contrast between the loads is less than the theoretical differences here, because of degradation in SNR caused by noise added by the receive signal chain.

(3) Experimental Measurements and Calibration Overview

We measure the noise produced by each load of interest in a cabled benchtop experiment, and use this to compute the corresponding noise temperature along with calibration parameters (i.e., receiver noise offset and gain). In addition to measuring the noise of the 50Ω load at room temperature, we also evaluate the 50Ω load when submerged in ice water (273K) and liquid nitrogen (77K). The three different temperatures for the 50Ω load are considered

to be *known* temperatures. One special load of interest we evaluate in addition to the communication loads is an LNA input, in order to characterize its effective temperature.

We collected 60 million measurements of the noise generated by each of the loads and found for each the mean squared voltage $\langle \widehat{V}_{s_{rx}}^2 \rangle$ using

$$\langle \widehat{V}_{s_{rx}}^2 \rangle = \frac{1}{N} \sum {s_{rx}}^2 \tag{5.7}$$

where s_{rx} is the received signal.[1] The hat (such as in $\widehat{V}_{s_{rx}}$) indicates that the voltages are in arbitrary SDR units. Voltages without a hat are in physical units (Volts). Using measurements of the 50Ω load at the three known temperatures, we performed a linear fit of $\langle \widehat{V}_{s_{rx}}^2 \rangle$ versus temperature. Using this linear fit, we extracted the noise temperature of the remaining loads (LNA Input, open, and short) from the observations of $\widehat{V}_{s_{rx}}$. Fig. 5.7 shows the results, which indicate that the LNA's noise temperature is cold (39K), much colder than a 50Ω load at room temperature (296K). The open and short circuit loads also appear colder than room temperature. The short circuit terminator has a physical temperature of 296K, but its noise temperature is 46K. The open circuit has a noise temperature of 40K.

(4) Comparing Theoretical and Observed Noise Distributions

To compare our experimental measurements with expected theoretical values, we also need to perform some calibration. In particular, we need to find the baseline or offset noise produced by the receive chain, as well as the gain value that converts between the SDR's arbitrary units and volts. This calibration of arbitrary SDR units to volts will allow us to compare the theory and experiment for the less obvious loads, namely the open and the short. Using our results from Fig. 5.7, we know that the receive chain introduces noise equivalent to an offset of 0.0212 in arbitrary SDR units. In other words, the offset is the $\langle \widehat{V}_{s_{rx}}^2 \rangle$ value at the y-intercept of the linear fit. The gain, which is constant across all measurements, can be computed by determining the ratio of the measured data for any particular load to the

[1] In particular, it is the real part of the complex values provided by the SDR; the imaginary part has the same statistics.

expected theoretical value, and then using this same value for all other datasets. The gain is given by,

$$G_{RX} = \frac{\langle \widehat{V}_{s_{rx}}^2 \rangle - offset}{\langle V^2 \rangle} \tag{5.8}$$

Now we can compare our theoretical model (based on impedance mismatch of Johnson noise) to measured values. For each load, we compare the histogram of our observed noise measurements with the theoretical distribution.

To find the distributions, recall that Eq. 5.5 predicts $\langle V^2 \rangle$ as a function of Temperature, Bandwidth, Johnson resistor value, and LNA input impedance. Since thermal noise is white, the distribution of the real or imaginary voltages is Gaussian with $\mu = 0$ and $\sigma = \sqrt{\langle V^2 \rangle} = V_{RMS}$. We know from the results in Fig. 5.5 that $R_1(short) = 0.254\Omega$ and $R_1(open) = 17k\Omega$ at our frequency of interest. We also must account for the gain and offset of the receive chain. As previously mentioned, the receive chain introduces noise equivalent to an offset of 0.0212 in arbitrary SDR units. We can also compute the receiver gain using Eq. 5.8. Now, the variance σ^2 (which we also refer to as $\langle \widehat{V}_{s_{rx}}^2 \rangle$) can be related to the theoretical value of $\langle V^2 \rangle$ using

$$\sigma^2 = \langle V^2 \rangle \cdot G_{RX} + offset \tag{5.9}$$

The Gaussian distribution is then computed for each case:

$$f(x) = \frac{1}{\sqrt{2\pi}\sqrt{\sigma^2}} \cdot e^{-\frac{1}{2}\left(\frac{x-\mu}{\sqrt{\sigma^2}}\right)^2} \tag{5.10}$$

Fig. 5.8 shows the histogram distribution of the measured data for each load and the the corresponding theoretical distribution. The bandwidth is determined by the bandwidth of the SDR, and in particular its anti-aliasing filters.

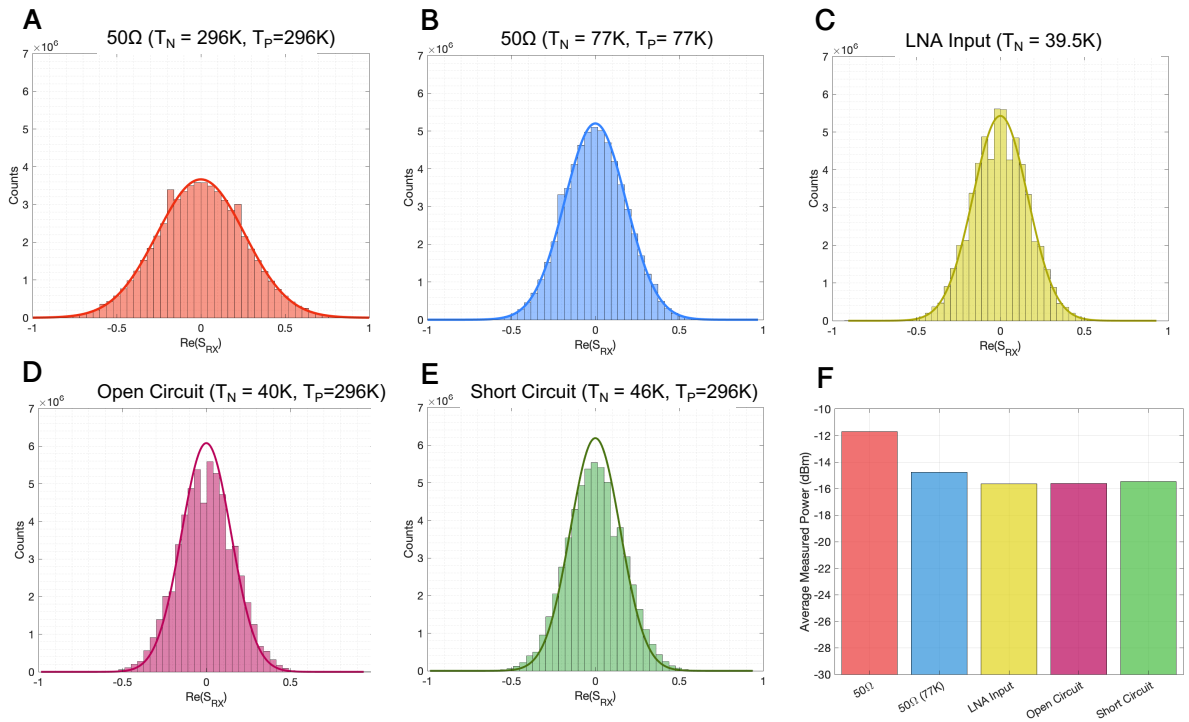

Figure 5.8: **Theoretical noise distributions compared to data.** For each load, the Noise Temperature T_N and Physical Temperature T_P is indicated at the top of the plot. **(A):** histogram distribution of the real component of the measured data for the 50Ω terminator at room temperature (296K). The theoretical Gaussian distribution is overlaid on the data. **(B):** 50Ω terminator in liquid nitrogen (77K) **(C):** the LNA input has a noise temperature T_N=39K. Its low temperature is important to the operation of the system. **(D):** the short circuit terminator has a low effective temperature T_N=46K because it is poorly matched to the LNA **(E):** the open circuit terminator has a low temperature of T_N=40K, also because of mismatch. **(F):** a comparison of the average measured power of each load in dBm after amplification.

5.4 Performance

The performance of the wireless communication system is evaluated in terms of achievable throughput and communication range by conducting wireless experiments using the previ-

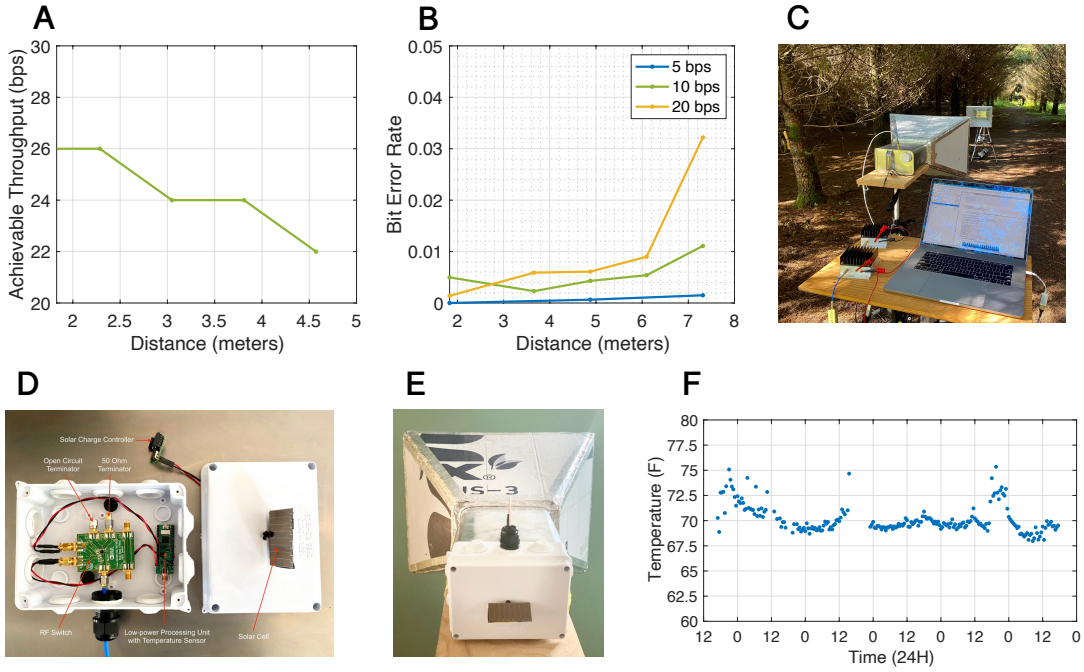

Figure 5.9: **Performance Evaluation.** **(A)** shows the achievable throughput with respect to distance, which was evaluated wirelessly inside of an anechoic chamber. **(B)** shows the bit-error-rate with respect to distance for three different data rates: 5bps, 10bps, and 20bps. Here, the communication range was evaluated wirelessly in an outdoor setting as shown in **(C)**. **(D)** and **(E)** show a battery-free transmitter prototype that collects ambient temperature data and transmits the information by means of modulated Johnson noise. The temperature plot in **(F)** shows the data that was collected wirelessly across 5 consecutive days in a residential building.

ously described hardware. Additionally, we demonstrate the potential for battery-free sensing by implementing a battery-free transmitter that transmits data by modulating Johnson noise.

5.4.1 Achievable Throughput

To evaluate achievable throughput, the measurements were performed inside of an anechoic chamber. This allows us to test the system in an interference free environment and to

achieve best possible performance given our hardware implementation and the modulation and demodulation scheme. Measurements were taken every 1.5 meters and up to 4.5 meters inside the anechoic chamber, and at each distance data packets were transmitted wirelessly. The data packet structure includes a 7-bit Barker code as the preamble, followed by a 13-bit data payload. At each distance, the maximum achievable throughput was determined by measuring which data rate would allow us to achieve a bit-error-rate (BER) less than or equal to 1%. Fig. 5.9A shows that the maximum achievable throughput is 26bps at 1.5 meters and goes down to 22bps at 4.5 meters. Moreover, the number of samples per bit was fixed to 5 while the subcarrier frequency was varied, which in turn varies the data rate. For example, at 2.3 meters the achievable throughput was 26bps and data was modulated by using a 130Hz subcarrier frequency. At 4.6 meters the achievable throughput was 22bps and data was modulated by using a 110Hz subcarrier frequency.

5.4.2 Communication Range

The wireless communication range is evaluated in an outdoor environment, as shown in Fig. 5.9C. Here, measurements were taken starting from 1.8 meters and up to 7.3 meters for three different data rates. We evaluate the communication range using data rates of 5, 10, and 20bps. Fig. 5.9B shows the range performance results. As expected, a higher data rate results in a short communication range. For example, with a data rate of 20bps the maximum communication range is approximately 6.1 meters. Whereas, using a data rate of 5bps the communication range increases to approximately 7.3 meters while having a BER = 0.15%.

5.4.3 Battery-free Transmitter

We demonstrate the potential for battery-free sensing using modulated Johnson noise, by developing a battery-free transmitter, as shown in Fig. 5.9D and E. The transmitter includes an ADG901 RF switch and a custom PCB that integrates a low-power MCU, temperature sensor, and solar energy harvester [91]. The custom PCB is used to collect temperature data,

packetize the data, and control the RF switch to modulate data. The transmitter is powered using a small 2.54cm. x 6.35cm. solar cell that charges a supercapacitor. We deploy the battery-free transmitter inside a residential building to monitor ambient temperature data and transmit the information by means of modulated Johnson noise. The data was received using the receiver hardware previously described. In Fig. 5.9F, we show the temperature collected across five days. Note that no error correcting codes were implemented which resulted in some anomalies in the ambient temperature data collection.

5.5 Conclusion

This paper opens up a new direction for low-power wireless communication by demonstrating for the first time that information bits can be wirelessly transmitted by modulating Johnson noise. By selectively connecting and disconnecting an impedance matched resistor to an antenna, data can be wirelessly transmitted with no active oscillator on either the transmit or the receive side, and no pre-existing RF carrier, as in backscatter and ambient backscatter communication. We have demonstrated that our system can achieve data rates up to 26bps and communication range as far as 7.3 meters. This method has many attractive features compared to previous passive communication schemes such as ambient backscatter or RFID, because it is not reliant on ambient or generated RF sources. One of the challenges faced by backscatter communication is the increased system complexity: conventional half-duplex radio transmission requires just two entities, a transmitter and a receiver. Backscatter requires three entities because it needs a carrier generator in addition to the data transmitter and receiver. (In a monostatic backscatter reader, the carrier generator is packaged with the receiver, but the system architecture is still more complex.) In addition to the deployment benefits of the simplified system architecture, the absence of a carrier means that the reader device does not need to contend with self-jamming caused by the carrier, which is a key driver of RFID reader complexity and cost. While the performance of the prototype system is modest, we hope that the techniques presented will lead to a new avenue of research and help realize the vision of ubiquitous computing.

Chapter 6

CONCLUSION AND FUTURE WORK

In conclusion, this dissertation presents several technological innovations that take important strides towards solving major challenges in order to enable ubiquitous computing. We focus on two key aspects: (1) designing reliable and affordable end-to-end wireless sensing systems and (2) pushing the boundaries of low-power wireless communication that can lead to new applications of intermittent and battery-free sensing. We first presented FarmBeats, an IoT platform for data-driven agriculture. Data driven agriculture techniques can help alleviate the world's food problem by reducing waste in resources, increasing yield, and ensuring sustainable farming practices. In particular, studies have shown that precision irrigation techniques can increase yield by 45% all the while reducing water intake by 37% [6]. Such results extend to other precision agriculture techniques as well. While the efficacy of data driven agriculture has been demonstrated, these techniques are sparsely adopted in today's farming practices. This is primarily due to the expensive cost of data collection and the challenging environment of typical farming locations. FarmBeats is an end-to-end AI and IoT system for agriculture that helps address these challenges and enable precision agriculture techniques. In particular, FarmBeats gathers data from sensors, cameras, and drones, to produce actionable insights for farmers. Moreover, it can extend Internet coverage to the farm and is resilient towards weather related outages or power variability. Thus far the FarmBeats system has been deployed world wide, including locations in the United States, India, Africa, and China to name a few. FarmBeats is being used for a variety of agriculture applications such as storage monitoring, generating precision sensor maps of crop fields, animal monitoring, and much more.

Next we focused on challenges related to connectivity, and discussed new methods for

low-power wireless communication. We present new methods to enable downlink communication for ambient backscatter systems. In Wireless QIM, we leverage data hiding techniques typically used for watermarking and apply them to wireless communication to create a new 'hidden' channel for downlink communication. We further extend this work by developing Glaze, an end-to-end system to enable downlink communication by introducing small amount of attenuation to an RF signal in order overlay additional data. Glaze can be integrated with existing RF systems without hardware modification and enabling a downlink communication for ambient backscatter systems, without interfering with existing transmissions and maintaining low power consumption. Finally, we looked beyond conventional methods for wireless communication and demonstrated that noise sources can be leveraged to modulated information bits. Conventional communication techniques view noise sources as something that harms wireless communication capacity and something that needs to be eliminated. On the contrary, it is possible to use noise to enable low-power communication on resource-constrained devices. We presented the design and implementation of the first system that uses modulated Johnson (thermal) noise to enable wireless communication. We wireless transmit information bits by selectively connecting and disconnecting an impedance matched resistor to an antenna. Compared to other passive wireless communication methods such as backscatter, this is system is not dependent on ambient RF sources.

Going forward there are several avenues for future research that build upon the work presented in this thesis. Some potential avenues are presented below.

- **New Methods for Wireless Communication:** Modulated Johnson Noise has opened new avenues of research for low-power wireless communication. Aside from resistors, other electronic components or networks of components can be utilized to enable zero-power or low-power communication. Beyond electronic components, there are many other sources of noise that we often seek to minimize or eliminate and instead could be leveraged to enable new forms of communication. These new methods of communication can lead to promising solutions in many different fields, from biomedi-

cal and stealth communication to space or satellite communication. Imagine enabling passive wireless communication in space. Satellites could potentially communicate amongst each other using existing noise sources while consuming minimal amounts of power.

- **Smart Agriculture and Food Systems:** With climate change underway, enabling widespread adoption of data-driven agriculture techniques has become increasingly relevant. FarmBeats has shown that affordable and seamless data collection systems for agriculture can be enabled. By utilizing many technological innovations from TV White Spaces to new machine learning and vision algorithms we can pave the way to more efficient and sustainable food production. Existing systems can be further improved by utilizing low-power communication and sensing techniques that can withstand long-term deployments. Moreover, using these low-power techniques in the context of robotics for agriculture is another promising area to explore. For example, robots are used to pick apples from trees or enable precision spraying of herbicides. Many of these applications require the use of sensors and with power consumption being a key constraint for such systems, enabling low-power sensing is relevant. Similarly, robots can be equipped with low-power wireless communication systems, such as passive thermal noise communication, to collect data from nearby battery-less sensors. For instance, daily drone flights are often necessary to monitor crops. These drones could be equipped with hardware to collect sensor data across the entire farm field as well.

BIBLIOGRAPHY

[1] IEEE 802.11af: https://standards.ieee.org/findstds/ standard/802.11af-2013.html.

[2] Khaled Q Abdelfadeel, Victor Cionca, and Dirk Pesch. Fair adaptive data rate allocation and power control in lorawan. In *2018 IEEE 19th International Symposium on" A World of Wireless, Mobile and Multimedia Networks"(WoWMoM)*, pages 14–15. IEEE, 2018.

[3] Khaled Q Abdelfadeel, Dimitrios Zorbas, Victor Cionca, and Dirk Pesch. Free-fine-grained scheduling for reliable and energy efficient data collection in lorawan. *IEEE Internet of Things Journal*, 2019.

[4] Adaptrum. http://www.adaptrum.com/.

[5] Talal Ahmad, Ranveer Chandra, Ashish Kapoor, Michael Daum, and Eric Horvitz. Wi-fly: Widespread opportunistic connectivity via commercial air transport. In *Proceedings of the 16th ACM Workshop on Hot Topics in Networks*, HotNets-XVI, pages 43–49, New York, NY, USA, 2017. ACM.

[6] M. H. Almarshadi and S. M. Ismail. Effects of precision irrigation on productivity and water use efficiency of alfalfa under different irrigation methods in arid climates. *Journal of Applied Science Research*, pages 299–308,, 2011.

[7] Mohammed H. Almarshadi and Saleh M. Ismail. Effects of Precision Irrigation on Productivity and Water Use Efficiency of Alfalfa under Different Irrigation Methods in Arid Climates. *Journal of Applied Sciences Research*, 2011.

[8] Chris Anderson. Relatively cheap drones with advanced sensors and imaging capabilities are giving farmers new ways to increase yields and reduce crop damage.

[9] Pedro Andrade-Sanchez, Francis J Pierce, and Todd V Elliott. Performance assessment of wireless sensor networks in agricultural settings. In *2007 ASAE Annual Meeting*, page 1. American Society of Agricultural and Biological Engineers, 2007.

[10] Arcturus. Arcturus uav. http://arcturus-uav.com.

[11] AutoPano. kolor.com.

[12] A. Baggio. Wireless sensor networks in precision agriculture. *ACM Workshop on Real-World Wireless Sensor Networks*, 2005.

[13] Paramvir Bahl, Ranveer Chandra, Thomas Moscibroda, Rohan Murty, and Matt Welsh. White Space Networking with Wi-Fi like Connectivity. *ACM SIGCOMM Computer Communication Review*, 2009.

[14] Paramvir Bahl, Ranveer Chandra, Thomas Moscibroda, Rohan Murty, and Matt Welsh. White space networking with wi-fi like connectivity. In *Proceedings of the ACM SIGCOMM 2009 Conference on Data Communication*, SIGCOMM '09, pages 27–38, New York, NY, USA, 2009. ACM.

[15] C Baugh. 4ipp traffic model for ieee 802.16. 3. *IEEE802*, 16.

[16] C. Baugh and J. Huang. Traffic model for 802.16 tg3 mac/phy simulations. *Contribution 802.16.3c-01/30r1*, March 2001.

[17] John Bicket. Bit-rate selection in wireless networks. Master's thesis, Massachusetts Institute of Technology, 2005.

[18] Bluetooth. Bluetooth core specification. 5:2536, September.

[19] Martin Bor and Utz Roedig. Lora transmission parameter selection. In *2017 13th International Conference on Distributed Computing in Sensor Systems (DCOSS)*, pages 27–34. IEEE, 2017.

[20] Tom Botterill, Steven Mills, and Richard Green. Real-time aerial image mosaicing. In *Image and Vision Computing New Zealand (IVCNZ), 2010 25th International Conference of*, pages 1–8. IEEE, 2010.

[21] Mic Bowman, Saumya K. Debray, and Larry L. Peterson. Reasoning about naming systems. *ACM Trans. Program. Lang. Syst.*, 15(5):795–825, November 1993.

[22] Johannes Braams. Babel, a multilingual style-option system for use with latex's standard document styles. *TUGboat*, 12(2):291–301, June 1991.

[23] Matthew Brown and David G. Lowe. Automatic Panoramic Image Stitching Using Invariant Features. *International Journal of Computer Vision*, 2007.

[24] J. Cai and I. Lee. The stitching of aerial videos from uavs. In *2013 28th International Conference on Image and Vision Computing New Zealand (IVCNZ 2013)*, pages 448–452, Nov 2013.

[25] Joseph Camp and Edward Knightly. Modulation rate adaptation in urban and vehicular environments: Cross-layer implementation and experimental evaluation. ACM MobiCom, 2008.

[26] Kenneth G. Cassman. Ecological Intensification of Cereal Production Systems: Yield Potential, Soil Quality, and Precision Agriculture. *Proceedings of the National Academy of Sciences (PNAS)*, 1999.

[27] B. Chen and G.W. Wornell. Quantization index modulation: a class of provably good methods for digital watermarking and information embedding. In *IEEE TRANSACTION ON INFORMATION THEORY*, 2001.

[28] Min Chen, Yiming Miao, Xin Jian, Xiaofei Wang, and Iztok Humar. Cognitive-lpwan: Towards intelligent wireless services in hybrid low power wide area networks. *IEEE Transactions on Green Communications and Networking*, 3(2):409–417, 2018.

[29] Zicheng Chi, Yan Li, Hongyu Sun, Yao Yao, Zheng Lu, and Ting Zhu. B2w2: N-way concurrent communication for iot devices. In *Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems CD-ROM*, SenSys '16, pages 245–258, New York, NY, USA, 2016. ACM.

[30] Ji chun Zhao, J. Zhang, Y. Feng, and J. Guo. The Study and Application of the lOT Technology in Agriculture. *IEEE International Conference on Computer Science and Information Technology*, 2010.

[31] Malcolm Clark. Post congress tristesse. In *TeX90 Conference Proceedings*, pages 84–89. TeX Users Group, March 1991.

[32] Advanced Television Systems Committee. https://www.atsc.org/.

[33] Connected Farm. https://www.connectedfarm.com/.

[34] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley, 2012.

[35] Francesca Cuomo, Manuel Campo, Alberto Caponi, Giuseppe Bianchi, Giampaolo Rossini, and Patrizio Pisani. Explora: Extending the performance of lora by suitable spreading factor allocations. In *2017 IEEE 13th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*, pages 1–8. IEEE, 2017.

[36] Francesca Cuomo, Julio César Carrasquel Gámez, Antonio Maurizio, Laura Scipione, Manuel Campo, Alberto Caponi, Giuseppe Bianchi, Giampaolo Rossini, and Patrizio Pisani. Towards traffic-oriented spreading factor allocations in lorawan systems. In *2018 17th Annual Mediterranean Ad Hoc Networking Workshop (Med-Hoc-Net)*, pages 1–8. IEEE, 2018.

[37] Y. Dashdorj and Y. Won. Modeling and analysis of wireless lan traffic. *Journal of Information Science and Engineering*, 25:1783–1801, October 2009.

[38] DataMapper. http://www.precisionhawk.com/datamapperinflight.

[39] Harsha de Silva and Dimuthu Ratnadiwakara. Using ict to reduce transaction costs in agriculture through better communication: A case-study from sri lanka. *LIRNEasia*, 2008.

[40] Decagon Devices. Decagon Devices Cellular Logger. https://www.decagon.com/en/data-loggers-main/data-loggers/em50g-wireless-cellular-data-logger/.

[41] John Deere. John deere field connect. https://www.deere.com/en_US/products/equipment /ag_management_solutions/field_and_crop_solutions /john_deere_field_connect/john_deere_field_conne

[42] John Deere. John deere sensors. https://www.deere.com/en_US/products/equipment/sensors/sensors

[43] Analog Devices. Adg901 datasheet. `https://semtech.my.salesforce.com/sfc/p/#E0000000JelG/a/2R0000001OJa/2BF2MTeiqIwkmxkcjjDZzalPUGlJ76lLdqiv.30prH8`.

[44] Min Ding, Kristian Lyngbaek, and Avideh Zakhor. Automatic registration of aerial imagery with untextured 3d lidar models. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008.

[45] DJI. http://developer.dji.com.

[46] J. Doerflinger and T. Gross. Sustainable ICT in Agricultural Value Chains. *IT Professional*, 2012.

[47] Thomas A. Doerge. In *International Plant Nutritional Institute*, 1999.

[48] Adwait Dongare, Revathy Narayanan, Akshay Gadre, Anh Luong, Artur Balanuta, Swarun Kumar, Bob Iannucci, and Anthony Rowe. Charm: Exploiting geographical diversity through coherent combining in low-power wide-area networks. IPSN 2018. IEEE, 2018.

[49] Precision Drone. Agriculture drones - precision drone. http://www.precisiondrone.com/.

[50] DroneDeploy. dronedeploy.com.

[51] Rashad Eletreby, Diana Zhang, Swarun Kumar, and Osman Yaundefinedan. Empowering low-power wide area networks in urban settings. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, SIGCOMM '17, page 309–321, New York, NY, USA, 2017. Association for Computing Machinery.

[52] J. F. Ensworth, A. T. Hoang, and M. S. Reynolds. A low power 2.4 ghz superheterodyne receiver architecture with external lo for wirelessly powered backscatter tags and sensors. In *2017 IEEE International Conference on RFID (RFID)*, pages 149–154, May 2017.

[53] J. F. Ensworth and M. S. Reynolds. Ble-backscatter: Ultralow-power iot nodes compatible with bluetooth 4.0 low energy (ble) smartphones and tablets. *IEEE Transactions on Microwave Theory and Techniques*, 65(9):3360–3368, Sept 2017.

[54] Farms.com. Precision maps: http://www.farms.com/precision-agriculture/precision-maps/.

[55] FCC. https://ecfsapi.fcc.gov/file/6518909731.pdf.

[56] Federal Communications Commission. https://www.fcc.gov/general/white-space-database-administration.

[57] C. D. Franco and G. Buttazzo. Energy-aware coverage path planning of uavs. In *International Conference on Autonomous Robot Systems and Competitions (ICARSC),* , 2015.

[58] FreeWave. http://www.freewave.com/.

[59] Akshay Gadre, Revathy Narayanan, Anh Luong, Anthony Rowe, Bob Iannucci, and Swarun Kumar. Frequency configuration for low-power wide-area networks in a heartbeat. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*, pages 339–352, Santa Clara, CA, February 2020. USENIX Association.

[60] Chuhan Gao, Mehrdad Hessar, 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.

[61] Antonio-Javier Garcia-Sanchez, Felipe Garcia-Sanchez, and Joan Garcia-Haro. Wireless Sensor Network Deployment for Integrating Video-surveillance and Data-monitoring in Precision Agriculture over Distributed Crops. *Computers and Electronics in Agriculture*, 2011.

[62] Gartner. http://www.gartner.com/newsroom/id/2636073.

[63] GNURadio. About gnu radio. `https://www.gnuradio.org/about/`.

[64] H. Charles J. Godfray, John R. Beddington, Ian R. Crute, Lawrence Haddad, David Lawrence, James F. Muir, Jules Pretty, Sherman Robinson, Sandy M. Thomas, and Camilla Toulmin. Food Security: The Challenge of Feeding 9 Billion People. *Science*, 2010.

[65] Shyamnath Gollakota et al. Clearing the rf smog. In *Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM)*. Association for Computing Machinery, 2011.

[66] Aditya Gudipati and Sachin Katti. Automatic rate adaptation. ACM HotNets, 2010.

[67] Xiuzhen Guo, Xiaolong Zheng, and Yuan He. Wizig: Cross-technology energy communication over a noisy channel. In *The 36th Annual IEEE International Conference on Computer Communications (INFOCOM)*, 2017.

[68] B. R. Hanson and S. Orloff. Monitoring soil moisture for irrigation water management. Technical report, UC Davis, 2007.

[69] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.

[70] Hauppauge. http://www.hauppauge.com/site/products/data$_h vr955q.html$.

[71] Jetmir Haxhibeqiri, Ingrid Moerman, and Jeroen Hoebeke. Low overhead scheduling of lora transmissions for improved scalability. *IEEE Internet of Things Journal*, 6(2):3097–3109, 2018.

[72] Maurice Herlihy. A methodology for implementing highly concurrent data objects. *ACM Trans. Program. Lang. Syst.*, 15(5):745–770, November 1993.

[73] T. Hirsch, P. Sengers, E. Blevis, R. Beckwith, and T. Parikh. Making food, producing sustainability. In *CHI'10 Extended Abstracts on Human Factors in Computing Systems (CHI EA '10)*, pages 3147–3150. ACM, 2010.

[74] Gavin Holland, Nitin Vaidya, and Paramvir Bahl. A rate-adaptive mac protocol for multi-hop wireless networks. ACM MobiCom, 2001.

[75] Jason Hsu, Sadaf Zahedi, Aman Kansal, Mani Srivastava, and Vijay Raghunathan. Adaptive Duty Cycling for Energy Harvesting Systems. In *International Symposium on Low Power Electronics and Design*, 2006.

[76] Y. Hu and P. C. Loizou. Evaluation of objective quality measures for speech enhancement. *IEEE Transactions on Audio, Speech, and Language Processing*, 16(1):229–238, Jan 2008.

[77] Y. Hu and P. C. Loizou. Evaluation of objective quality measures for speech enhancement. *IEEE Transactions on Audio, Speech, and Language Processing*, 16(1):229–238, Jan 2008.

[78] Texas Instruments. http://www.ti.com/product/msp432p401r.

[79] International Food Policy Research Institute. Agricultural technologies could increase global crop yields as much as 67 percent and cut food prices nearly in half by 2050, 2014.

[80] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *CoRR*, abs/1502.03167, 2015.

[81] 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 ACM SIGCOMM Conference*, SIGCOMM '16, pages 356–369, New York, NY, USA, 2016. ACM.

[82] Junsu Jang and Fadel Adib. Underwater backscatter networking. In *Proceedings of the ACM Special Interest Group on Data Communication*, SIGCOMM '19, page 187–199, New York, NY, USA, 2019. Association for Computing Machinery.

[83] J. B. Johnson. Thermal agitation of electricity in conductors. *Phys. Rev.*, 32:97–109, Jul 1928.

[84] O. Visser K. Langendoen, A. Baggio. Murphy loves Potatoes: Experience from a Pilot Sensor Network Deployment in Precision Agriculture. In *International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 2006.

[85] Kansas State University. Weak internet connectivity in rural areas hindering agricultural production, 2016.

[86] Z. Kapetanovic, R. Chandra, T. Chakraborty., and A. Nelson. Improving farm productivity using data-driven agriculture. *SIAM News*, 2019.

[87] Z. Kapetanovic, V. Talla, A. Parks, J. Qian, and J. R. Smith. Wireless quantization index modulation: Enabling communication through existing signals. In *2018 IEEE International Conference on RFID (RFID)*, pages 1–8, April 2018.

[88] Zerina Kapetanovic, Ali Saffari, Ranveer Chandra, and Joshua R. Smith. Glaze: Overlaying occupied spectrum with downlink iot transmissions. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 3(4), dec 2019.

[89] Zerina Kapetanovic and Joshua R. Smith. Communication by means of modulated johnson noise. *CoRR*, abs/2111.08629, 2021.

[90] Zerina Kapetanovic, Deepak Vasisht, Jongho Won, Ranveer Chandra, and Mark Kimball. Experiences deploying an always-on farm network. *GetMobile: Mobile Comp. and Comm.*, 21(2):16–21, aug 2017.

[91] Mohamad Katanbaf, Ali Saffari, and Joshua R. Smith. Multiscatter: Multistatic backscatter networking for battery-free sensors. In *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems*, SenSys '21, page 69–83, New York, NY, USA, 2021. Association for Computing Machinery.

[92] Kathleen McLaughlin . Gaps in 4G Network Hinder High-tech Agriculture: FCC Prepares to Release 500 Million to Improve Coverage. http://www.bendbulletin.com/newsroomstafflist/4535283-151/gaps-in-4g-network-hinder-high-tech-agriculture.

[93] Bryce Kellogg, Aaron Parks, Shyamnath Gollakota, Joshua R. Smith, and David Wetherall. Wi-fi backscatter: Internet connectivity for rf-powered devices. *SIGCOMM Comput. Commun. Rev.*, 44(4):607–618, August 2014.

[94] Bryce Kellogg, Vamsi Talla, Shyamnath Gollakota, and Joshua R Smith. Passive wi-fi: Bringing low power to wi-fi transmissions. In *NSDI*, volume 16, pages 151–164, 2016.

[95] Hak-Jin Kim, Kenneth A. Sudduth, and John W. Hummel. Soil Micronutrient Sensing for Precision Agriculture. *Journal of Environmental Monitoring, The Royal Society of Chemistry*, 2009.

[96] Song Min Kim and Tian He. Freebee: Cross-technology communication via free side-channel. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*, MobiCom '15, pages 317–330, New York, NY, USA, 2015. ACM.

[97] Song Min Kim, Shuai Wang, and Tian He. cetx: Incorporating spatiotemporal correlation for better wireless networking. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*, SenSys '15, pages 323–336, New York, NY, USA, 2015. ACM.

[98] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014.

[99] Laszlo B Kish. Stealth communication: Zero-power classical communication, zero-quantum quantum communication and environmental-noise communication. *Applied physics letters*, 87(23):234109, 2005.

[100] R. Kumar, H. Sawhney, S. Samarasekera, S. Hsu, Hai Tao, Yanlin Guo, K. Hanna, A. Pope, R. Wildes, D. Hirvonen, M. Hansen, and P. Burt. Aerial video surveillance and exploitation. *Proceedings of the IEEE*, 89(10):1518–1539, Oct 2001.

[101] Vivek Kwatra, Arno Schödl, Irfan Essa, Greg Turk, and Aaron Bobick. Graphcut textures: image and video synthesis using graph cuts. In *ACM Transactions on Graphics (ToG)*, volume 22, pages 277–286. ACM, 2003.

[102] Leslie Lamport. *LaTeX User's Guide and Document Reference Manual*. Addison-Wesley Publishing Company, Reading, Massachusetts, 1986.

[103] "W.S. Lee, V. Alchanatis, C. Yang, M. Hirafuji, D. Moshou, and C. Li. Sensing Technologies for Precision Specialty Crop Production. *Computers and Electronics in Agriculture*, 2010.

[104] Y. Lee, G. Kim, S. Bang, Y. Kim, I. Lee, P. Dutta, D. Sylvester, and D. Blaauw. A modular 1mm3 die-stacked sensing platform with optical communication and multi-modal energy harvesting. *IEEE International Solid-State Circuits*, pages 402 – 404, 2012.

[105] K. B. Letaief, W. Chen, Y. Shi, J. Zhang, and Y. A. Zhang. The roadmap to 6g: Ai empowered wireless networks. *IEEE Communications Magazine*, 57(8):84–90, 2019.

[106] Philip Levis. Experiences from a decade of tinyos development. In *Proceedings of the 10th USENIX Conference on Operating Systems Design and Implementation*, OSDI'12, pages 207–220, Berkeley, CA, USA, 2012. USENIX Association.

[107] Hengyang Li, Usman Raza, and Aftab Khan. How agile is the adaptive data rate mechanism of lorawan? *CoRR*, 2018.

[108] Yan Li, Zicheng Chi, Xin Liu, and Ting Zhu. Chiron: Concurrent high throughput communication for iot devices. In *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services*, MobiSys '18, pages 204–216, New York, NY, USA, 2018. ACM.

[109] Z. Li and V. Isler. Large Scale Image Mosaic Construction for Agricultural Applications. *IEEE Robotics and Automation Letters*, 2016.

[110] Vincent Liu, Aaron Parks, Vamsi Talla, Shyamnath Gollakota, David Wetherall, and Joshua R. Smith. Ambient backscatter: Wireless communication out of thin air. In *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM*, SIGCOMM '13, pages 39–50, New York, NY, USA, 2013. ACM.

[111] Digital Logger. PoE Switch. http://www.digital-loggers.com/poe48.html.

[112] LoRa Technology. https://www.lora-alliance.org/what-is-lora/technology.

[113] Jess Lowenberg-DeBoer. The Precision Agriculture Revolution: Making the Modern Farmer. https://www.foreignaffairs.com/articles/united-states/2015-04-20/precision-agriculture-revolution.

[114] Alan Mainwaring, David Culler, Joseph Polastre, Robert Szewczyk, and John Anderson.

Wireless sensor networks for habitat monitoring. In *Proceedings of the 1st ACM international workshop on Wireless sensor networks and applications*, pages 88–97. ACM, 2002.

[115] J. Matos, S. Gama, H. Ruskin, A. Al Sharkasi, and M. Crane. Time and scale hurst exponent analysis for financial markets. *Physica A: Statistical Mechanics and its Applications*, 387:3910–3915, June 2008.

[116] A. McBratney and MJ Pringle. Estimating average and proportional variograms of soil properties and their potential use in precision agriculture. *Precision Agriculture*, 1999.

[117] A. McBratney, Brett Whalen, Tihomir Ancev, and Johan Bouma. Future Directions of Precision Agriculture. *Precision Agriculture*, 2005.

[118] Xiangyun Meng, Wei Wang, and Ben Leong. Skystitch: A cooperative multi-uav-based real-time video surveillance system with stitching. In *Proceedings of the 23rd ACM International Conference on Multimedia*, MM '15, pages 261–270, New York, NY, USA, 2015. ACM.

[119] Microseven. Ip cameras. http://www.microseven.com/product/IP-Cameras.html.

[120] Microsoft. Azure IoT Hub. https://azure.microsoft.com/en-us/services/iot-hub/.

[121] Microsoft. Azure IoT Suite. https://www.microsoft.com/en-us/cloud-platform/internet-of-things-azure-iot-suite.

[122] Microsoft. https://news.microsoft.com/2003/01/09/microsoft-presents-smart-personal-objects-technology-spot-based-wristwatches-at-ces/.

[123] Microsoft. Project premonition, http://research.microsoft.com/en-us/um/redmond/projects/projectpremonition/default.aspx.

[124] Microsoft Research. Image Composite Editor. http://research.microsoft.com/en-us/um/redmond/projects/ice/.

[125] MiniCircuits. Vbf-1445+ bandpass filter. https://www.minicircuits.com/pdfs/VBF-1445+.pdf.

[126] MiniCircuits. Zhl-1217hln+ low noise amplifier. https://www.minicircuits.com/pdfs/ZHL-1217HLN+.pdf.

[127] Pierre Moulin and Ralf Koetter. *Data Hiding: Theory and Algorithms*, pages 131–176. WORLD SCIENTIFIC, 2011.

[128] Nathaniel D. Mueller, James S. Gerber, Matt Johnston, Deepak K. Ray, and Jonathan A. Foley Navin Ramankutty. Closing Yield Gaps through Nutrient and Water management. *Nature*, 2012.

[129] Razvan Musaloiu-E and Andreas Terzis. Minimising the effect of wifi interference in 802.15. 4 wireless sensor networks. *International Journal of Sensor Networks*, 3(1):43–54, 2008.

[130] S. Naderiparizi, A. N. Parks, Z. Kapetanovic, B. Ransford, and J. R. Smith. Wispcam : A battery-free rfid camera. In *in Proc. IEEE RFID*, 2015.

[131] Saman Naderiparizi, Aaron N. Parks, Zerina Kapetanovic, Benjamin Ransford, and Joshua R. Smith. Wispcam: A battery-free rfid camera. In *2015 IEEE International Conference on RFID (RFID)*, pages 166–173, 2015.

[132] L. Nagl, R. Schmitz, S. Warren, T. S. Hildreth, H. Erickson, and D. Andresen. Wearable Sensor System for Wireless State-of-health Determination in Cattle. In *Engineering in Medicine and Biology Society*, 2003.

[133] NooElec. Nesdr smart xtr. https://www.nooelec.com/store/sdr/sdr-receivers/nesdr-smart-xtr.html.

[134] Nooelec. Nesdr smart xtr sdr. https://www.nooelec.com/store/nesdr-smart-xtr-sdr.html.

[135] Nooelec. Sawbird+ h1. https://www.nooelec.com/store/sdr/sdr-addons/sawbird/sawbird-h1.html.

[136] Tamoghna Ojha, Sudip Misra, and Narendra Singh Raghuwanshi. Wireless sensor networks for agriculture. *Computers and Electronics in Agriculture*, 2015.

[137] OpenWeatherMap, Inc. OpenWeather API. http://openweathermap.org/.

[138] Brenda Ortiz, Joey Shaw, John fulton, and Amy Winstead. Management Zones II – Basic Steps for Delineation. *Alabama Precision Ag Extension: Precision Agriculture Series*, 2011.

[139] Timothy J O'Shea, Johnathan Corgan, and T. Charles Clancy. Convolutional radio modulation recognition networks, 2016.

[140] J. Petajajarvi, K. Mikhaylov, A. Roivainen, T. Hanninen, and M. Pettissalo. On the coverage of lpwans: range evaluation and channel attenuation model for lora technology. In *2015 14th International Conference on ITS Telecommunications (ITST)*, pages 55–59, 2015.

[141] PIX4D. Pix4d. https://pix4d.com/.

[142] Mary Pols. To run their businesses, farmers need reliable internet connections.

[143] David M. Pozar. *Microwave Engineering*. John Wiley & Sons, Inc., 2012.

[144] Precision Planting. Ag Insights. http://www.precisionplanting.com/ag_insights/.

[145] Jacopo Primicerio, Salvatore Filippo Di Gennaro, Edoardo Fiorillo, Lorenzo Genesio, Emanuele Lugato, Alessandro Matese, and Francesco Primo Vaccari. A flexible unmanned aerial vehicle for precision agriculture. *Precision Agriculture*, 13(4):517–523, 2012.

[146] PyPi. Python rtl-sdr library. https://pypi.org/project/pyrtlsdr/.

[147] Ranch Systems. http://marius.ranchsystems.com/wp/rm210/.

[148] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.

[149] B Razavi. *RF Microelectronics*. Prentice Hall.

[150] Dark Reading. Emerging long-range wan networks vulnerable to hacking, compromise. `https://www.darkreading.com/iot/emerging-long-range-wan-networks-vulnerable-to-hacking-compromise/d/d-id/1336899`.

[151] Etus Research. https://www.ettus.com/product/details/x300-kit.

[152] Brecht Reynders, Wannes Meert, and Sofie Pollin. Power and spreading factor control in low power wide area networks. In *2017 IEEE International Conference on Communications (ICC)*, pages 1–6. IEEE, 2017.

[153] Brecht Reynders, Qing Wang, Pere Tuset-Peiro, Xavier Vilajosana, and Sofie Pollin. Improving reliability and scalability of lorawans through lightweight scheduling. *IEEE Internet of Things Journal*, 5(3):1830–1842, 2018.

[154] S. Roberts, P. Garnett, and R. Chandra. Connecting Africa Using the TV White Spaces: From Research to Real World Deployments. In *IEEE LANMAN*, 2015.

[155] Mark W. Rosegrant, Jawoo Koo, Nicola Cenacchi, Claudia Ringler, Richard D. Robertson, Myles Fisher, Cindy M. Cox, Karen Garrett, Nicostrato D. Perez, and Pascale Sabbagh. *Food security in a world of natural resource scarcity: The role of agricultural technologies*. International Food Policy Research Institute (IFPRI), 2014.

[156] L. Ruiz-Garcia, L. Lunadei, P. Barreiro, and J. Ignacio Robla. A Review of Wireless Sensor Technologies and Applications in Agriculture and Food Industry: State of the Art and Current Trends. *Sensors*, 2009.

[157] Naseer Sabri, S. A. Aljunid, M. S. Salim, S. Fouad, and R. Kamaruddin. *Recent Trends in Physics of Material Science and Technology*, chapter Wireless Sensor Network Wave Propagation in Vegetation, pages 283–298. Springer Singapore, Singapore, 2015.

[158] A. Sahai, J. Sanz, V. Subramanian, C. Tran, and K. Vodrahalli. Blind interactive learning of modulation schemes: Multi-agent cooperation without co-design. *IEEE Access*, 8:63790–63820, 2020.

[159] S.L. Salas and Einar Hille. *Calculus: One and Several Variable*. John Wiley and Sons, New York, 1978.

[160] A. P. Sample, D. J. Yeager, P. S. Powledge, A. V. Mamishev, and J. R. Smith. Design of an rfid-based battery-free programmable sensing platform. *IEEE Transactions on Instrumentation and Measurement*, 57(11):2608–2615, Nov 2008.

[161] Aaron R. Schepers, John F. Shanahan, Mark A. Liebig, James S. Schepers, Sven H. Johnson, and Ariovaldo Luchiari. Appropriateness of Management Zones for Characterizing Spatial Variability of Soil Properties and Irrigated Corn Yields across Years. *Agronomy Journal*, 2004.

[162] David Schimmelpfennig. Cost Savings From Precision Agriculture Technologies on U.S. Corn Farms. *USDA Amber Waves*, 2016.

[163] Campbell Scientific. https://www.campbellsci.com/products.

[164] Leo Selavo, Anthony Wood, Qing Cao, Tamim Sookoor, Hengchang Liu, Aravind Srinivasan, Yafeng Wu, Woochul Kang, John Stankovic, Don Young, et al. Luster: wireless sensor network for environmental research. In *Proceedings of the 5th international conference on Embedded networked sensor systems*, pages 103–116. ACM, 2007.

[165] Semtech. Lora modulation basics. `https://semtech.my.salesforce.com/sfc/p/#E0000000JelG/a/2R0000001OJa/2BF2MTeiqIwkmxkcjjDZzalPUGlJ76lLdqiv.30prH8`.

[166] Semtech. Low Power Digital I and Q RF Multi-PHY Mode Analog Front End 860-1000 MHz. https://www.semtech.com/products/wireless-rf/lora-gateways/sx1257.

[167] Semtech. Semtech sx1276 137 mhz to 1020 mhz long range low power transceiver. https://www.semtech.com/products/wireless-rf/lora-transceivers/sx1276.

[168] Semtech. Sx1261/sx1262 datasheet. https://semtech.my.salesforce.com/sfc/p/#E0000000JelG/a/2R000000HT76/7Nka9W5WgugoZe.xwIHJy6ebj1hW8UJ.USO_Pt2CLLo.

[169] Semtech. Sx1301 datasheet. https://semtech.my.salesforce.com/sfc/p/#E0000000JelG/a/44000000MDnR/Et1KWLCuNDI6MDagfSPAvqqp.Y869Flgs1LleWyfjDY.

[170] Semtech. Understanding adr. https://lora-developers.semtech.com/library/tech-papers-and-guides/understanding-adr/.

[171] LoRa SemTech. https://lora-developers.semtech.com/documentation/tech-papers-and-guides/lora-and-lorawan/.

[172] SenseFly. Sensefly. sensefly.com.

[173] Sentera. https://sentera.com/.

[174] A. Shahid, J. Fontaine, M. Camelo, J. Haxhibeqiri, M. Saelens, Z. Khan, I. Moerman, and E. D. Poorter. A convolutional neural network approach for classification of lpwan technologies: Sigfox, lora and ieee 802.15.4g. In *2019 16th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*, pages 1–8, 2019.

[175] Michael Shamos. *Computational Geometry*. Yale University, 1978.

[176] Navin Sharma, Jeremy Gummeson, David E. Irwin, and Prashant J. Shenoy. Cloudy Computing: Leveraging Weather Forecasts in Energy Harvesting Sensor Systems. In *IEEE SECON*, 2010.

[177] SIGFOX. http://www.sigfox.com/.

[178] Sudipta N Sinha and Marc Pollefeys. Pan–tilt–zoom Camera Calibration and High-resolution Mosaic Generation. *Computer Vision and Image Understanding*, 2006.

[179] Skyfront. http://skyfront.com.

[180] Joshua R. Smith and Barrett O. Comiskey. Modulation and information hiding in images. In Ross Anderson, editor, *Information Hiding*, pages 207–226, Berlin, Heidelberg, 1996. Springer Berlin Heidelberg.

[181] M. Soderstrom, Gustav Sohlenius, Lars Rodhe, and Kristin Piikki. Adaptation of Regional digital soil mapping for precision agriculture. *Precision Agriculture*, 2016.

[182] John V. Stafford. *Precision Agriculture Book*. 2005.

[183] STMicroelectronics. Stm32l151/152. `https://www.st.com/en/microcontrollers-microprocessors/stm32l151-152.html`.

[184] Harry Stockman. Communication by means of reflected power. *Proceedings of the IRE*, 36(10):1196–1204, 1948.

[185] Richard Szeliski. Image Alignment and Stitching: A Tutorial. Technical report, Microsoft Research, 2004.

[186] Richard Szeliski, Matthew Uyttendaele, and Drew Steedly. Fast poisson blending using multi-splines. In *Computational Photography (ICCP), 2011 IEEE International Conference on*, pages 1–8. IEEE, 2011.

[187] Wark T., P. Csiro, P. Corke, Sikka, L. Klingbeil, Y. Guo, C. Crossman, P. Valencia, D. Swain, and G. Bishop-Hurley. Transforming Agriculture through Pervasive Wireless Sensor Networks. *IEEE Pervasive Computing*, 2007.

[188] Vamsi Talla, Mehrdad Hessar, Bryce Kellogg, Ali Najafi, Joshua R. Smith, and Shyamnath Gollakota. Lora backscatter: Enabling the vision of ubiquitous connectivity. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 1(3):105:1–105:24, September 2017.

[189] Vamsi Talla, Bryce Kellogg, Shyamnath Gollakota, and Joshua R. Smith. Battery-free cellphone. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 1(2), jun 2017.

[190] Joe Taylor. Moonbounce at w2pu - adaptive polarization at 432 mhz. `https://physics.princeton.edu/pulsar/K1JT/Moonbounce_at_W2PU.pdf`.

[191] LoRa Technology. https://www.lora-alliance.org/what-islora/technology.

[192] M. Tennøe, E. Helgedagsrud, M. Næss, H. K. Alstad, H. K. Stensland, V. R. Gaddam, D. Johansen, C. Griwodz, and P. Halvorsen. Efficient implementation and processing of a real-time panorama video pipeline. In *Multimedia (ISM), 2013 IEEE International Symposium on*, pages 76–83, Dec 2013.

[193] Tensorflow. An end-to-end open source machine learning platform. `https://www.tensorflow.org`.

[194] The Weather Channel. https://weather.com/.

[195] N. Thomos, N. V. Boulgouris, and M. G. Strintzis. Optimized transmission of jpeg2000 streams over wireless channels. *IEEE Transactions on Image Processing*, 15(1):54–67, Jan 2006.

[196] United Nations General Assembly. Food Production Must Double by 2050 to Meet Demand from World's Growing Population, Innovative Strategies Needed to Combat Hunger, Experts Tell Second Committee, 2009. http://www.un.org/press/en/2009/gaef3242.doc.htm.

[197] UP-Board. Up xtreme. `https://up-shop.org/up-xtreme-series.html`.

[198] Gabriel Vasilescu. *Electronic noise and interfering signals: principles and applications*. Springer Science and Business Media, 2006.

[199] Deepak Vasisht, Zerina Kapetanovic, Jongho Won, Xinxin Jin, Ranveer Chandra, Sudipta Sinha, Ashish Kapoor, Madhusudhan Sudarshan, and Sean Stratman. FarmBeats: An IoT

platform for Data-Driven agriculture. In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*, pages 515–529, Boston, MA, March 2017. USENIX Association.

[200] Christopher M. Vigorito, Deepak Ganesan, and Andrew G. Barto. Adaptive Control of Duty Cycling in Energy-harvesting Wireless Sensor Networks. In *IEEE SECON*, 2007.

[201] Mythili Vutukuru, Hari Balakrishnan, and Kyle Jamieson. Cross-layer wireless bit rate adaptation. In *Proceedings of the ACM SIGCOMM 2009 Conference on Data Communication*, SIGCOMM '09, page 3–14, New York, NY, USA, 2009. Association for Computing Machinery.

[202] Mythili Vutukuru, Hari Balakrishnan, and Kyle Jamieson. Cross-layer wireless bit rate adaptation. ACM SIGCOMM, 2009.

[203] Anran Wang, Vikram Iyer, Vamsi Talla, Joshua R. Smith, and Shyamnath Gollakota. Fm backscatter: Enabling connected cities and smart fabrics. In *Proceedings of the 14th USENIX Conference on Networked Systems Design and Implementation*, NSDI'17, pages 243–258, Berkeley, CA, USA, 2017. USENIX Association.

[204] WebSDR. http://www.websdr.org.

[205] Geoff Werner-Allen, Konrad Lorincz, Jeff Johnson, Jonathan Lees, and Matt Welsh. Fidelity and Yield in a Volcano Monitoring Sensor Network. In *Operating Systems Design and Implementation*, 2006.

[206] WVURail. Dspira horn assembly. `https://wvurail.org/cra/Files_uploaded/DSPIRA_Horn_Assembly_2019b.pdf`.

[207] WVURail. Dspira horn assembly. `http://wvurail.org/cra/Files_uploaded/DSPIRA_Horn_Assembly_2019.pdf`.

[208] X. Xie, Z. Xu, and H. Xie. Channel capacity analysis of spread spectrum watermarking in radio frequency signals. *IEEE Access*, PP(99):1–1, 2017.

[209] W. Xu and J. Mulligan. Performance evaluation of color correction approaches for automatic multi-view image and video stitching. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 263–270, June 2010.

[210] Chunhua Zhang and John M. Kovacs. The application of small unmanned aerial systems for precision agriculture: a review. *Precision Agriculture*, 13(6):693–712, 2012.

[211] X. Zhang and E. W. Knightly. Watch: Wifi in active tv channels. *IEEE Transactions on Cognitive Communications and Networking*, 2(4):330–342, Dec 2016.