A Low-Cost, Scalable Platform for Sub-Centimeter UHF RFID Positioning

by

Isaac Perper

S.B., Electrical Engineering and Computer Science
S.B., Mechanical Engineering
Massachusetts Institute of Technology, 2020

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2021

Author .....
Department of Electrical Engineering and Computer Science
August 15, 2021

Certified by .....
Fadel Adib
Associate Professor
Thesis Supervisor

Accepted by .....
Katrina LaCurts
Chair, Master of Engineering Thesis Committee

A Low-Cost, Scalable Platform for Sub-Centimeter UHF RFID Positioning

by

Isaac Perper

Submitted to the Department of Electrical Engineering and Computer Science
on August 15, 2021, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

High-precision localization is a technology that has started to make areas from autonomous vehicles to packing robots more efficient and accurate. While there are many approaches to localization, RFID micro-location is a growing technology that has been shown to be fast and robust, and can leverage the existing infrastructure of billions of RFID tags. However, many prior RFID positioning systems lack portability, scalabiltiy, and cost-effectiveness. In this thesis, I explore how low-cost software-defined radios can be leveraged to overcome those three key issues with RFID localization. I contribute a low-cost, scalable, and portable RFID micro-location platform that can overcome real-world deployment issues such as RFID orientation. Finally, I conclude with a characterization of the platform and a novel application of the system for robotic grasping.

Thesis Supervisor: Fadel Adib
Title: Associate Professor

Acknowledgments

I would like to thank all the members of the Signal Kinetics group for including me in the lab family over the past year, supporting me before deadlines, and sharing wonderful memories. In particular, my year would not have been the same without working late nights alongside my colleague Tara Boroushaki, and I will miss my adventures with Mergen Nachin in and out of the lab. Most importantly, I'd like to thank my supervisor, Fadel Adib, for his continual guidance and encouragement throughout the entire year. I am also grateful for the support of my family and my girlfriend, all of whom have helped me grow as student, researcher and person.

The research is sponsored by the National Science Foundation, the MIT Media Lab, NTT DATA, Toppan, Toppan Forms, and Abdul Latif Jameel Water and Food Systems Lab (J-WAFS).

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Finding objects with high precision, also known as micro-localization, is an important and growing area of research and deployment across many industries. While conventional localization, such as outdoor GPS or indoor beacon-based ranging devices, has been improved for many decades to provide meter-scale positioning [1], the latest research in micro-localization produces centimeter-scale positioning that could enable entirely new tasks. For example, with micro-localization, autonomous robots would be able to quickly navigate warehouses, precisely track inventory locations, and safely interact with humans [33]. In fulfillment centers, high-accuracy positioning could overcome many of the biggest challenges faced when deploying and coordinating packing robots [36]. However, designing micro-localization technologies is a challenging task that must consider many factors that affect real-world usage including power requirements, bulkiness, weight, ease-of-setup, cost, etc.

One technology that has made many in-roads in recent years is radio-frequency identification (RFID) positioning. RFIDs themselves are small, batteryless, 3 cent tags that can stick onto almost any object. RFIDs can store data, such as a unique ID or SKU, that can be read wirelessly (even through other objects) with an RFID reader. Most importantly, billions of RFIDs are already used across many industries for tracking all types of goods [19]. Thus, as radio-frequency (RF) technologies have improved, there has been a natural shift towards using these same tags to track location as well, and state-of-the-art research can locate RFIDs with sub-centimeter

accuracy [18]. However, few of the RFID micro-localization systems have made their way into the real-world because they often rely on expensive RF hardware, require extensive calibration, and have limited scalability. Additionally, RFID localization is impacted by real-world factors, such as how the RFID tags are oriented, RF power limits, and sensing range.

This thesis aims to bring RFID micro-localization closer to real-world deployments by addressing some of the challenges above. In particular, I build upon the state-of-the-art RFID positioning system and advance it through three key contributions:

- **Developing a Low-Cost Portable Platform:** I demonstrate how centimeter-scale RFID micro-localization can be implemented on portable, low-cost RF hardware in a scalable way.

- **Understanding the Impact of Orientation:** In the real-world, RFID tags can be placed on many different objects in different orientations. This thesis shows how the orientation of a RFID tag impacts the localization accuracy of the system.

- **System Characterization:** Finally, I characterize the performance of the localization system across different RFID tags, RF power levels, and operation frequencies. Here, I also explore how the micro-location system can operate under FCC regulations, which impact the allowed frequencies and transmit powers.

Beyond the above contributions, I demonstrate my low-cost platform in a novel application for robotic sensing and grasping via sensor fusion.

Chapter 2

Related Work

In this chapter, I discuss an overview of RFID localization approaches and cost-effective software-defined radios.

2.1 RFID Localization

RFID localization has been a topic of research over the past two decades, but only in the last few years have multipath-robust and centimeter-accurate positioning of UHF tags been developed. Early attempts at localization focused on signal strength (e.g., RSSI) [3, 7, 27, 23, 13, 38] and angle-of-arrival based approaches [2, 37, 14], but many of these systems also required an extensive infrastructure of reference tags [4, 35, 22, 8]. More recent work has investigated other uses of RFID signal phase [28, 17, 32, 21, 12, 29] or tag/antenna movement [6, 20, 35] to improve the accuracy. Still, for real world deployments, issues such as multipath significantly impact the robustness and accuracy of these systems.

This thesis builds on the foundational work in Turbotrack [16] and RFind [18], which use a wide emulated bandwidth and super-resolution to achieve sub-centimeter positioning accuracy even in multipath-rich environments. Although Clester made improvements to the speed and accuracy of the RFind techniques [9], all these systems still suffered from issues by using expensive radio hardware, relying on fixed calibration techniques, and failing to be robust to tag orientation changes. This the-

sis builds on this line of work and is the first to demonstrate it on a low-cost, portable hardware platform, as well as overcome past issues with calibration and orientation.

2.2 Low-Cost SDRs

Although many papers have demonstrated RFID positioning with low-cost, off-the-shelf RFID readers, these RFID readers don't provide the raw data or signal processing capabilities that the multipath-robust, wideband approaches from RFind, Turbotrack and similar papers need. Instead, RFind and others used software defined radios (SDRs), such as the commonly used USRPs from Ettus Research [26]. Yet, USRPs are limited in their scalability and portability given that they cost thousands of dollars, require additional hardware for multi-device synchronization, and are too heavy and large for portable systems.

In contrast, for the past several years, small, low-cost SDR platforms have become available for hobby-communities, such as bladeRF, HackRF, and LimeSDR [25, 11, 15]. Many of these are open-source and can be used for general purpose wireless systems including RFID localization. They also can be an order of magnitude cheaper. Despite the benefits in size and cost, these systems tend to have less-developed capabilities than those of more expensive research SDRs. This thesis focuses on overcoming the challenges of multi-device synchronization, phase alignment, and other issues that come with working on a low-cost SDR.

Chapter 3

Portable, Low-Cost RFID Localization with BladeRFs

In this chapter, I explain how I implemented a portable, low-cost, and robust RFID localization system with sub-centimeter accuracy. I began by selecting the Nuand bladeRF to serve as the SDR platform for this research. However, working with low-cost, general purpose radio hardware required overcoming multiple challenges, which I highlight in this chapter, along with their solutions.

3.1 RFID Localization Primer

Before I describe my system, I will explain the basics of RFID operation and localization.

RFIDs are batteryless devices that operate wirelessly via a principle called backscatter. Figure 3-1 shows the traditional operation of a standard RFID tag on the left and wideband RFID localization on the right. There are three main steps to standard RFID operation:

1. The in-band (IB) TX antenna transmits a signal to the RFID tag while operating within the ISM band, which in the United States is 902-928 MHz.

2. The RFID absorbs the energy of the transmitted signal and powers an internal

circuit.

3. The RFID responds to the RX antenna via backscatter. To do so, the RFID tag switches between reflecting the transmitted signal and absorbing the transmitted signal (non-reflective region). This on-off procedure encodes bits of information that can be processed by the receiver.

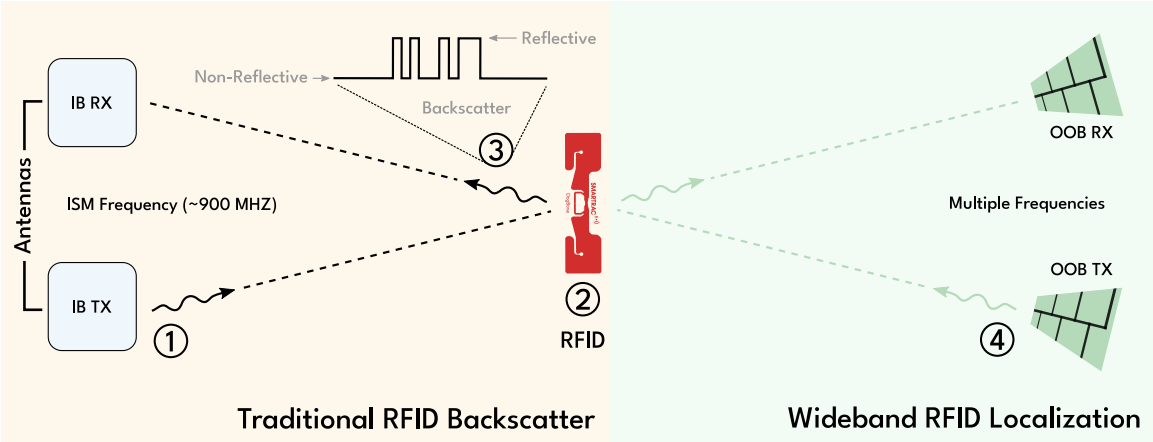

Figure 3-1: RFID Backscatter and Localization. RFIDs operate via backscatter, which is essentially reflecting and not reflecting a transmitted signal to communicate data to the receiver. Wideband RFID localization takes advantage of the frequency-agnostic nature of backscatter to measure the wireless channel at multiple frequencies.

The standard operation and uses of RFID backscatter are well researched, and RFID hardware is continually improving in both range and sensitivity. Emulated wideband RFID localization, like the technique found in RFind [18], takes advantage of the frequency-agnostic nature of backscatter. Once the RFID has powered up, step (4) in Figure 3-1 illustrates how the RFID will also reflect/not reflect any signal that is simultaneously transmitted, not just the ISM-band signal. To connect this with real data, Figure 3-2 shows the time-domain response of an RFID tag. Both the in-band (IB) and out-of-band (OOB) signals contain the reflective and non-reflective segments of the response. Note that the IB signal is not as noisy because the RFID is optimized for the IB frequencies.

At a high-level, RFind showed that one can receive the backscatter response of the RFID tag at multiple frequencies to form a wideband channel estimate. Combined with super-resolution techniques, this wideband channel estimate enables us to

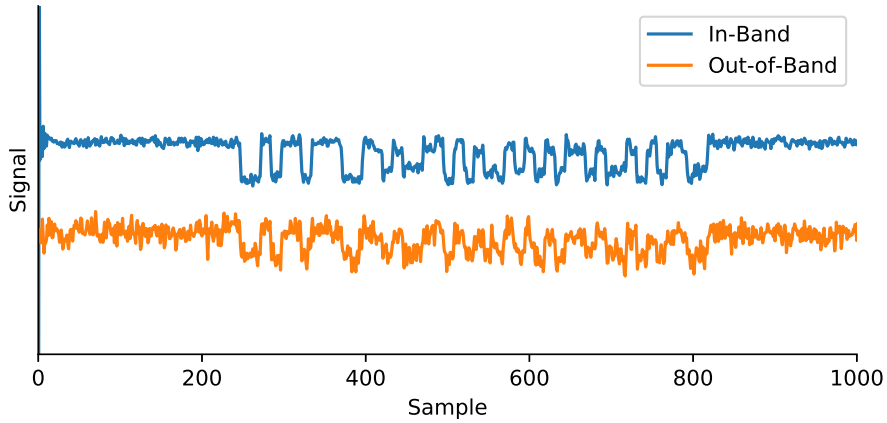

Figure 3-2: Sample IB and OOB Backscatter. This is a backscatter response of an RFID as received by both the in-band and out-of-band antennas. The RFID is optimized for the IB signal, so the response is cleaner with less noise compared to the OOB signal.

position RFIDs with sub-centimeter accuracy even in multipath-rich environments. The focus of this chapter will be on implementing the RFind approach on bladeRFs, so additional details on the wideband localization techniques can be found within the RFind paper.

3.2 Overview

With the general approach to RFID localization in mind, I begin by explaining high-level system architecture and challenges. Figure 3-3 shows the system architecture. Like the architecture shown in the primer, this architecture has one bladeRF to send in-band signals that power up the RFID and another bladeRF to send out-of-band signals for wideband channel estimates and localization. The faded bladeRF device depicts how the system can be scaled to include additional OOB bladeRFs to provide more accuracy and robustness to the localization system or increase the spatial coverage of the system to new areas. Research SDRs (e.g., USRPs) often have a straightforward way of synchronizing signals across frequency, time, and phase. However, this synchronization is more challenging with the bladeRF platform, and the dashed grey boxes highlight the clock synchronization and phase alignment sections

discussed in this chapter. Finally, Chapters 4 & 5 will discuss the importance of antennas and their impact on orientation and frequency.

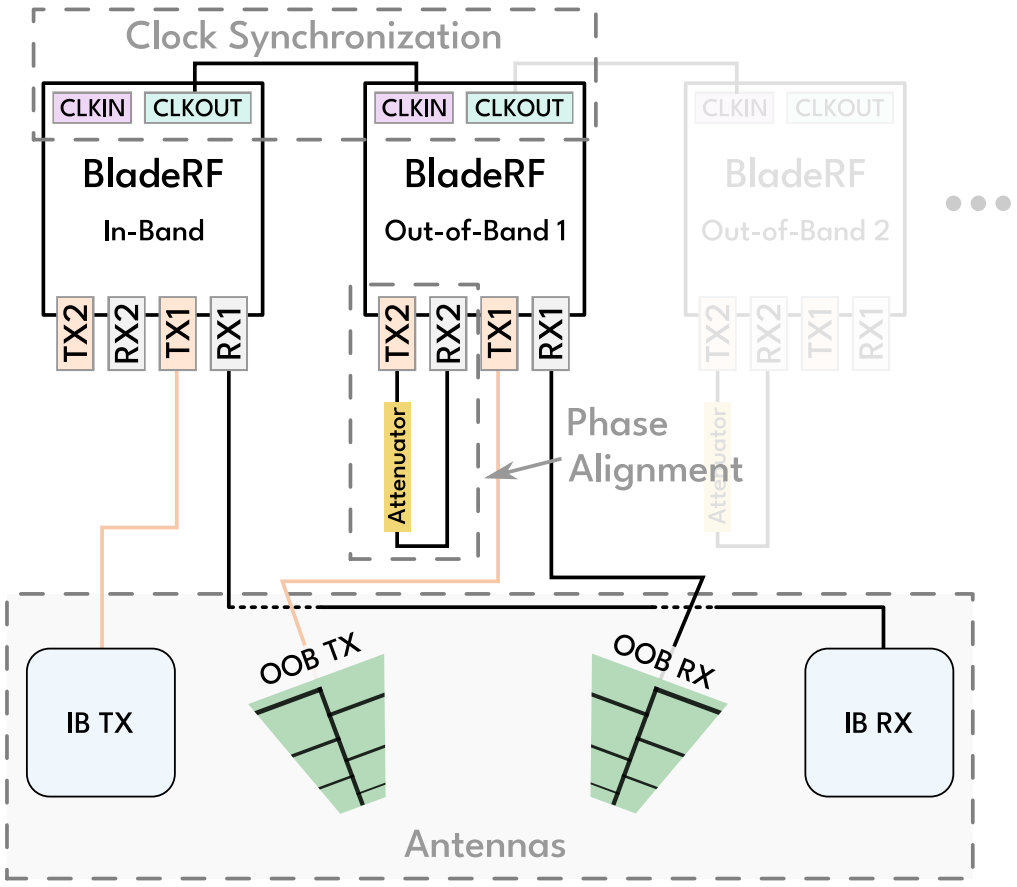

Figure 3-3: System Diagram. The overall system architecture contains multiple bladeRFs for powering up RFIDs and sending synchronized out-of-band signals.

3.3 Frequency Synchronization across BladeRFs

The wideband channel estimation techniques needed for localization assume that all the OOB signals are frequency-synchronized, for this is important for accurate channel estimation. Specifically, we want the local oscillator on the RF frontend of each bladeRF device to share the same reference signal. If reference signals are not shared, receiving a signal transmitted by one bladeRF on a different bladeRF could

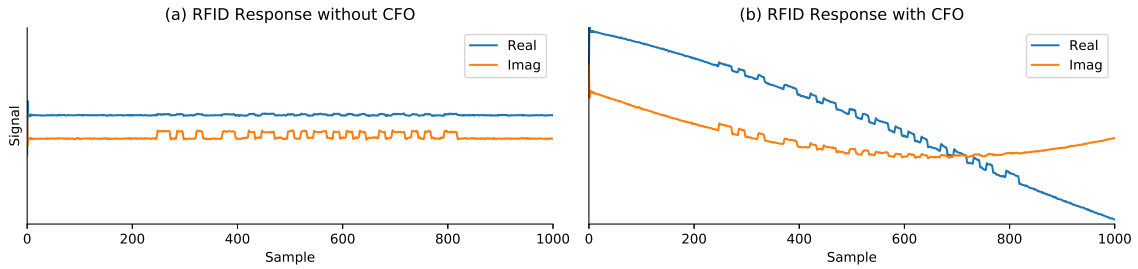

Figure 3-4: Carrier Frequency Offset. The plots show the same RFID backscatter response with and without a carrier frequency offset (CFO). When bladeRFs are not frequency synchronized, CFO makes estimating the wideband channel response more difficult and increases errors.

result in something called carrier-frequency offset (CFO). Fig. 3-4 shows the same complex RFID signal with and without a carrier frequency offset. As apparent in Fig. 3-4(b), CFO makes it difficult to measure the difference between the reflective and non-reflective samples. CFO could be corrected in post-processing, but doing so will still introduce new error into the system.

Luckily, the bladeRFs have a fairly standard solution to this problem. Figure 3-5 shows the clock distribution of the bladeRF across multiple devices. Each bladeRF can send a 38.4 MHz clock signal to the RF frontend via two methods: an internal oscillator signal or an external oscillator signal from CLKIN [25]. Additionally, a bladeRF can output this same reference signal to the CLKOUT port.

As such, I can set up the bladeRF clocks in a master-slave format. The in-band bladeRF uses its internal oscillator to send a clock signal to both its RF frontend and CLKOUT port. The first OOB bladeRF connects its CLKIN port to the CLKOUT of the master IB bladeRF. Now these two bladeRFs share the exact same clock signal, and thus both RF frontends should tune to same frequencies (within the frontend error specification). Moreover, additional OOB bladeRFs can be daisy-chained together to share the same clock signal, which allows infinite expansion of the architecture for clock synchronization purposes.

There are two alternatives to this approach that also achieve frequency synchronization. First, the bladeRFs support sharing a 10 MHz reference signal to align their internal oscillator. However, this approach requires additional hardware. Sec-

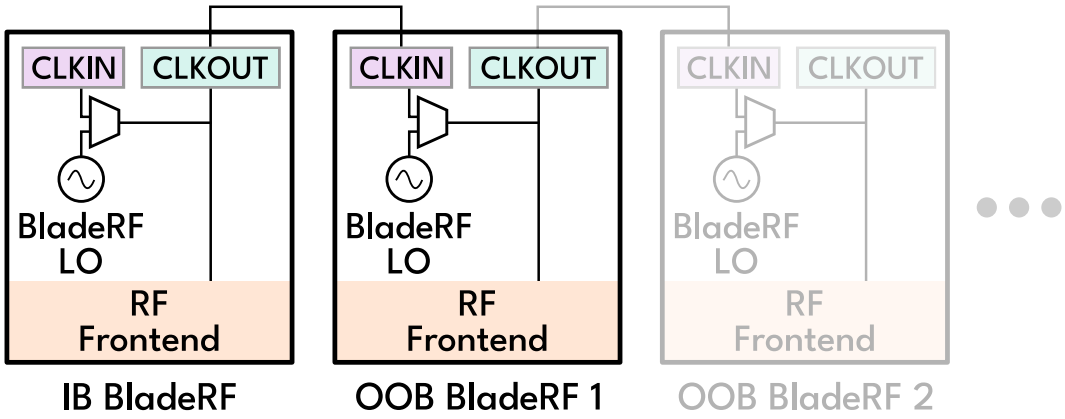

Figure 3-5: BladeRF Clock Synchronization. To ensure accurate frequency tuning across devices, all the bladeRFs must share a clock signal. The in-band bladeRF device serves as the master clock signal and shares its reference signal with the other slave devices.

ond, the RF frontend can directly accept an external LO signal at the desired carrier frequency instead of using the clock from the bladeRF. Again, this approach requires external hardware that is capable of switching between multiple OOB frequencies quickly. Overall, sharing the master bladeRF clock to other devices is simple and robust. Now that the RF frontend frequencies are synchronized, the next section looks at aligning RF signals in time.

3.4 Time Synchronization Across RX Channels

The key to synchronizing the RF frequencies across bladeRFs was sharing the clock frequency. However, this sharing does not give all the bladeRFs the same sense of absolute time. Instead, each bladeRF has an internal counter that starts to increment at the clock rate as soon as it is powered on. As a result, if each bladeRF is scheduled to receive a signal at $t = 100$ms, these signals will not actually be aligned in real-time, which causes two problems. First, the RFID only backscatters signal within a certain time window, so I must ensure that each bladeRF is receiving the OOB signal within that window. Second, the decoding and localization algorithm requires that

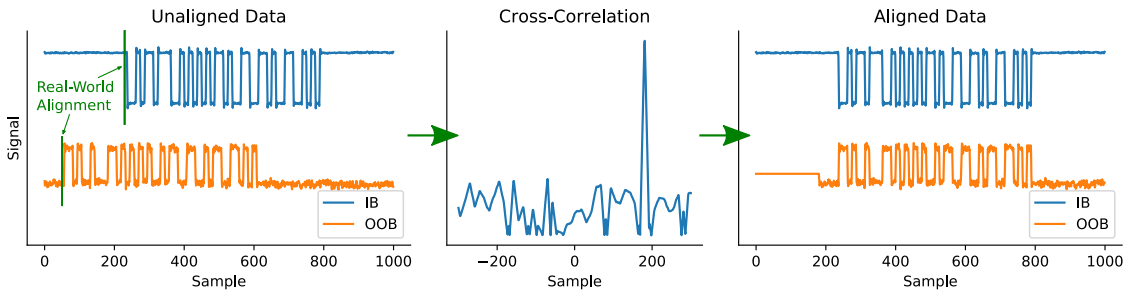

Figure 3-6: Time Alignment Pipeline. The OOB bladeRF started receiving about 150 real-time samples earlier than the IB bladeRF, resulting in out-of-alignment data. The two signals can be correlated and aligned in post-processing to ensure the channel estimates are accurate.

the IB and OOB signals are all perfectly aligned. Doing so allows us to assume that a reflecting sample in the IB data corresponds to a reflecting sample in the OOB data.

The first plot in Figure 3-6 demonstrates these issues occurring across bladeRFs without any additional time synchronization. I overcome these limitations with two steps.

1. First, I roughly align the transmit and receive calls based on Python execution timing. That is, I rapidly check and store the internal clock time for each bladeRF. Then, I schedule the transmit and receive executions for some small delay in the future from these clock times. Thus, each bladeRF executes the RF frontend at about the same time plus or minus any overhead from the bladeRF and Python communication. However, the alignment needs to be within 1-2 samples of data to get accurate channel estimates, which brings us to step two.

2. After each bladeRF has received all the necessary data, I align the data in post processing with a cross-correlation via FFT. Since all the data should represent the same real-world signal, a cross-correlation works robustly even in high-noise signals. Note that step 1 has sufficient accuracy for aligning the transmitted signal data, so step 2 is only for processing received data.

Like the frequency synchronization, this approach to time synchronization can be easily scaled to additional OOB bladeRF devices by just taking additional cross-correlations with the IB signal. I will note that there is a hardware-based way to

trigger multiple bladeRFs to execute simultaneously. This approach should be explored by future bladeRF researchers for speed and robustness, but it was not in scope of this thesis work.

3.5 Phase Synchronization

Phase comprises the last remaining synchronization. Figure 3-7 shows the a high-level layout of the RF frontend on the bladeRF. Notice that although the RF frontends of all devices receive the same reference clock signal because of our frequency synchronization step (Figure 3-5), the TX and RX channels each operate on different local oscillators within the RF frontend. The TX and RX channels can tune to the exact same frequency based on the reference signal, but there will be a random phase offset between them due to the phase-lock-loops. Phase-lock-loops are devices that can generate a wide range of new oscillator frequencies from a fixed oscillator frequency (in this case being the reference clock), which allows the OOB bladeRF to send data across a wide band of frequencies. Inherent to their design, however, is the need to "lock" to the reference signal, which happens at a random phase offset each time.

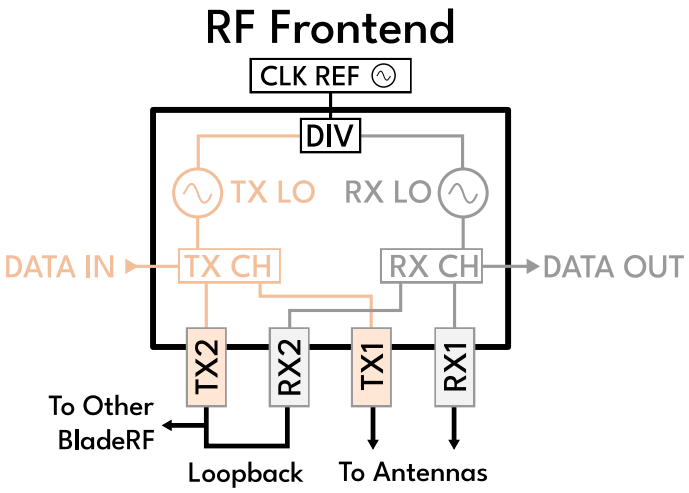

Figure 3-7: RF Frontend Oscillator Layout. The TX and RX channels each operate on a different oscillators which contain phase-lock loops. This layout means that although the frequency is the same between the transmit and receive channels, the phase of the signal is not and leads to inaccuracies in localization.

A random phase offset impacts localization because phase has a direct correlation with a distance measurement via the relationship: $d = \lambda\frac{\theta}{2\pi}$. To overcome this issue, I take advantage of the dual RX and TX outputs on each bladeRF. Both TX outputs operate via the same PLL and thus have the same random phase offset (and likewise for the RX ports). I use the extra ports to measure this offset while still transmitting and receiving RFID signals on the other two ports. In Figure 3-7, the loopback placed on the TX2 and RX2 ports is used to measure the phase offset between the PLLs so I can correct for it in post-processing.

Unlike the previous two synchronizations, phase is hard to understand in the time-domain. Instead, I can verify that the phase of the wideband channel estimate is consistent across trials. Figure 3-8 shows the difference in phase of the wideband channel estimate across two identical trials with and without the loopback correction. The random phase offset is apparent without the correction, as each time the frequency changes, a new random phase offset occurs.

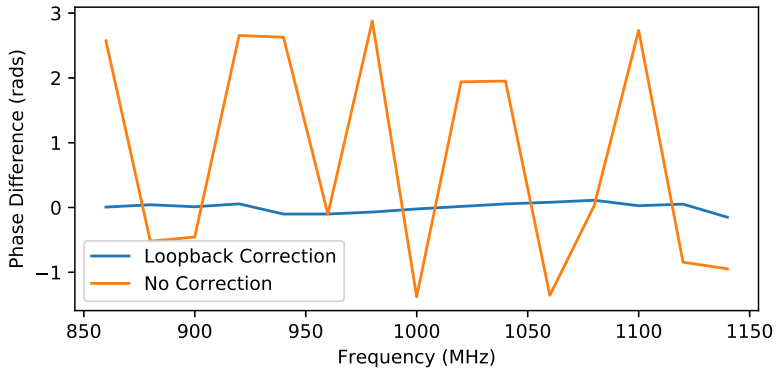

Figure 3-8: Phase Synchronization. The difference in the phase of the wideband channel estimate across two identical trials is corrected using a loopback. Without the loopback, the random phase offset impacts all the channel estimates.

I can also verify this approach mathematically. Let's denote the phase of the signal I want to transmit as θ_i, but due to a random phase offset of θ_{rng1} in the TX channel PLL, the actual phase of the signal at the TX port is $\theta_{TX} = \theta_i + \theta_{rng1}$. Similarly, I receive a signal at the RX port with a phase of θ_{RX}, but this has an offset from the

RX channel PLL as well before I process it outside the RF frontend.

$$\theta_{TX} = \theta_i + \theta_{rng1} \qquad \theta_f = \theta_{RX} + \theta_{rng2}$$

. Since the loopback is a short direct path from TX to RX, I know that $\theta_{TX} \approx \theta_{RX}$. However, I can only actually measure the phase difference between the signal we send to the TX channel and the signal we receive from the RX channel: $\theta_f - \theta_i$.

$$\begin{aligned} \theta_f - \theta_i &= (\theta_{RX} + \theta_{rng2}) - (\theta_{TX} - \theta_{rng1}) \\ \theta_f - \theta_i &= \theta_{TX} + \theta_{rng2} - \theta_{TX} + \theta_{rng1} \\ \theta_f - \theta_i &= \theta_{rng2} + \theta_{rng1} \end{aligned} \tag{3.1}$$

This procedure thus measures the random phase offset between the RX and TX PLLs. I can phase-shift the RX data from the other port by this amount to arrive at a corrected signal, which is the same procedure that achieves the correction in Figure 3-8. In addition, if I have additional bladeRF devices to phase synchronize, I can share the TX loopback signal with the additional RX ports via a splitter.

3.6 Dealing with Outliers

In theory, each OOB bladeRF provides a time-of-flight distance measurement to the RFID tag. Given the locations of the bladeRFs themselves, a simple trilateration of the RFID tag in 3D space is straightforward. However, in practice, RF signals suffer from noise and interference, which can lead to outlier measurements that may result in poor localization accuracy. For example, if one of bladeRFs returns a distance measurement off by 10s of centimeters, the trilateration would be inaccurate or fail to converge on a solution. Thus, I also implemented a robust least squares approach to localization using the RANSAC algorithm. RANSAC picks the most likely set of accurate measurements while ignoring any outliers and increases the robustness of the system.

3.7 A Software-Hardware Implementation

In the previous sections, I have described how to overcome the challenges of synchronization in time, frequency, and phase on bladeRF SDRs. In this section, I describe my software-hardware implementation and evaluation of the low-cost platform.

3.7.1 A Fully-Integrated Platform

Software: The software for the system is a combination of Python and C++ code which comprises mostly of new custom code for the bladeRFs to interface with the previously developed channel estimation and localization algorithms. Nuand develops a libbladeRF API for the bladeRF, which is used extensively in both languages. Python is primarily used for scripting, while C++ is used for any heavy processing tasks.

To get a wideband channel estimate, the OOB bladeRFs must change between multiple frequencies. The standard approach on bladeRFs is to have the RF frontend tune to a new frequency on each change. However, the tuning process is time consuming because it requires a feedback control loop to determine the right circuitry settings to achieve this frequency. No data can be sent or received during this time, and this tuning time can otherwise double the time it takes acquire a wideband channel estimate [24]. A slower channel estimate means a lower framerate, less accuracy on moving objects, and under-utilized CPU resources. To overcome these issues, the RF frontend on the bladeRF offers a feature called quick-tune. Quick-tune stores the tuning settings for a given frequency ahead of time and loads them when that frequency is requested again. Since the localization algorithm can take channel measurements at pre-determined frequencies, quick-tune decreases the total time spent tuning to new frequencies by an order of magnitude. The main drawback is that there is a limit to the number of quick-tune profiles which can be stored ahead of time, which could render frequency-adaptive channel estimates more difficult in the future.

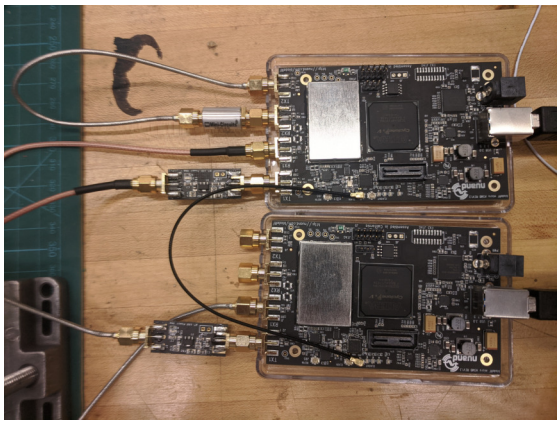

(a) BladeRF Setup

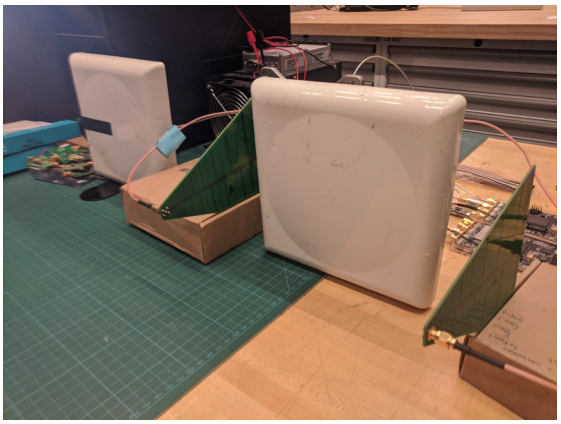

(b) Antenna Setup

Figure 3-9: Real-World Low-Cost Platform Setup. The bladeRF devices were synchronized together and amplifiers were added to the transmit signals. The antenna hardware was placed in multipath-rich environments.

Hardware: For hardware, the bladeRFs contain almost all the components needed for the system, a fact which allows for scalability. Figure 3-9a shows a real bladeRF setup consisting of one IB bladeRF and one OOB bladeRF. U.Fl cables are used to connect the bladeRF clocks, and a 20 dB attenuator on the phase correction loopback prevents overpowering the RX channels. Both the IB and OOB bladeRFs use Nuand BT-100 power amplifiers on the TX port to increase the transmit power of the RFID queries.

For antennas, MTi Wireless Edge MT-242025/TRH/A (865-956 MHz) circularly polarized patch antennas transmit and receive the in-band signals. Most experiments used WA5VJB Log Periodic Antennas (850–6,500 MHz) for out-of-band. Chapters 4 and 5 discuss antennas in more detail. I used the Gen 2 RFID protocol along with a variety of UHF RFID tags across different experiments. The experiments were conducted in several multipath-rich areas of a lab environment to replicate realistic deployment settings.

3.7.2 Results

Here, I focus on localization results in a multipath-rich indoor environment. I used the same parameters as prior implementations [9] and fixed the RFID orientation

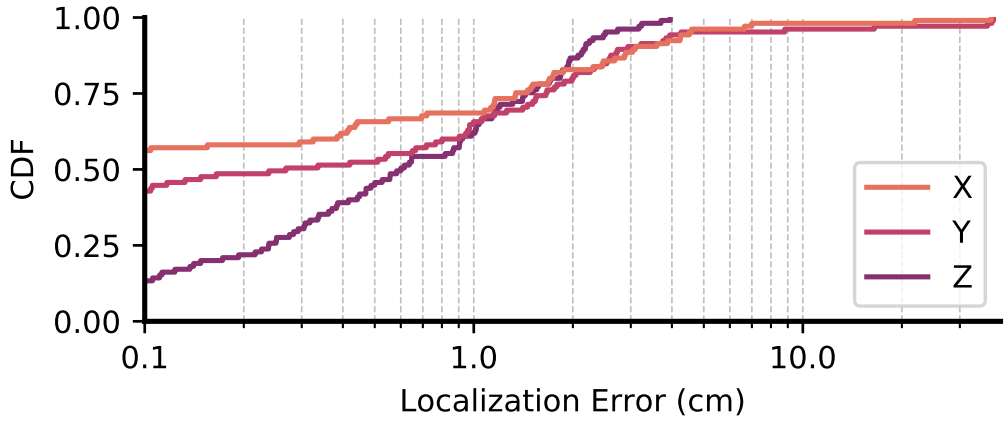

Figure 3-10: Localization Accuracy. The plot shows a CDF of localization accuracy along the X, Y, and Z dimensions.

for fair comparison. In my evaluation, the localization (out-of-band) antennas were moved to at least 3 different locations to obtain time-of-flight measurements, and the measurements were combined to perform 3D positioning.

I ran multiple experiments with the bladeRF system placed in a single part of the lab. I placed an RFID tag in over a dozen locations, both visible and hidden by other objects on the table. Figure 3-10 shows the CDF of the localization error along the x, y, and z dimensions across all of these experiments. The figure shows that the median errors along x, y, and z dimensions are 0.1 cm, 0.26 cm, and 0.6 cm, respectively. The figure also shows that the bladeRF platform achieves 90^{th} percentile errors of 3.07 cm, 2.85 cm, and 2.16 cm along the x, y, and z dimensions, respectively. These results demonstrate that a low-cost, portable RFID localization platform can achieve sub-centimeter level positioning accuracy that matches or exceeds that of state-of-the-art systems [18, 16].

Chapter 4

RFID Orientation and Localization

RFID tag orientation changes pose a challenge to maintaining high-accuracy localization of the RFID tag. Some localization systems have tried to overcome this challenge by using more tags, reference tags, or additional antennas [34, 30, 5]. In this chapter, I discuss the mechanisms by which orientation impacts localization accuracy. In particular, I discuss polarization and phase and conclude with results showing the impacts of tag orientation for different antenna setups.

4.1 Polarization and Phase

First I discuss background on polarization and its impact on RFIDs. Figure 4-1 shows how RF waves propagate with circular polarization and linear polarization. As the name suggests, circular polarization propagates the electromagnetic waves in a spiral pattern, while linear polarization propagates as a sine wave in a single plane. The polarization depends on the design of the antenna, and circular polarization can actually be generated by two linearly polarized antennas placed perpendicularly to each other and offset by 90 degrees in phase.

Polarization is important because relates to how RF waves excite a signal (the electrons) in antennas. Figure 4-2 shows the same linearly polarized signal being received by two linearly polarized antennas that are perpendicular to each other. When an antenna maintains the same polarization as the signal it is receiving, the

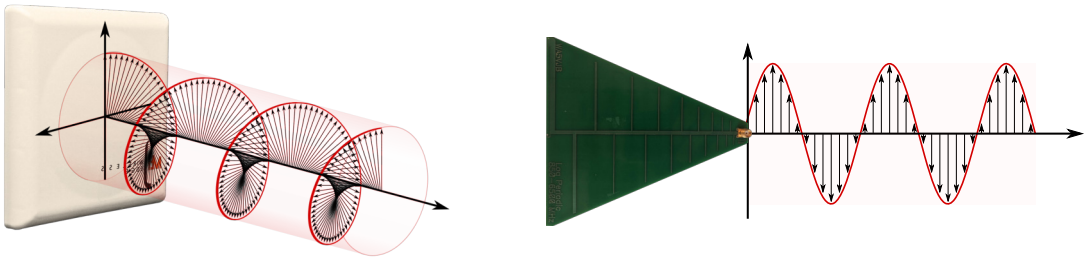

(a) Circular Polarization (b) Linear Polarization

Figure 4-1: Types of Polarization. Polarization defines how the electromagnetic RF waves propagate through space. (a) modified from [10].

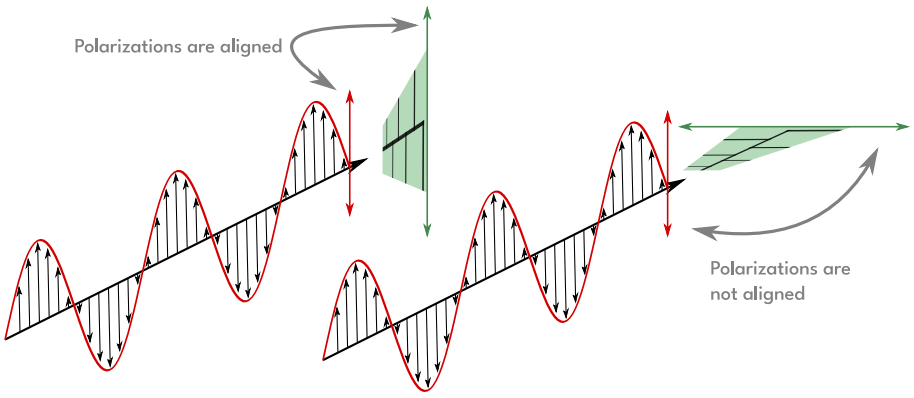

Figure 4-2: Signal-Antenna Polarization Alignment. When the signal polarization is aligned with the antenna polarization, the signal strength is maximized. The signal strength is weaker when polarization does not align.

signal is able to excite a strong current in the antenna. In contrast, when the signal and antenna polarization are perpendicular, the EM waves that form the signal do not resonate the electrons in the antenna in the correct direction, so the signal power captured by the antenna is significantly attenuated. Polarization relates to RFIDs in several ways. First, the antenna of the RFID itself is generally linearly polarized. Thus, if an RFID reader has a linearly polarized antenna, and the RFID is perpendicular to this reader antenna, the reader would be unable to power-up or read the RFID. This problem is why many commercial RFID antennas are circularly polarized as a circularly polarized EM wave can excite a signal in a RFID antenna no matter the orientation. However, simply using circularly polarized antennas for RFID localization is not straightforward because polarization impacts signal phase.

In fact, when sending a circularly polarized signal, the phase of the signal received at an RFID changes based on the orientation of the RFID relative to the signal[1]. For visualization, you can imagine that the change in phase is related to the angle of the polarization spiral in Figure 4-1a when the signal reaches the RFID.

Some prior work has attempted to correct for this phase difference [31]. However, this work required an additional antenna and was not generalized to multipath-rich environments. In this chapter, I experimentally verify this prior work, extend it to multipath environments, and propose a new solution that uses a single linearly polarized antenna to remove the impacts of tag orientation.

4.2 Orientation and Accuracy

Figure 4-3 shows the three main antenna setups of interest when understanding how RFID orientation impacts localization accuracy.

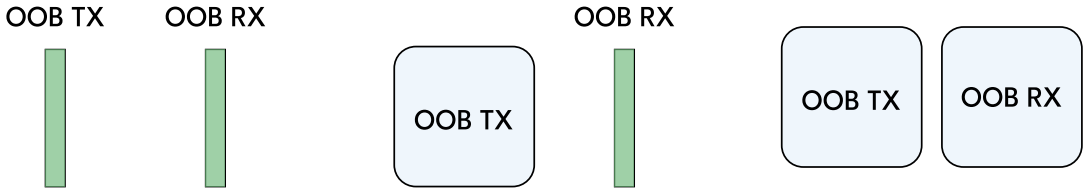

(a) Linear TX and RX (b) Linear TX, Circular RX (c) Circular TX and RX

Figure 4-3: Antenna Setups for Orientation Experiments. The setups consist of a combination of vertically oriented, linearly polarized log-periodic antennas and circularly polarized patch antennas.

For these experiments, I held the RFID at a constant distance from the antennas and rotated it about each axis, as shown in Figure 4-4, which allowed me to measure how the change in orientation affects the received data. The experiments were done in a multipath-poor environment to limit extraneous factors in the channel estimates. Before describing the results, I formulate expectations for each scenario based on the work of [31].

[1]Refer to [31] for a in-depth analysis of why this is the case.

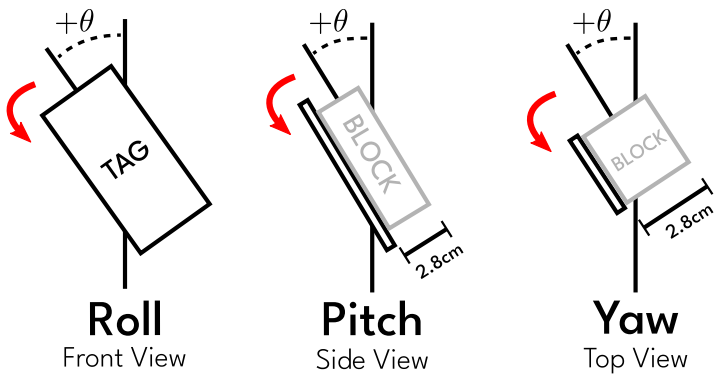

Figure 4-4: Tag Orientation Changes. That RFID tag was rotated around each axis to determine the impact of orientation on channel phase.

Assume that all antennas are oriented vertically, which can be any direction for circularly polarized antennas. Also assume that we use an RFID with a linearly polarized antenna. Define the orientation of the RFID relative to the antennas as θ. From [31], I know that the circularly polarized antenna should introduce a phase offset that is a function of the tag orientation $\phi_0(\theta)$. Thus, for an RFID tag placed at a fixed roundtrip distance away from the antennas, I expect the total change in phase ϕ to be

$$\Delta\phi = \phi_0(\theta) \tag{4.1}$$

For ideal circularly and linearly polarized antennas, [31] formulates $\phi_0(\theta)$, and the expectations for each antenna setup are shown in the roll column of Table 4.1. Additionally, neither pitch nor yaw affects the relative angle between the signal and RFID polarization for these particular experiments.

Antennas	Expected $\Delta\phi$		
	Roll	Pitch	Yaw
Linear TX, Linear RX	0	0	0
Circular TX, Linear RX	$\boldsymbol{\theta}$	0	0
Circular TX, Circular RX	$\boldsymbol{2\theta}$	0	0

Table 4.1: Expected Change in Phase with Orientation

There are few caveats to these expectations. First, whenever the relative angle

between polarization is nearly perpendicular, I expect no response from the RFID, such that any phase calculated from this channel estimate will be meaningless since it is just the phase of noise. Additionally, as indicated in Figure 4-4 for the pitch and yaw experiments, the rotation axis is not exactly at the center of the tag. In fact, there is a $2.8 \times 2 = 5.6$cm increase in the roundtrip distance between the tag and the antennas when the tag is rotated between 0 and 180 degrees in these two directions. Thus, rather than zero phase change, I expect a phase change up to $2\pi\frac{0.056}{\lambda} = \frac{0.35}{\lambda}$. For frequencies 700 MHz - 1200 MHz, this is in the range of 0.82 to 1.406 radians.

4.3 Results

I first ran the orientation experiments with the linearly polarized TX and RX antennas. Figure 4-5 shows the change in channel phase when adjusting the roll, pitch, and yaw orientation of the tag. Each data point with error bars represents the median, minimum, and maximum change in channel phase across all frequencies in the wideband channel estimate. The red regions highlight orientations where there was a large discrepancy in phase across the different frequencies, which suggests that the tag did not power up, and the channel estimates are meaningless at these points. The dashed grey line shows the expected change in phase as I calculated in the previous section, taking into account any distance changes for the pitch and yaw experiments.

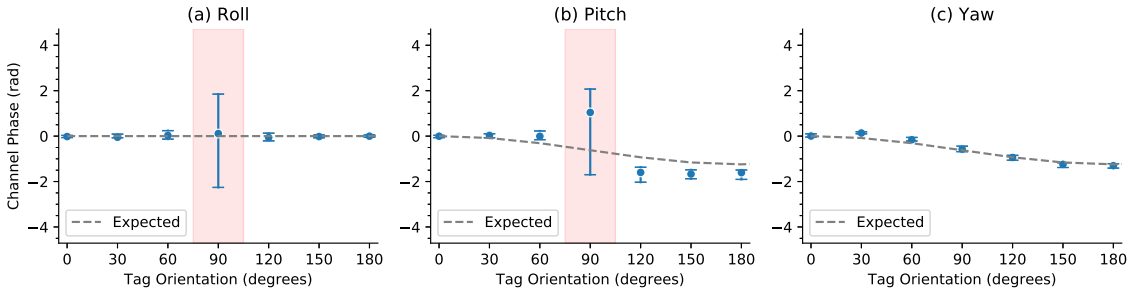

Figure 4-5: Channel Phase and Orientation for Linear TX, Linear RX. The plots show the change in channel phase as a function of tag orientation. The roll orientation has no impact on channel phase.

I make the following remarks. First, when the RFID tag is perpendicular to the

antenna polarization (for roll) or faced flat downward (for pitch), the tag fails to power up and respond. This result is expected since the effective antenna area or polarization alignment is almost zero, so the RFID is unable to absorb much RF energy. I would have expected the same issue in the yaw experiment, but the experimental setup was not perfectly perpendicular in this case. Additionally, the expected changes in phase match the measured changes in phase well for the roll and yaw experiments. The pitch experiments show the same general trend, but with some more error. Most importantly, the channel phase does not change at all for roll orientation changes of the RFID tag. This indicates that, assuming the tag can power up, orientation will have no impact on localization accuracy.

Next, I show the results for circularly polarized TX and linearly polarized RX in Figure 4-6. With a circularly polarized TX antenna in the setup, I expect to see a change in channel phase when the RFID's roll orientation changes. For the first 90 degrees, the results do in fact match the expectation well. However, after 90 degrees the channel phase jumps by π before continuing with the same expected trend. This is explained by the fact that the tag itself it symmetrical, such that the antenna phase should wrap around every 180 degrees instead of every 360 degrees. For the pitch and yaw experiments, the results are similar to the previous antenna setup. The red regions also concur with expectations, as the effective antenna area is zero when the RFID is at 90 degrees of pitch rotation.

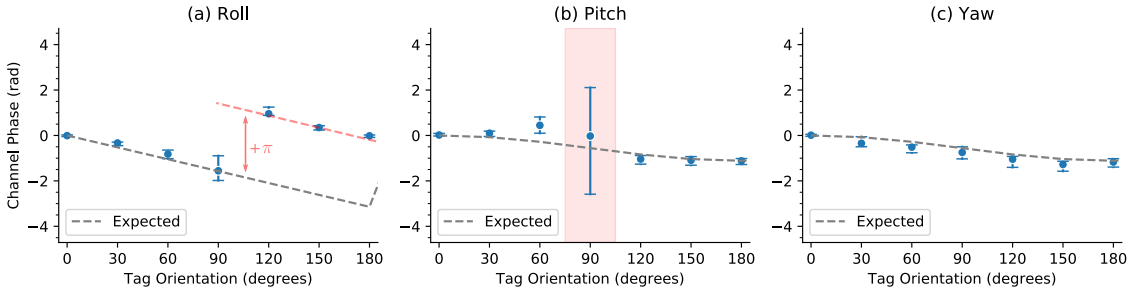

Figure 4-6: Channel Phase and Orientation for Circular TX, Linear RX. The circular antenna introduces a θ shift in channel phase for (a), and the symmetry of the tag cause the phase to jump by π at 90 degrees.

Lastly, Figure 4-7 shows the same set of results for the circularly polarized TX

and RX experiment. For the roll experiment, I note the jump in phase at 90 degrees is 2π, which is due to the expected 2θ change wrapping around at $-\pi$. Additionally, while the yaw experiment again matches my expectations, there is a discrepancy with the pitch experiment. In particular, the tag does not respond at 60 degrees of pitch orientation, and I would expect the tag to power-up as it did in the previous two antenna setups. This could be a combination of destructive interference and a smaller effective antenna on the RFID at this orientation.

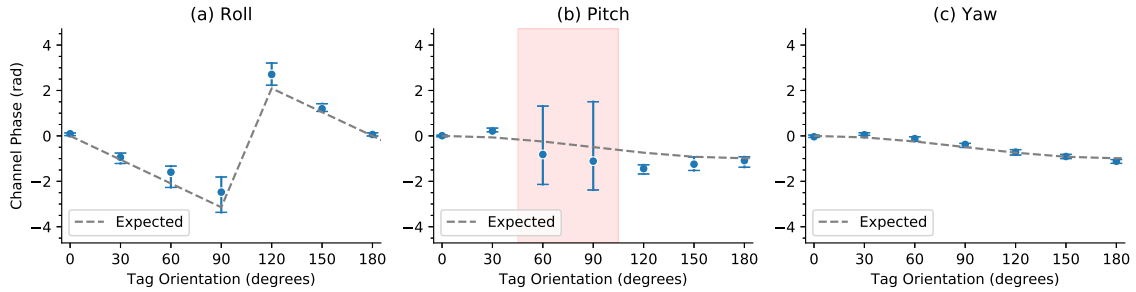

Figure 4-7: Channel Phase and Orientation for Circular TX and RX. The roll orientation results match the expected 2θ change in phase.

In summary, I show that the experimental results generally match the expected changes in phase due to tag orientation. Additionally, I find that a pair of lineally polarized transmit antennas is able to overcome the impacts of tag orientation on channel phase. Thus, linearly polarized antennas can be used to localize the RFID tag accurately despite orientation changes. The only issue is that when an RFID tag is near perpendicular to the antennas, no signal will be received. This issue can easily be detected during signal processing, and either adding additional antennas in stationary setups or moving the antenna in mobile setups would overcome this issue.

Chapter 5

System Characterization and Robustness

In Chapter 3, I showed how the low-cost platform is able to achieve sub-centimeter accuracy in a multipath-rich lab environment. However, these experiments were at a relatively close range, using a single RFID tag, and operating within a specific frequency band. As I continue to push the performance of the bladeRF platform, it is important to understand the limits of bladeRF sensitivity, bladeRF transmit power, and RFID signal strength across a larger frequency spectrum. As mentioned multiple times, RFID backscatter is (almost) frequency-agnostic, and utilizing a wider range of frequencies for localization allows the system to avoid interference with other signals, comply with regulations, and improve overall accuracy. In this chapter, I discuss how I characterize the frequency response of bladeRFs, antennas, and RFID tags. For the following experiments, a signal generator (Keysight MXG N5183B) and a spectrum analyzer (Keysight MXA N9020B) are used to characterize the system components accurately.

5.1 BladeRF TX Performance

First, I will show how I measured the frequency response of the bladeRF transmit channel. Figure 5-1 shows the experimental setup. One of the TX ports on the OOB

bladeRF is connected via a coaxial cable to the spectrum analyzer. The bladeRF was programmed to transmit a constant signal for 100 ms, and sweep across frequencies 200 MHz to 6000 MHz every 20 MHz. The digital sampling rate was set to 1 MSps. Additionally, the internal gain settings for the bladeRF were adjusted every 10 dB from -14 dB to 66 dB. For reference, a bladeRF TX gain setting of 60 dB is roughly equal to 0 dBm of signal power. The spectrum analyzer was set to max hold and a frequency range of 200-6000 MHz as well, sweeping in increments of 20 MHz. The resolution bandwidth (RBW) was set to 1.05 MHz, and a -6dB Gaussian detector was used. The amplitude reference level was set to 20 dBm. The experiment was conducted with and without a Nuand BT-100 amplifier attached to the TX port of the bladeRF.

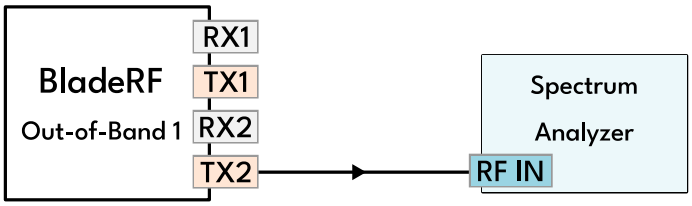

Figure 5-1: TX Characterization Setup. The transmit power of the bladeRF was measured using a spectrum analyzer.

Figure 5-2 shows the signal power received by the spectrum analyzer across the frequency range. Each line represents a different TX gain setting on the bladeRF. In both plots, it can be noted that the signal power is generally much weaker at higher frequencies than that at lower frequencies; dropping over 10 dB from 200 MHz to 6000 MHz without the BT-100 amplifier, and dropping almost 30 dB with the BT-100 amplifier. Nonetheless, the BT-100 amplifier adds about 15 dBm of signal power at lower frequencies, and about 5 dBm at higher frequencies. Finally, note that for certain TX gain settings the transmitted power is actually lower than the environmental noise floor. For example, if the TX gain setting were set to 6 dB with the BT-100 amplifier attached, it would be impossible to read any signal from the bladeRF above 2000 MHz. Thus, for certain frequencies and power levels, the bladeRF will not be able to perform RFID localization.

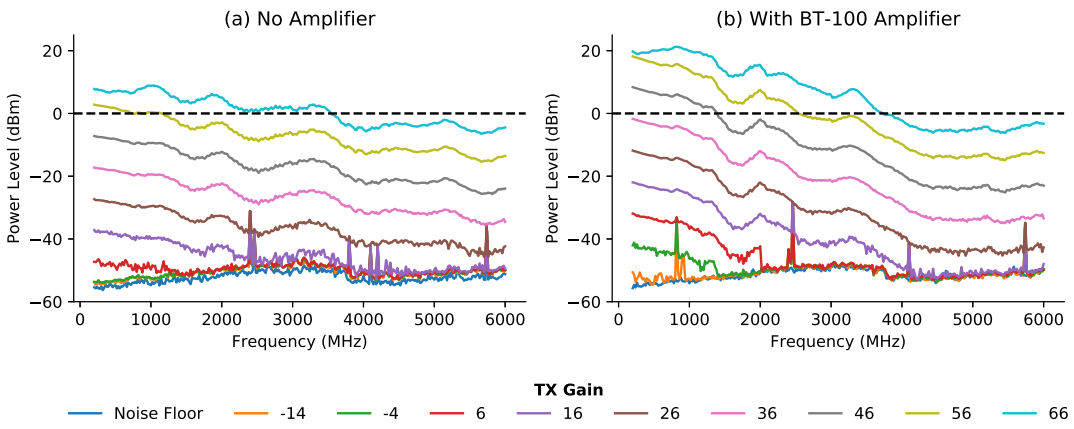

Figure 5-2: BladeRF TX Power. The plots show the power output of the bladeRF at different TX gain settings.

5.2 BladeRF RX Performance

Next, I characterize the receive sensitivity of the bladeRF. Figure 5-3 shows the experimental setup, where the signal generator transmits a signal of known power to an RX port on the bladeRF via a coaxial cable. The bladeRF settings were the same as those in the TX experiment, except that only 2ms of samples were received at each frequency and the bladeRF RX gain settings were set at 0, 20 and 40 dB. On the bladeRF, the gain setting adjusts the internal LNA in the RF frontend. The signal generator was manually adjusted to a every frequency from 200 - 6000 MHz at 20 MHz intervals. The signal generator power was also changed between -20, -15, and -10 dBm.

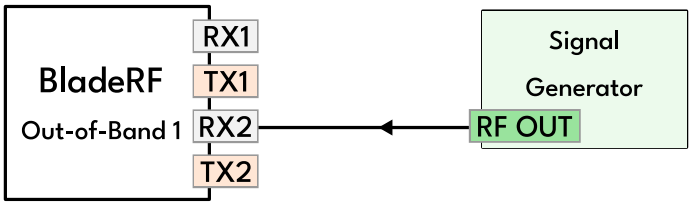

Figure 5-3: RX Characterization Setup. A signal generator transmits a known-power signal to the bladeRF to measure the sensitivity of the RX channel.

Figure 5-4a shows the frequency response of the bladeRF for different signal gen-

erator power levels at a bladeRF RX gain setting of 0 dB. The received signal power is calculated by taking the average power of all samples $x[i]$ in that frequency band, as shown in Eq. 5.1.

$$\text{Power} = \frac{1}{N}\sum_{i}^{N} x[i]x^*[i] \tag{5.1}$$

Additionally, the bladeRF ADC generates complex integers between -2048 and 2048. Thus, the maximum power before saturation is $10\log\left(2 \cdot 2048^2\right) = 69.2dB$.

As expected, the received power increases by 5 dB for every 5 dBm increase in signal generator power. However, there is still a large decrease in received power at higher frequencies compared to the lower frequencies. In the -20 dBm signal generator power trial, the received power drops from above 60 dB at 200 MHz to below 45 dB at 6000 MHz.

Next, Figure 5-4b shows the frequency response across different bladeRF RX gain settings for a fixed signal generator power of -20 dBm. The plot shows that the 20, 40 and 60 dB gain settings cause ADC saturation for most of the spectrum.

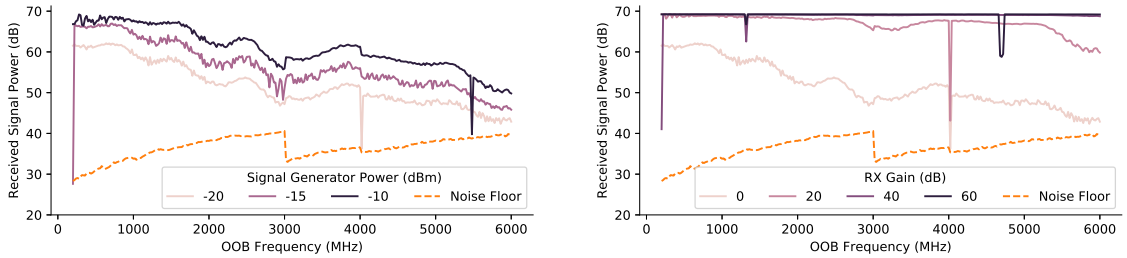

(a) RX Power vs. Signal Generator Power

(b) RX Power vs. RX Gain

Figure 5-4: BladeRF RX Power. The plots show that received signal power decreases at higher frequencies.

Now, I look at the noise floor in more detail to determine the minimum power signal the bladeRF can receive. Figure 5-5 shows the received power of signals near the noise floor with the bladeRF RX gain set to 0 dB. Below 1000 Mhz, the bladeRF is able to receive all the tested signals down to -40 dBm. The noise floor is still 10 dB below the weakest signal, suggesting a sensitvity down to even -50 dBm as well. From 3000-6000 MHz, the bladeRF received power decreases by around 5 dB, and the noise floor also increases by a similar amount. Thus, at 6000 MHz the bladeRF

is only able to receive signals at least as strong as -24dBm. Overall, this section shows that the bladeRF devices are not very sensitive to weak signals, especially at higher frequencies, and thus the bladeRFs may struggle to localize RFIDs using high frequencies or at far ranges where the signal has strongly attenutated.

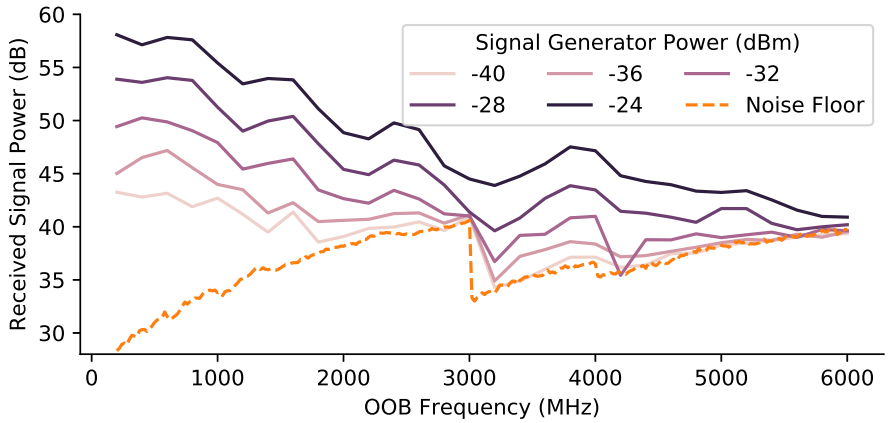

Figure 5-5: RX Sensitivity. The plots show the received signal power for weak signals near the noise floor. At higher frequencies, signals below -8 dBm can not be read by the bladeRF.

Another thing to note is that the noise floor jumps around 3000 MHz. This is because the bladeRF uses the Analog Devices AD9361 chip which can span a wide frequency, but has a different transceiver chain for the 70-3000 MHz range than for the 3000-6000 MHz range.

5.3 Antenna Performance

Besides the bladeRF device itself, the specific antennas used during localization also impact the system performance. Antennas are optimized for different frequencies, and antennas also have different gain settings which impact the amount of power pointed in a certain direction. In this section, I characterize the frequency response of three different antennas that can be used in the system. These antennas are shown in Figure 5-6 along with the experimental setup. For these experiments, the spectrum analyzer was given the same settings as in Section 5.1, except that the amplitude reference level was 0 dBm. The signal generator was set to sweep between 200-6000

MHz with a 20 MHz step size and a 100 ms dwell time at each interval. The power level was set to 0 dBm. The front tips of each pair of antennas were placed facing each other 69.5 cm apart for all experiments.

(a) MTi Wireless Edge MT-242025/TRH/A

(b) WA5JVB Log Periodic 850-6500 MHz

(c) WA5JVB Log Periodic 2-11 GHz

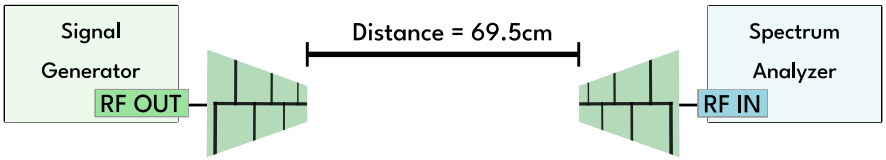

(d) Antenna Characterization Setup

Figure 5-6: Antennas Characterization. Each antenna has a different frequency response that impacts the localization accuracy of the system.

Figure 5-7 shows the signal power received after transmitting through the transmit and receive antennas over-the-air. The dashed black line at 0 dBm indicates the transmitted signal power. The plot also includes for reference the isotropic free-space path loss that is expected based on the antenna separation.

A few points are worth noting:

- The free-space path-loss has a downward trend, which is expected since the loss is inversely proportional to frequency. In addition, the path-loss does not account for the antenna gain which explains the discrepancy between the plot and the data obtained with the antennas.

- Both log periodic antennas produce a stronger response at higher frequencies than the circular patch antenna. Outside of the 865-956 MHz band that the patch antenna is optimized for, the patch antenna rapidly drops off in signal strength. Additionally, all antennas show a similar trend to the free-space path

loss based on frequency, but I don't expect the absolute level of the path loss to match the antenna response as all the antennas have different directional gains.

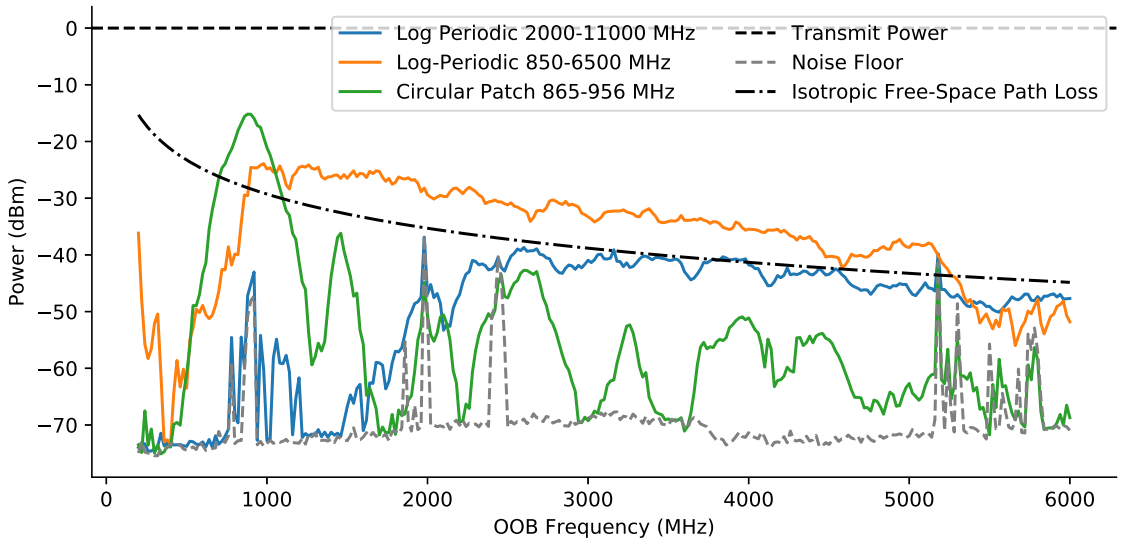

Figure 5-7: Antenna Characterization. The log periodic antennas show a much more consistent wideband response, as is expected, while the patch antenna is optimized for the RFID ISM band.

5.4 Amplifier Performance

As discovered in Sections 5.1- 5.3, all the components of the localization system have weaker responses at higher frequencies. Receive gain settings can only improve results to a certain extent, as higher gain settings correspond with a higher noise floor as well. Thus, additional amplification on the transmit side may be needed for certain deployment scenarios at higher frequencies. In this section, I characterize the amplification power of a Mini Circuits ZHL-4W-422+ power amplifier. Figure 5-8 shows the experimental setup, where an amplifier is connected between the signal generator and spectrum analyzer. The signal generator is power output set to -20 dBm.

Figure 5-9 plots the increase in power level that the amplifier produces at different frequencies. The amplifier increases the signal power by over 20 dB on average until

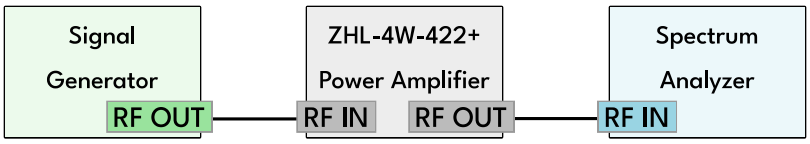

Figure 5-8: Power Amplifier Characterization Setup. The power amplifier boosts the signal from the signal generator, and the amplification level is measured by the spectrum analyzer.

4500 MHz, after which the power drops off rapidly. With such a strong amplifier, the bladeRF should be able to receive an RFID signal at higher frequencies.

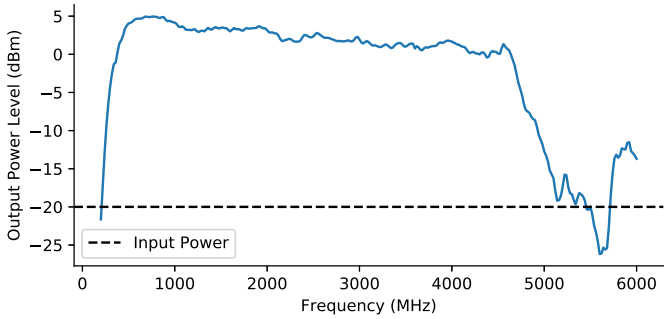

Figure 5-9: ZHL Amplifier Gain. The power amplifier can boost the transmitted signal by over 20 dB up to 4500 MHz.

5.5 One-Way Verification

Now that the bladeRF, antennas, and amplifiers have all been characterized, I will verify the characterizations by combining these components into an end-to-end system, as shown in Figure 5-10. The bladeRF TX port transmits a signal to the 850-6500 MHz log-periodic antenna, and the bladeRF RX port receives that signal from another log-periodic antenna placed 69.5 cm away from the first. The bladeRF settings are the same as those in Sections 5.1 and 5.2. The goal of this section is to show that the true received signal power matches what I expect based on the characterizations from the previous sections.

Figure 5-11a shows the received signal power on the bladeRF. The actual signal power (calculated from the received data) is in blue, the expected signal power (cal-

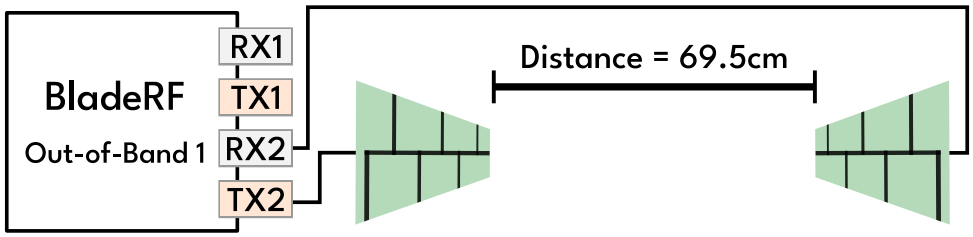

Figure 5-10: One Way Characterization Setup. The system is set-up end-to-end, such that the characterizations of TX, antennas, and RX can be verified in a closed loop.

culated from transmitted data) is the dashed black line, and the receive channel noise floor is in yellow. I make the following remarks:

- The expected power matches the real power from 200 MHz to 3500 MHz. This means that the previous characterizations were done correctly, and this also indicates that I can now calculate the inverse relationship as well (i.e. determine transmitted power from received power).
- Above 3500 MHz, the expectation diverges from the actual. Instead, the actual power matches the power of the noise floor at those frequencies. This indicates that the signal power itself at these frequencies is below the noise floor, so we are unable to measure that transmitted signal there.
- At around 1000 MHz, the actual received power appears to flatten briefly. This is due to the received signal being very strong at this frequency, and thus it saturates the bladeRF ADCs.

Figure 5-11b shows the power received at the RX port for the same experiment, which differs from the power of the processed signal as plotted in Figure 5-11a. Figure 5-11b was calculated by taking the inverse of the RX characterization (Figure 5-4b), such that the processed signal power is mapped back into input power in dBm. Power in dBm can be compared across devices more easily.

Given these results, the likely next question is how can the bladeRF read signals above 3500 MHz? One solution is to increase the RX gain settings. The results for RX gain settings 10, 20, 30 and 40 dB are shown in Figure 5-12. As the RX gain

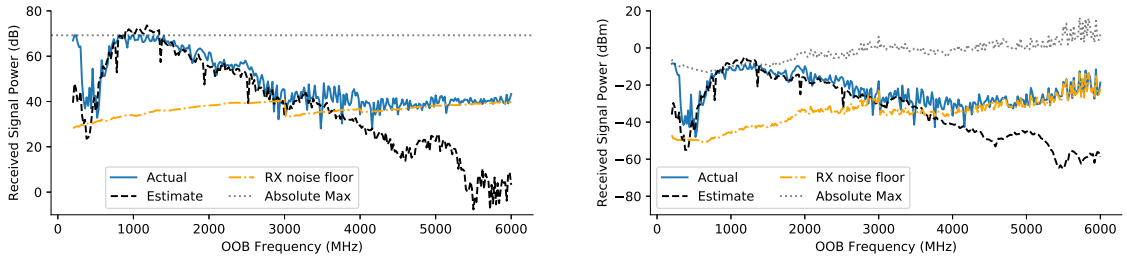

(a) Signal Power (dB)

(b) Signal Power at RX Port (dBm)

Figure 5-11: One Way Verification of Characterizations. The plots show the actual signal power received compared to the estimated power received based on the transmit power and path loss.

increases, the signal becomes readable at higher frequencies (e.g., up to 4000 MHz with a RX gain of 20 dB). However, the maximum input power to the system before saturation (indicated by the dotted grey line) decreases with an increase in RX gain settings, reducing the dynamic range of the system. This is expected since a large RX gain will boost strong signals beyond the saturation value. Thus, increasing the transmit power instead may be a better solution.

Figure 5-13 shows the received signal power after adding the ZHL power amplifier to the TX channel. Even with keeping the RX gain setting at just 0 dB, the bladeRF can read the signal all the until 5000 MHz before dropping below the noise floor again. This is a clear improvement, but it does introduce new hardware and increased power requirements.

5.6 UWB RFID Backscatter

By now I have characterized all the components of the system, verified the characterizations, and introduced amplification to improve the response at higher frequencies. The last part of the system is the RFID itself. Although RFIDs are optimized for the roughly 900-950 MHz ISM band, I have previously discussed how backscatter is frequency agnostic. In this section, I explore how RFIDs respond at different frequencies, and discuss how this impacts localization.

Figure 5-14 shows the setup for these experiments. The RFID is placed at a

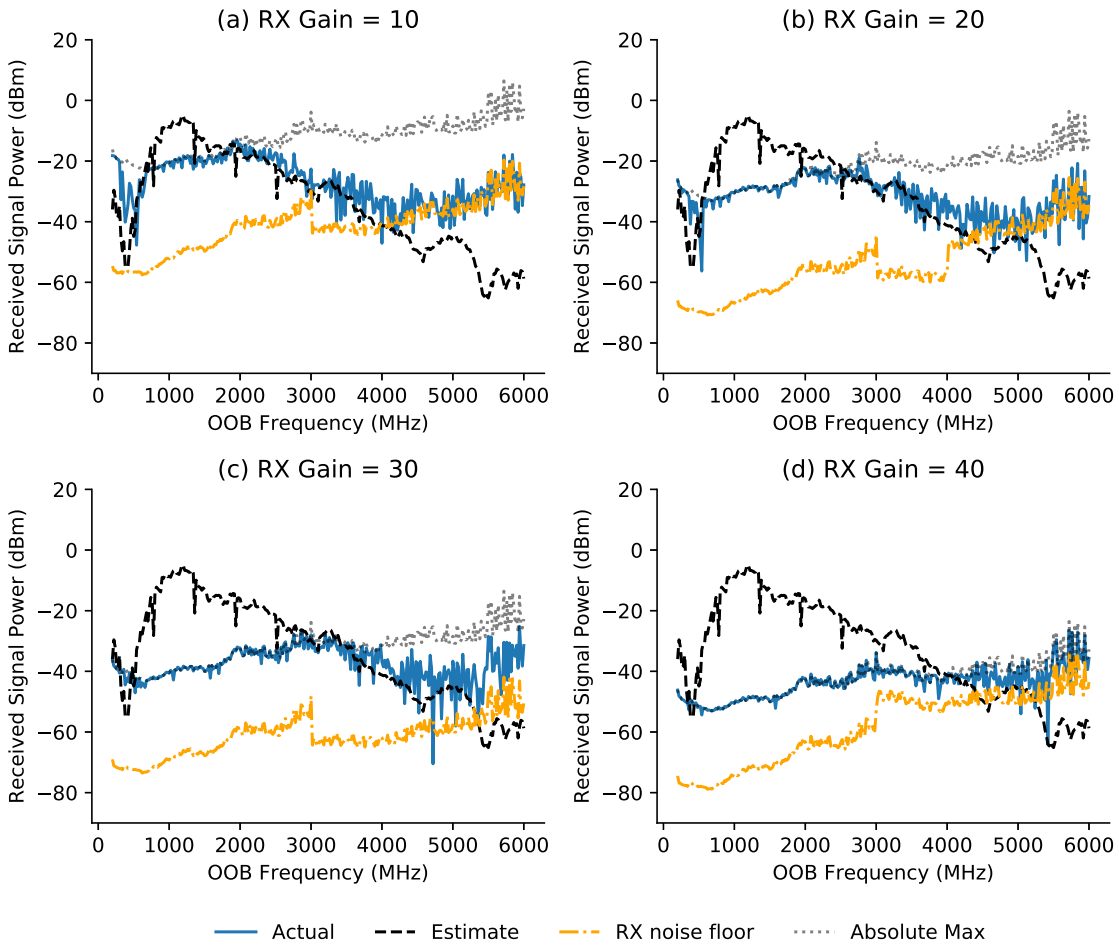

Figure 5-12: Adjusting RX Gain. Increasing the RX can boost the signal strenght above the noise floor at higher frequencies. However, it also increases the noise floor itself which reduces the system's dynamic range.

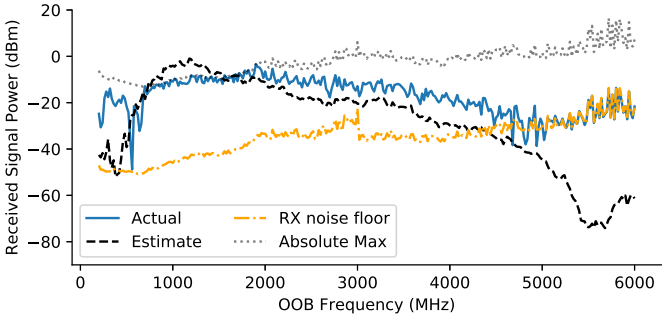

Figure 5-13: ZHL Amplifier Improving Response. The ZHL power amplifier enables the bladeRFs to read signals at higher frequencies without increasing the RX noise floor.

location such that the roundtrip distance from the TX antenna to the RFID and back to the RX antenna is 69.5 cm. The log-periodic 850-6500 MHz antennas are pointed at the RFID, and the IB bladeRF (not pictured) is used to power up the RFID tag. The experiments were conducted by powering up the RFID and transmitting an OOB signal 10 times every 20 MHz from 200-6000 MHz. Finally the data was processed to produce channel estimates at each frequency.

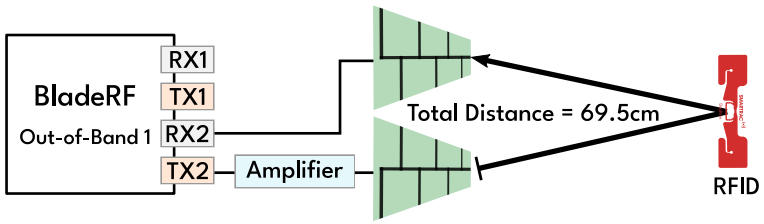

Figure 5-14: RFID Characterization Setup. The OOB bladeRF is connected to the TX and RX antennas, and placed a distance away from the RFID tag. The IB bladeRF is not pictured, but it is nearby and synchronized with the OOB bladeRF.

The first experiment was conducted using the Nuand BT-100 amplifier. Fig 5-15 shows the the signal-to-noise ratio (SNR) of the OOB channel estimates at different frequencies. Each frequency was repeated 10 times, and the blue line represents the median of the SNRs across those 10 trials. Additionally, the green dashed line represents the SNR of the channel determined by all 10 trials together. Specifically, that means estimating the channel from all 10 trials together then calculating the SNR, rather than calculating the SNR then taking the median. If the RFID reponse is readable, I expect the increased length estimate to have a 10dB increase in SNR. In the figure, this is true from 200 MHz - 2000 MHz, which means the RFID tag is backscattering frequencies in that range. Above 2000 MHz, however, it is unknown whether the lack of response is a failure of the RFID to respond, or a failure of the bladeRF to read such a weak signal. Next, I add the Mini Circuits ZHL amplifier to determine this.

Figure 5-16 shows the SNR of the OOB channel estimates when using the ZHL power amplifier. The RFID response is clearly noticeable up until 4000 MHz, after which it is unclear again. However, refering back to Figure 5-13, the bladeRF received

power begins approaching the noise floor after 4000 MHz. Thus, it's likely that the RFID signal is again too weak for the bladeRF to measure at these high frequencies. The failure in RFID response from 800-1500 MHz was unexpected, but this could be due to various reasons (for example, noise from the amplifier interfering with the IB bladeRF signals in this frequency range, or the RFID itself being overpowered with the transmit frequency).

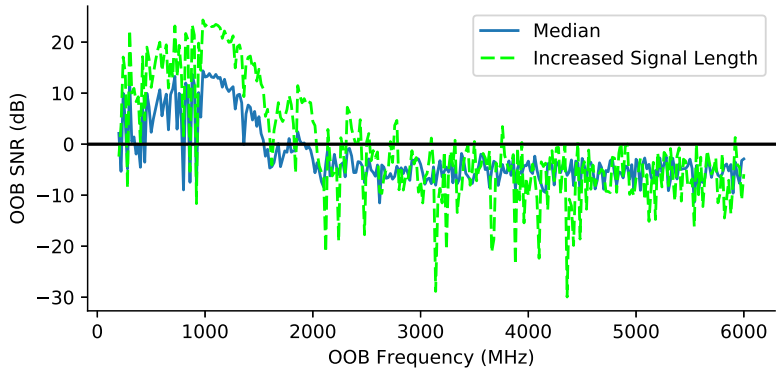

Figure 5-15: BT-100 Amplified BLadeRF RFID Tag Response

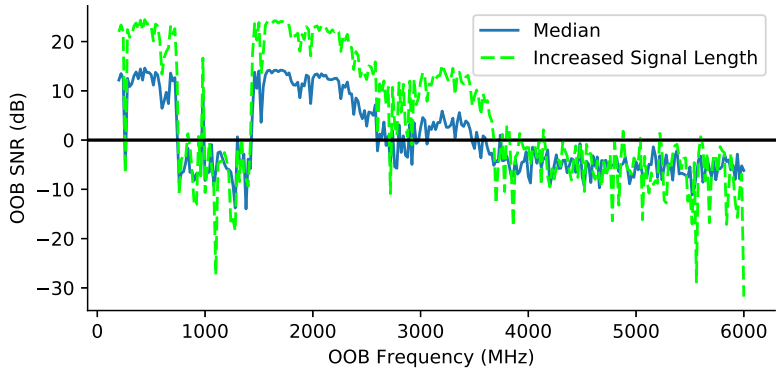

Figure 5-16: BT-100 Amplified BLadeRF RFID Tag Response

Chapter 6

Conclusion

In this thesis, I implemented a portable, scalable, and low-cost RFID localization platform by leveraging the latest technology in RFID micro-location and software-defined radios. I showed that the system can match state-of-the-art accuracy, overcome the impacts of tag orientation, and perform well with different RFID tags, environments, and scenarios. Most importantly, this system can help extend RFID localization beyond research environments by making it easy and cost-effective to set up a micro-location system that is robust to real-world scenarios.

6.1 Applications for Robotic Grasping

To understand the potential of the localization system better, I demonstrate a novel application of how the low-cost localization platform can enable entirely new functionality and tasks. The application was completed as joint work with my Signal Kinetics colleagues for a recently submitted conference paper, and all of the work for this thesis contributed to the localization system on the robot. Figure 6-1 shows the high-level setup. There is an RFID tagged item on the table, which can be hidden from sight, and the robot holds a wrist-mounted camera and antenna by its gripper. The system fuses RF and visual sensor data (from the camera and antenna) to locate, maneuver toward, and grasp items in line-of-sight and non-line-of-sight.

While robotic search and grasping is a well-researched field, these system con-

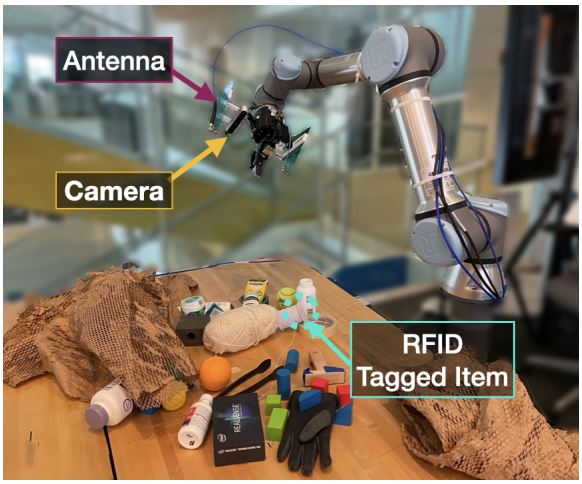

Figure 6-1: Robotic Grasping Setup. The robot has a wrist-mounted camera and antenna that are used to located a target RFID-tagged item.

ventionally use vision camera which have two significant problems. First, vision only works when the object of interest is visible, but many robotic tasks require finding objects that are hidden behind others, placed out of line-of-sight, or just in non-optimal lighting. Second, vision algorithms are prone to error in object identification, which can cause serious issues for supply chains and fulfillment centers that rely on robotic picking and packing.

A portable, low-cost RFID localization system can address both these problems and more. Specifically, RFIDs can be localized in non-line-of-sight and they also provide a highly accurate system of verification. Additionally, my thesis work introduced a system that is much smaller and more portable than previous RFID localization systems, making it well suited for mobile robots and robot-mounted systems like this one. Lastly, by overcoming issues of orientation, my system can provide continuous sensor data of all the objects in the workspace without blind spots due to tag placement. In summary, my portable, affordable RFID platform can help bridge the gap into real-world robotic systems.

Figure 6-2 summarizes the high-level idea behind the robotic system. The RFID localization system is used in all three steps. First, the RFID location is determined by taking multiple RF measurements of the workspace, where the target object is in an unknown position. Once the target item is found, any obstructing items in the

way are removed and placed to the side. Lastly, the target object is picked up and the RF system can verify that it is in fact the correct object. This approach improves the efficiency and robustness of robotic grasping tasks, and is just one example of how a portable RFID localization platform can improve many industries.

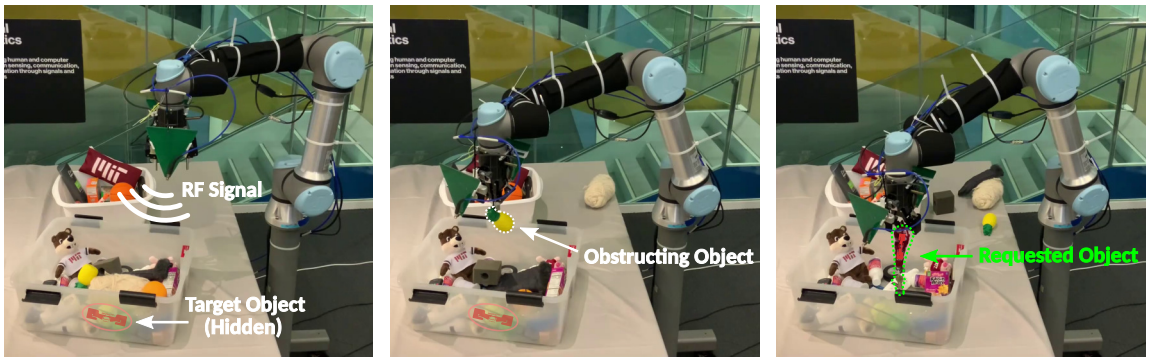

(a) Locating (b) Decluttering (c) Verification

Figure 6-2: RF and Vision Sensor Fusion for Robotic Grasping. The robotic system has three main steps to find and retrieve hidden items tagged with RFIDs.

6.2 Future Directions

Although this thesis made significant progress in bringing RFID localization to real-world deployments, there are additional improvements that can be made. Primarily, the speed and frame rate of the system can be increased by streamlining the localization code, pre-processing data with the FPGA, adding hardware timing synchronization. I briefly explored using a hardware trigger to initiate transmit and receive on multiple bladeRF devices simultaneously, and getting this fully functional would remove the need for post-collection software alignment. Additionally, it would be useful to explore the limits of frequency hopping rate, which would decrease the wait times between channel estimates. Furthermore, much of the data processing pipeline is spent doing channel estimation from the raw RF data. Moving these data-heavy, straightforward tasks to the built-in FPGA would free up compute resource for the rest of the data pipeline. Doing this would also allow the RFID localization platform to be use in resource-constrained tasks such as edge computing on mobile robots.

Apart from the system frame rate, localization range is also a limiting factor for many potential use cases. The bladeRF hardware is a good start, but future iterations of this system can explore custom hardware that will increase device sensitivity and range. In particular, the hardware can be optimized for a sufficient bandwidth of frequencies in which RFIDs have the strongest response. These hardware improvements would also make it easier to perform robust localization under FCC regulations.

Overall, RFIDs are an accurate and robust approach to micro-location for Industry 4.0, and I expect to see more widespread adoption as the systems get even smaller and more cost effective.

Bibliography

[1] Abeeway. Tracking for warehouses and indoor facilities. https://www.abeeway.com/tracking-for-warehouses-and-indoor-facilities/, 2021.

[2] Salah Azzouzi, Markus Cremer, Uwe Dettmar, Rainer Kronberger, and Thomas Knie. New measurement results for the localization of UHF RFID transponders using an Angle of Arrival (AoA) approach. In *2011 IEEE International Conference on RFID*, pages 91–97, 2011.

[3] Mathieu Bouet and Aldri L. dos Santos. RFID tags: Positioning principles and localization techniques. In *2008 1st IFIP Wireless Days*, pages 1–5, 2008.

[4] Mathieu Bouet and Guy Pujolle. L-VIRT: A 3-D Range-Free Localization Method for RFID Tags Based on Virtual Landmarks and Mobile Readers. In *2009 6th IEEE Consumer Communications and Networking Conference*, pages 1–5, 2009.

[5] A. Buffi, M. R. Pino, and P. Nepa. Experimental Validation of a SAR-Based RFID Localization Technique Exploiting an Automated Handling System. *IEEE Antennas and Wireless Propagation Letters*, 16:2795–2798, 2017.

[6] Alice Buffi, Paolo Nepa, and Fabrizio Lombardini. A Phase-Based Technique for Localization of UHF-RFID Tags Moving on a Conveyor Belt: Performance Analysis and Test-Case Measurements. *IEEE Sensors Journal*, 15(1):387–396, 2015.

[7] Kirti Chawla, Christopher McFarland, Gabriel Robins, and Connor Shope. Real-time RFID localization using RSS. In *2013 International Conference on Localization and GNSS (ICL-GNSS)*, pages 1–6, 2013.

[8] Jae Hyung Cho and Myeong-Woo Cho. Effective Position Tracking Using B-Spline Surface Equation Based on Wireless Sensor Networks and Passive UHF-RFID. *IEEE Transactions on Instrumentation and Measurement*, 62(9):2456–2464, 2013.

[9] Ian Clester. *RFID Localization for Interactive Systems*. PhD thesis, MIT, September 2020.

[10] Dave3457. Circular.Polarization.Circularly.Polarized.Light Without.Components Right.Handed.svg. https://en.wikipedia.org/wiki/Circular_polarization#/media/File:Circular.Polarization.Circularly.Polarized.Light_Without.Components_Right.Handed.svg, March 2010.

[11] HackRF One Great Scott Gadgets. https://greatscottgadgets.com/hackrf/one/, 2021.

[12] Cory Hekimian-Williams, Brandon Grant, Xiuwen Liu, Zhenghao Zhang, and Piyush Kumar. Accurate localization of RFID tags using phase difference. In *2010 IEEE International Conference on RFID (IEEE RFID 2010)*, pages 89–96, 2010.

[13] Aikaterini D. Koutsou, Fernando Seco, Antonio R. Jimenez, Javier O. Roa, Joao L. Ealo, Carlos Prieto, and Jorge Guevara. Preliminary Localization Results With An RFID Based Indoor Guiding System. In *2007 IEEE International Symposium on Intelligent Signal Processing*, pages 1–6, 2007.

[14] Rainer Kronberger, Thomas Knie, Roberto Leonardi, Uwe Dettmar, Markus Cremer, and Salah Azzouzi. UHF RFID localization system based on a phased array antenna. In *2011 IEEE International Symposium on Antennas and Propagation (APSURSI)*, pages 525–528, 2011.

[15] LimeMicro, LimeSDR. https://limemicro.com/products/boards/limesdr/, 2021.

[16] Zhihong Luo, Qiping Zhang, Yunfei Ma, Manish Singh, and Fadel Adib. 3D backscatter localization for fine-grained robotics. In *16th ${$USENIX$}$ Symposium on Networked Systems Design and Implementation (${$NSDI$}$ 19)*, pages 765–782, 2019.

[17] Haishu Ma and Kesheng Wang. Fusion of RSS and Phase Shift Using the Kalman Filter for RFID Tracking. *IEEE Sensors Journal*, 17(11):3551–3558, 2017.

[18] Yunfei Ma, Nicholas Selby, and Fadel Adib. Minding the Billions: Ultrawideband Localization for Deployed RFID Tags. *ACM MobiCom*, 2017.

[19] A McWilliams. RFID: Technology, Applications, and Global Markets. Technical report, BCC Research, 2016. https://www.bccresearch.com/market-research/instrumentation-and-sensors/rfid-technology-applications-markets-report.html.

[20] Robert Miesen, Fabian Kirsch, and Martin Vossiek. UHF RFID Localization Based on Synthetic Apertures. *IEEE Transactions on Automation Science and Engineering*, 10(3):807–815, 2013.

[21] Andrea Motroni, Paolo Nepa, Alice Buffi, and Bernardo Tellini. A Phase-Based Method for Mobile Node Localization through UHF-RFID Passive Tags. In *2019 IEEE International Conference on RFID Technology and Applications (RFID-TA)*, pages 470–475, 2019.

[22] L.M. Ni, Yunhao Liu, Yiu Cho Lau, and A.P. Patil. LANDMARC: Indoor location sensing using active RFID. In *Proceedings of the First IEEE International Conference on Pervasive Computing and Communications, 2003. (PerCom 2003).*, pages 407–415, 2003.

[23] Theresa Nick, Jürgen Götze, Werner John, and Gerhard Stoenner. Localization of UHF RFID labels with reference tags and Unscented Kalman Filter. In *2011 IEEE International Conference on RFID-Technologies and Applications*, pages 168–173, 2011.

[24] Nuand. libbladeRF: Frequency Tuning on the bladeRF. https://nuand.chttps://nuand.com/libbladeRF-doc/v2.2.1/tuning.htmlom/libbladeRF-doc/v2.2.1/relnotes_2_0.html, 2021.

[25] Nuand, BladeRF 2.0 Micro. https://www.nuand.com/bladerf-2-0-micro/, 2021.

[26] Ettus Research. Usrp n210. http://www.ettus.com, 2017.

[27] Samer S. Saab and Zahi S. Nakad. A Standalone RFID Indoor Positioning System Using Passive Tags. *IEEE Transactions on Industrial Electronics*, 58(5):1961–1970, 2011.

[28] Simo Sarkka, Ville V. Viikari, Miika Huusko, and Kaarle Jaakkola. Phase-Based UHF RFID Tracking With Nonlinear Kalman Filtering and Smoothing. *IEEE Sensors Journal*, 12(5):904–910, 2012.

[29] Martin Scherhäufl, Markus Pichler, Erwin Schimbäck, Dominikus J. Müller, Andreas Ziroff, and Andreas Stelzer. Indoor Localization of Passive UHF RFID Tags Based on Phase-of-Arrival Evaluation. *IEEE Transactions on Microwave Theory and Techniques*, 61(12):4724–4729, 2013.

[30] Martin Scherhäufl, Markus Pichler, and Andreas Stelzer. Robust localization of passive UHF RFID tag arrays based on phase-difference-of-arrival evaluation. In *2015 IEEE Topical Conference on Wireless Sensors and Sensor Networks (WiSNet)*, pages 47–49, 2015.

[31] Nicholas Selby. *Rapid RFID Location and Orientation Recovery.* PhD thesis, MIT, Cambridge, MA, June 2018.

[32] Anastasios Tzitzis, Spyros Megalou, Stavroula Siachalou, Emmanouil Tsardoulias, Traianos Yioultsis, and Antonis G. Dimitriou. 3D Localization of RFID Tags with a Single Antenna by a Moving Robot and "Phase ReLock". In *2019 IEEE International Conference on RFID Technology and Applications (RFID-TA)*, pages 273–278, 2019.

[33] Goran Vasiljević, Damjan Miklić, Ivica Draganjac, Zdenko Kovačić, and Paolo Lista. High-accuracy vehicle localization for autonomous warehousing. *Robotics and Computer-Integrated Manufacturing*, 42:1–16, 2016. https://www.sciencedirect.com/science/article/pii/S0736584515300314.

[34] Jue Wang, Fadel Adib, Ross Knepper, Dina Katabi, and Daniela Rus. RF-Compass: Robot Object Manipulation Using RFIDs. In *Proceedings of the 19th Annual International Conference on Mobile Computing & Networking*, MobiCom '13, pages 3–14, New York, NY, USA, 2013. Association for Computing Machinery. https://doi.org/10.1145/2500423.2500451.

[35] Jue Wang and Dina Katabi. Dude, where's my card?: RFID positioning that works with multipath and non-line of sight. *ACM SIGCOMM Computer Communication Review*, 43(4):51–62, September 2013. https://dl.acm.org/doi/10.1145/2534169.2486029.

[36] Han Wu, Shizhen Qu, Dongdong Xu, and Chunlin Chen. Precise Localization and Formation Control of Swarm Robots via Wireless Sensor Networks. *Mathematical Problems in Engineering*, 2014:1–12, 2014. http://www.hindawi.com/journals/mpe/2014/942306/.

[37] Junru Zhou, Hongjian Zhang, and Lingfei Mo. Two-dimension localization of passive RFID tags using AOA estimation. In *2011 IEEE International Instrumentation and Measurement Technology Conference*, pages 1–5, 2011.

[38] Junyi Zhou and Jing Shi. RFID localization algorithms and applications—a review. *Journal of Intelligent Manufacturing*, 20(6):695–707, December 2009. http://link.springer.com/10.1007/s10845-008-0158-5.