

# Minding the Billions: Ultra-wideband Localization for Deployed RFID Tags

Yunfei Ma, Nicholas Selby, Fadel Adib  
Massachusetts Institute of Technology  
{yunfeima,nselby,fadel}@mit.edu

## ABSTRACT

State-of-the-art RFID localization systems fall under two categories. The first category operates with off-the-shelf narrowband RFID tags but makes restrictive assumptions on the environment or the tag's movement patterns. The second category does not make such restrictive assumptions; however, it requires designing new ultra-wideband hardware for RFIDs and uses the large bandwidth to directly compute a tag's 3D location. Hence, while the first category is restrictive, the second one requires replacing the billions of RFIDs already produced and deployed annually.

This paper presents RFind, a new technology that brings the benefits of ultra-wideband localization to the billions of RFIDs in today's world. RFind does not require changing today's passive narrowband RFID tags. Instead, it leverages their underlying physical properties to emulate a very large bandwidth and uses it for localization. Our empirical results demonstrate that RFind can emulate over 220MHz of bandwidth on tags designed with a communication bandwidth of only tens to hundreds of kHz, while remaining compliant with FCC regulations. This, combined with a new super-resolution algorithm over this bandwidth, enables RFind to perform 3D localization with sub-centimeter accuracy in each of the x/y/z dimensions, without making any restrictive assumptions on the tag's motion or the environment.

## CCS CONCEPTS

• **Networks** → **Cyber-physical networks**; *Mobile networks*; *Sensor networks*;

## KEYWORDS

RFID; Localization; Battery-free; UWB; Smart Environments

## ACM Reference format:

Yunfei Ma, Nicholas Selby, Fadel Adib. 2017. Minding the Billions: Ultra-wideband Localization for Deployed RFID Tags. In *Proceedings of MobiCom '17, Snowbird, UT, USA, October 16–20, 2017*, 13 pages. <https://doi.org/10.1145/3117811.3117833>

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*MobiCom '17, October 16–20, 2017, Snowbird, UT, USA*

© 2017 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN 978-1-4503-4916-1/17/10...\$15.00

<https://doi.org/10.1145/3117811.3117833>

## 1 INTRODUCTION

Accurate RFID localization can be a game-changer for many industries ranging from virtual reality to factory automation. For example, virtual reality systems, like the HTC Vive [28] and Facebook's Oculus Rift [44], rely on relatively large trackers like handheld motion controllers. Accurate RFID localization would enable us to replace these handheld trackers with on-body RFID stickers that can track multiple user limbs. Another application that can benefit from fine-grained RFID localization is packaging quality control in factories and warehouses. For example, today's packaging control ends once a box is sealed. However, since many of today's packaged items are already tagged with RFIDs, accurate RFID localization would enable employees to check the number of items in a box or whether the right item is in the right box even after the box is sealed. More generally, absolute RFID localization can enable many applications in retail stores, factories & warehouses, virtual & augmented reality, and smart environments.

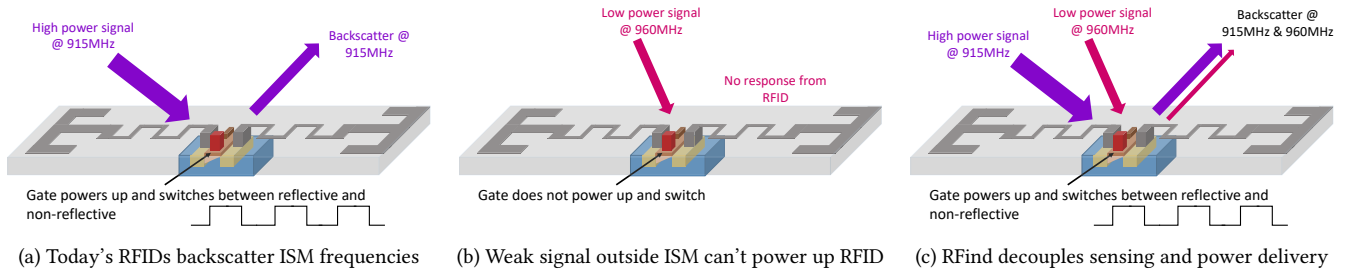
Indeed, the topic of RFID localization has gained much attention from the academic community over the past decade [31, 41, 48, 52, 53, 57]. However, none of the past proposals can enable ubiquitous localization and deliver on the applications described above. In particular, early proposals in this space relied on measuring the received signal strength (RSS) [16, 18, 38, 41, 59] and the angle of arrival (AoA) [14, 32, 60] and demonstrated a median accuracy of the order of tens of centimeters. Recent proposals have demonstrated finer centimeter-scale localization accuracy. However, these proposals either require furnishing the environment with a dense, surveyed grid of reference tags and localize by matching to reference tags [17, 53], and/or they require the tag or the reader to move over multiple wavelengths on a predefined trajectory at a predefined speed [40, 48, 57].

In this paper, we investigate whether we can achieve sub-centimeter RFID localization by measuring the time-of-flight (TOF) – i.e., the time it takes the signal to travel between a reader and an RFID. Accurate TOF measurements would allow us to localize RFIDs without reference tags and without prior trajectory knowledge. In particular, the TOF can be mapped to the distance traveled by taking into account the speed at which the RF signals travel.

The fundamental challenge in realizing this goal, however, is that accurate TOF-based localization hinges on the ability to measure time at a very fine granularity. In particular, achieving centimeter-scale localization would require hardware that can support very high sampling rate or very large bandwidth, often multiple GHz of bandwidth [11, 19, 34].<sup>1</sup> In contrast, RFIDs communicate at only tens of kHz, i.e., five orders of magnitude lower than such systems. Hence, TOF-based localization with such a communication

---

<sup>1</sup>Time resolution is inversely proportional to the bandwidth.



**Figure 1: RFInd leverages the physics underlying RFID switching.** For simplicity, the figure only shows an RFID's frontend circuitry, consisting of an antenna inlay and a switching transistor. (a) An RFID powers up when the reader transmits a high power signal in the ISM band, e.g., at 915MHz. To respond, it modulates the voltage of the switching transistor, resulting in two backscatter states: reflective and non-reflective. (b) An RFID cannot power up when a reader transmits an extremely low power signal outside ISM. (c) Decoupling power delivery and sensing: When a reader transmits both a high power signal in ISM and an extremely low power signal outside ISM, the RFID powers up (due to ISM-frequency) and reflects both frequencies simultaneously, allowing RFInd to sense the channel even outside the ISM band.

bandwidth would result in an accuracy of multiple kilometers. One option to overcome this challenge is to design new RFIDs that have large enough bandwidth to enable accurate localization [20, 23, 37]. Not only would such an approach require designing new hardware for RFID tags – making them significantly more expensive and non-compliant with today's FCC regulations and RFID communication protocols – but it would also leave out the billions of RFIDs already deployed in today's world [39].

We present RFInd, a system that can achieve sub-centimeter RFID localization by directly measuring the time-of-flight *without any hardware modification to passive narrowband RFID tags* and despite their bandwidth limitations. Using a single antenna, RFInd can compute the exact distance to an RFID tag. To achieve 2D or 3D localization, it uses two or three antennas respectively in close proximity to the first.

RFInd's localization algorithm is based on a realization that RFID modulation is *frequency agnostic*. Specifically, RFIDs communicate with a wireless device called a reader through backscatter technology. In backscatter systems, the reader transmits a continuous wave at some frequency, and the RFID switches its internal impedance between two states – reflective and non-reflective – to communicate bits to the reader, as shown in Fig. 1(a). By sensing subtle changes in the reflected signal due to the RFID's impedance changes, a reader can decode the bits communicated by an RFID tag. RFInd's realization is that such impedance changes may also be sensed at various frequencies. This is because reflectivity is fundamentally a physical process, similar to turning a mirror on and off.

RFInd harnesses the above realization to generate a *virtual* localization bandwidth that can be multiple orders of magnitude larger than the bandwidth of RFID communication. In particular, rather than transmitting a continuous wave at a single frequency, RFInd transmits continuous waves at multiple frequencies, as shown in Fig. 1(c). When an RFID switches its internal impedance to "reflective", it will reflect all the transmitted frequencies. On the other hand, when it changes its internal impedance to "non-reflective", it would absorb all the frequencies.<sup>2</sup> This effectively enables RFInd's

reader to estimate the RFID's channel at *all* the reflected frequencies. A large bandwidth enables the reader to compute the time-of-flight and use it to localize the RFID tag.

Translating this high level idea into a practical system, faces multiple challenges:

- First, the ISM band for UHF RFID is only 26 MHz wide [9]. Hence, even if a reader transmits at all of the frequencies within the ISM band, it would still be much smaller than the GHz-wide bandwidth necessary for centimeter-scale localization [11, 19, 34]. On the other hand, if it wishes to transmit outside the ISM band, it must limit itself to extremely low power due to FCC regulations; such power would not be sufficient to power up the tag as shown in Fig. 1(b).
- Second, transmitting and receiving over such a large bandwidth would require expensive hardware (high-speed ADCs, high-throughput I/Os, etc.) that can support Giga-samples/sec throughput.
- Finally, an RFID tag's signal does not arrive on a single path to the reader, but rather bounces off many reflectors in the environment before it arrives at the reader. This results in a classical challenge called multi-path, where an RF device receives several copies of the signal from the various reflectors. To overcome this challenge, past techniques require deploying a dense, surveyed grid of reference tags and localize by matching to the reference tags [53].

RFInd introduces multiple innovations that enable it to deal with the above challenges:

- First, it decouples the frequency for communication from that for localization. Specifically, it only uses frequencies within the ISM band for communication and powering up, and at the same time, it transmits very low power (sub-milliWatt) frequencies outside the ISM band, as shown in Fig. 1(c). This allows it to measure the channel over a large bandwidth while remaining compliant to FCC regulations.
- Second, instead of acquiring the entire bandwidth at once, it performs frequency hopping to emulate a large *virtual* localization bandwidth in the time domain. In particular, at every point in time, it only transmits at two frequencies (one inside the ISM

<sup>2</sup>Note that the amount of reflection/absorption can vary over frequency and is taken into account in RFInd's design.

and another outside the ISM band). Over time, it varies the carrier wave of the sensing frequency and estimates the channel at that frequency. Then, it stitches the channels at the various frequencies obtained from an RFID over time.<sup>3</sup> This enables it to transmit at a very narrow bandwidth at every point in time and operate entirely within the bandwidth (and sampling rate) capabilities of RFID readers on the market.

- Finally, it leverages the large virtual localization bandwidth to tease apart the various multi-paths in the environment (recall that time resolution is inversely proportional to bandwidth), and identify the path that arrives earliest in time as the line-of-sight path for localization. Then, it incorporates a super-resolution technique that enables it zoom in on the line-of-sight path to achieve sub-centimeter localization accuracy.

We implemented a prototype of RFind using USRP N210 software radios [8] and off-the-shelf Alien Squiggle RFID tags [12]. Our results demonstrate that RFind can emulate over 220 MHz of bandwidth on passive RFIDs. Combined with RFind’s super-resolution algorithm, this bandwidth enables it to perform 3D localization with median and 90<sup>th</sup> percentile errors less than 1 cm and 3 cm respectively in each of the  $x/y/z$  dimensions. Such accuracy matches or exceeds those reported by past state-of-the-art techniques [37, 53, 57]; however, in contrast to these past techniques, RFind does not require any reference tags and does not require any assumptions about the tag’s motion.

RFind has few additional desirable features. First, it is modulation independent – i.e., it can work with any backscatter modulation: FM0, Miller-8, etc. [3]. Second, it is fully compliant with the RFID communication protocol (the EPC Gen2 [3]). And third, it can operate in both line-of-sight and non-line-of-sight environments. However, one drawback of our current implementation is that it requires few seconds in order to localize. We would also like to note that, similar to any RFID reader, even though RFind’s reader is FCC compliant, marketing it commercially requires formal FCC certification [10].

**Contributions:** This paper presents RFind, the first system that can emulate a large bandwidth on passive narrowband RFIDs. RFind’s design introduces two key innovations: first, it introduces a technique that decouples the frequencies for power delivery and sensing in RFID communications; second, it presents a new super-resolution algorithm that operates over a large emulated bandwidth enabling RFind to achieve very high accuracy in 3D localization. The paper also presents a prototype implementation and evaluation of RFind demonstrating its accuracy in localizing RFIDs in multipath-rich environments, without reference tags and without requiring tag or reader motion.

## 2 RFIN D OVERVIEW

RFind is a system that enables sub-centimeter localization of UHF (ultra-high frequency) RFIDs. RFind’s localization works in line-of-sight, non-line-of-sight, and highly cluttered settings. As a result, it can operate in multipath-rich indoor environments. Moreover,

<sup>3</sup>This is possible because there is no carrier frequency offset (CFO) across time measurements since passive tags don’t generate their own RF signal but rather reflect the reader’s signal.

RFind is fully compliant with today’s standard UHF RFID protocol (the EPC-Gen2) as well as with FCC regulations for consumer electronics.

At a high level, RFind operates by estimating the time-of-flight from an RFID reader to an RFID tag. It then maps the time-of-flight to distance by taking into account the propagation speed of RF signals. To perform 1D localization, RFind leverages one receive antenna. To enable 2D or 3D localization, it employs two or three antennas respectively and performs trilateration.

RFind’s time-of-flight estimation has two components:

- **Emulating a Large Bandwidth on RFIDs:** The first component consists of a technique that enables RFind to emulate a large bandwidth on off-the-shelf RFIDs. The technique operates by decoupling the frequencies for power delivery and sensing in RFID communication. By varying the carrier wave of the sensing frequency over time, RFind can estimate the channel at each of these carriers. Then, it stitches the channel estimates at the various frequencies obtained from an RFID over time, realizing a large virtual bandwidth.
- **Localization by Multi-path Suppression and Super-resolution:** The second component of RFind is an algorithm that accurately localizes RFIDs using the large virtual bandwidth. The algorithm first identifies the line-of-sight path from RFind’s reader to an RFID, and eliminates multi-path reflections. Then, it zooms in on the line-of-sight path through a super-resolution algorithm to achieve sub-centimeter localization accuracy.

In §3 and §4, we describe the above components in details.

## 3 EMULATING A LARGE BANDWIDTH ON NARROWBAND RFIDS

In this section, we explain how RFind can emulate a large bandwidth on passive RFIDs.


### 3.1 Frequency Agnostic Modulation

We start by providing intuition into why RFID backscatter is frequency agnostic, then delve into how RFind leverages this property to emulate a large bandwidth on passive RFIDs. We describe the different states of RFID modulation at a high level and refer the reader to [22] for a more detailed explanation. Fig. 2(a) shows the components of a typical RFID tag with an antenna, an impedance switch, and a power harvesting unit.<sup>4</sup> The switch is controlled by the RFID’s logic (sequence of 0’s and 1’s), resulting in two states:

- *Non-reflective.* When the switch is off (i.e., logic=0), it acts as an open terminal as shown in Fig. 2(b). Since the input impedance of the power harvesting unit is designed to match the antenna, the received signal can flow into the circuit, enabling the tag to harvest as much power as possible. In this state, the RFID is absorptive or non-reflective.
- *Reflective.* When the switch is on (i.e., logic=1), it acts as a short terminal as shown in Fig. 2(c). Hence, the antenna is connected to ground rather than to a matched circuit.<sup>5</sup> Since the impedance is not matched, there is electromagnetic inhomogeneity which

<sup>4</sup>The impedance switch is usually implemented as a transistor between the antenna and the analog ground.

<sup>5</sup>Recall from basic circuit principles that the equivalent impedance of a circuit in parallel with ground is zero [25].



**Figure 2: RFID tag antenna impedance control by antenna switch.** (a) RFID tag circuit diagram consisting of an antenna, antenna switch, power harvesting unit, and logic & memory circuitry. (b) RFID tag in non-reflective state. Antenna switch is turned "off", resulting in an open terminal. RF power flows into the power harvesting unit. (c) RFID tag in reflective state. Antenna switch is turned "on", resulting in a short terminal. RF power gets reflected by the ground.

results in total reflection. In this state, the RFID is reflective, and all received power by the antenna is reflected (re-radiated) back to the reader.

In the above discussion, we described the two states as perfectly reflective or perfectly absorptive. In practice, however, the matching is not perfect. As a result, an RFID switches between more-reflective and less-reflective states. Nonetheless, the fact that an RFID switches between two states holds, and is sufficient for RFind to sense the channel at different frequencies. We call this the *frequency agnostic* property of backscatter modulation. RFind extends this concept to sensing reflectivity changes in the *complex* domain rather than only in the amount of reflection power, as we describe in detail in Appendix A.

### 3.2 Decoupling Sensing & Power Delivery


RFind leverages the frequency agnostic property of RFID modulation to estimate a tag's channel over a wide bandwidth. Specifically, rather than transmitting a single frequency  $f_1$  as in today's RFID protocol, it transmits at multiple frequencies, e.g.,  $f_1$  and  $f_2$ . When the switch is open, the tag absorbs both frequencies as shown in Fig. 2(b); when it is closed, the tag reflects both frequencies as shown in Fig. 2(c). In what follows, we describe how RFind expands on this idea to emulate a very large bandwidth on passive RFIDs.

#### 3.2.1 How large of a bandwidth can RFind sense?

A natural question is: over how large of a bandwidth can RFind communicate with an off-the-shelf RFID? Traditionally, due to FCC regulations, this bandwidth is limited to only 26 MHz. Specifically, recall that passive RFIDs are batteryless and harness power from the reader's RF signal to power up and communicate. To deliver sufficient power, today's RFID readers transmit around 36dBm in the ISM band [9].

Outside the ISM band, however, this power is limited to  $-13.3$  dBm and  $6.7$  dBm for average and peak power respectively,<sup>6</sup> i.e.,  $1,000\times$  lower than the peak power of frequencies within the ISM band and  $100,000\times$  lower than their allowed average power as shown in Fig. 3.

To overcome this bandwidth limitation, RFind decouples RFID power delivery and sensing. In particular, it transmits at two frequencies simultaneously:  $f_p$  inside the ISM band (to power up the



**Figure 3: Decoupling Sensing and Power Delivery.** By decoupling the frequencies of sensing and power delivery, RFind can transmit at an extremely low-power sensing frequency  $f_s$  outside the ISM band. This enables it to emulate a large bandwidth while remaining compliant to FCC regulations.


tag) and  $f_s$  outside the ISM band (at much lower power) as shown in Fig. 3. An RFind reader uses  $f_p$  to power up the tag, and  $f_s$  to sense the tag's backscatter reflection. RFind exploits the fact that RFID backscatter is frequency agnostic to sense the channel at  $f_s$  despite its very low power. Finally, in order to sense the channel over a wide bandwidth, RFind varies  $f_s$  over time and collects channel measurements.

To investigate the feasibility of this idea, we run an experiment where we transmit at one high-power frequency and sweep the low-power sensing frequency  $f_s$  over 300 MHz. In this experiment, the tag is in direct line-of-sight of RFind's antenna. Fig. 4 shows the output of our experiment in blue when we plot the SNR (Signal-to-Noise Ratio) over the sensing frequency. We make few observations about this figure:

- First, RFind can sense the channel over a bandwidth of over 300 MHz, i.e., around  $10\times$  larger than the ISM band for UHF communication, despite the low-power sensing frequency  $f_s$ .
- Second, the SNR is highest at the ISM band (between 902-926 MHz). Such result is expected since an RFID tag's antenna and matching circuit are designed to be flat over this frequency range. Outside this frequency range, the SNR fluctuates since the RFID's circuitry is not matched.
- Third, despite the fact that  $f_s$  is more than 30 dB lower than  $f_p$  in the ISM band, the SNR degradation can be relatively insignificant in some parts of the spectrum. For example, at 950 MHz, the SNR is the same as within the ISM band, despite the significantly lower-power signal. The underlying reason is that by

<sup>6</sup>FCC Part 15.231 rules that the average radiation limits for unlicensed low-power transmitter shall not exceed  $12,500\mu V/m$  at 3 meter distance, so in SI base units, the EIRP =  $0.3 \times E^2 = -13.3$  dBm [1]. The peak power can be higher when applying duty cycle, with a maximum 20dB peak-to-average ratio as ruled by Part 15.35 [9].





**Figure 4: SNR of a tag’s response when the sensing frequency is varied over a wide bandwidth.** The SNR (in blue) is consistently higher than 10dB over more than 300 MHz despite that  $f_s$  is extremely lower power outside the ISM band.

reducing the transmit power at  $f_s$ , an RFind reader also experiences significantly less noise. This owes to the fact that the main source of noise in RFID readers is the phase noise induced by the self-leakage of the transmitted signal [37].<sup>7</sup> In particular, recall that RFID readers are full-duplex systems, and they transmit and receive on the same frequency. Moreover, the power of the phase noise due to self-leakage is proportional to the power of the transmitted signal itself.<sup>8</sup> This is why the noise floor is reduced by the same ratio as the sensing frequency  $f_s$ , resulting in negligible SNR degradation over certain parts of the spectrum.

- Finally, we observe that the envelope of the SNR (red line in Fig. 4) degrades as we move further from the ISM band. This owes primarily to the beam pattern of the antennas of the RFID and the reader.

### 3.2.2 Channel Estimation

Now that RFind has a mechanism to sense the tag’s response at different sensing frequencies, it can apply standard channel estimation techniques to recover the channels at each of these frequencies. In particular, it uses the known preamble  $p_t$  of the tag’s response  $y_t$  to obtain an estimate of the channel  $h_k$  at a given sensing frequency  $f_k$  as follows:

$$h_k = \sum_t y_t p_t^* \quad (1)$$

By repeating this operation over different sensing frequencies, RFind can obtain channel estimates  $\{h_1 \dots h_K\}$  over a wide bandwidth.

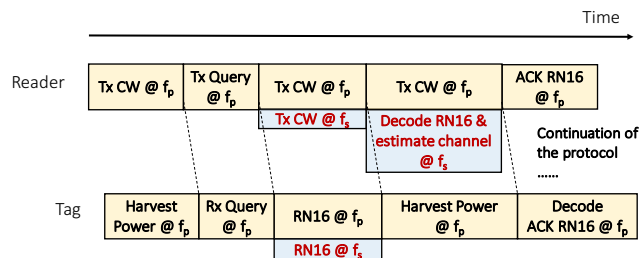
### 3.2.3 Integrating with the EPC-Gen2 Protocol

Next, we describe how RFind integrates its dual-frequency approach into the EPC-Gen2 Protocol. Fig. 5 shows the different stages of RFID communication:

- (1) A reader initiates a session by transmitting a continuous wave on a high-power signal at some frequency  $f_p$  in the ISM band. The tag harnesses power from the signal to power up and start decoding.
- (2) Then, the reader sends a Query command at  $f_p$  to the tag providing it with communication information (modulation, data rate, etc.). After a tag successfully decodes the Query, it

<sup>7</sup>In today’s RFID receivers, the phase noise floor is more than 80dB higher than its thermal noise floor [29, 33].

<sup>8</sup>In our implementation, we also filter out  $f_p$  in the analog domain at the receiver, which eliminates its effect on the phase noise.



**Figure 5: RFind is compatible with the RFID communication protocol.** To incorporate RFind into the EPC-Gen2 protocol, a reader only needs to transmit at two frequencies rather than one at the third stage of a communication session.

starts modulating its antenna impedance to transmit a 16-bit number called RN16.

- (3) At this moment, the reader transmits continuous waves at two frequencies  $f_p$  and  $f_s$ . As discussed in the previous section, because the backscatter modulation is frequency agnostic, both frequencies are reflected by the RFID. The reader calculates the channel at both frequencies
- (4) Then, the reader sends an ACK at  $f_p$  and proceeds with the remainder of the communication session exactly as per today’s EPC-Gen2 protocol [3].

Note that the fundamental difference between the above procedure and today’s RFID protocol is that during the RFID backscatter stage (i.e., step 3), the reader transmits at two frequencies ( $f_p$  and  $f_s$ ) rather than at just one frequency ( $f_p$ ). One might wonder if transmitting at  $f_s$  might be sufficient during this stage if all the reader needs is to measure the channel at  $f_s$ . Recall, however, the tag still needs to harvest power when it is in the absorptive (or non-reflective) state of its backscatter modulation as described in §3.1. In fact, our experiments demonstrated that without transmitting at  $f_p$ , the communication session is interrupted since the tag runs out of power during the backscatter stage. This is why the reader should transmit at both frequencies during the stage of backscatter modulation.

### 3.2.4 Putting it together

RFind’s technique to emulate a large bandwidth on passive RFIDs operates in three key stages:

- The first is a dual-frequency excitation technique, whereby the reader transmits at both a frequency  $f_p$  inside the ISM band and another extremely low power frequency  $f_s$  outside the ISM band. It uses  $f_p$  to power up the tag and communicate with it, and it uses  $f_s$  to estimate the channel of the tag at the corresponding frequency.
- Second, RFind repeats the above operation at different  $f_s$  carrier waves. In our implementation, RFind hops  $f_s$  over  $K$  consecutive carriers  $\{f_1, f_2, \dots, f_K\}$  as depicted in Fig. 3, where the spacing between adjacent carriers is equal to  $\Delta f$ . However, our technique is more general and can be used with a randomized hopping pattern.
- Finally, RFind leverages the fact that there is no carrier frequency offset (CFO) across the measurements. Hence, it can juxtapose the channel measurements obtained at different points in time

with the assumption that they were captured simultaneously.<sup>9</sup> This enables it to measure a large bandwidth of 300 MHz on off-the-shelf RFIDs.

It is worth noting that despite its large bandwidth of operation, incorporating RFind’s design in today’s readers likely requires only firmware modifications. Specifically, even though RFID readers can only transmit at a very narrow communication bandwidth, the center frequency of their communication can be varied over more than 100 MHz. This is because RFID readers are designed to accommodate for differences in regulations on the UHF band across regions and countries [3]. For example, while the UHF ISM band in the US is 902-928 MHz, China assigns two bands: 840.25-844.75 MHz and 920.25-924.75 MHz. In fact, this is also why RFID tag antennas are designed to be wideband despite that passive RFIDs can only support very narrow communication bandwidth of tens to hundreds of kHz. These underlying design properties make RFind amenable to practical use with today’s passive RFIDs and readers.

## 4 LOCALIZATION USING LARGE VIRTUAL BANDWIDTH

So far, we have demonstrated how RFind can emulate a large bandwidth on passive off-the-shelf RFID tags. Next, we discuss how RFind uses the large bandwidth in order to localize an RFID tag. In this section, we focus on how it can localize a single tag. However, the technique generalizes to any number of tags in the environment.

RFind’s localization algorithm operates in two stages. First, it leverages the bandwidth to tease apart the different paths traversed between an RFID and the reader, and identify the line-of-sight (LOS) path. Second, it zooms into the LOS path through super-resolution technique to achieve sub-centimeter localization accuracy.

### 4.1 Identifying the LOS path

In indoor environments, RF signals bounce off different obstacles (such as ceilings, walls, and furniture) before arriving at a receiver. This phenomenon is called the multipath effect. In order for RFind to localize an RFID, it first needs to identify the LOS path among all these paths. Below, we describe how RFind can identify the LOS path and obtain a rough time-of-flight estimate of that path.


Recall from §3.2.4 that RFind obtains the channel estimates in the frequency domain. To identify the LOS path, RFind needs to transform the channels from the frequency domain to the time domain, i.e., it needs to perform an inverse fourier transform. While there are various ways to implement an inverse fourier transform, RFind leverages the Inverse Fractional Fourier Transform (IFRFT) [15] since it incorporates an interpolation mechanism; hence, it provides RFind with a finer-granularity initial estimate of the time-of-flight.

Mathematically, let us denote the channel estimates as  $h_1, \dots, h_K$  at  $K$  different carrier frequencies. To obtain the time domain representation, RFind performs the following IFRFT operation:

$$S(\tau) = \sum_{k=1}^K h_k e^{j2\pi(k-1)\Delta f \tau} \quad (2)$$

where  $\tau$  denotes the delay in the time domain. The solid black line in Fig. 6 shows the output of this operation when it is performed over

<sup>9</sup>In contrast, this assumption does not hold for WiFi transmissions. Hence, past WiFi-based proposals that perform frequency hopping introduce various techniques to address offsets such as packet detection delay and CFO [51, 55, 56].



**Figure 6: Delay profiles at different bandwidths.** The delay profile computed with 220 MHz of bandwidth (in black) enables RFind to tease apart the different channel taps and identify the LOS as the path that arrives earliest in time. In contrast, a delay profile computed only with the 26 MHz ISM band does not provide sufficient resolution in isolating the paths.

an emulated bandwidth of 220 MHz. The figure plots the power as a function of the delay  $\tau$ . The delay profile exhibits multiple peaks, each of which corresponds to a different path traversed by the RFID’s signal. To identify the LOS path, RFind leverages the fact that the LOS arrives earliest in time, since all multi-path reflections traverse indirect paths as they bounce off reflectors, hence traversing a longer distance and experiencing a longer delay. It then uses the delay corresponding to the first peak as an initial estimate for the distance to an RFID.

To demonstrate the significance of a large bandwidth, we repeat the IFRFT operation over only the 26MHz ISM band and plot the delay profile as a dashed red line in Fig. 6. The plot shows that with such a small bandwidth, we are unable to tease apart the LOS path from the indirect paths. Intuitively, this is because when the bandwidth is smaller, the different paths merge into each other. Mathematically, at the output of an inverse fourier transform, each path is convolved with a sinc function whose width is inversely proportional to the bandwidth. Specifically, if there are  $L$  paths with delays  $\{\tau_1 \dots \tau_L\}$ , we can write the output of the IFRFT as:

$$\sum_{i=1}^L a_i \text{sinc}(B(t - \tau_i)) \quad (3)$$


where  $B$  is the bandwidth and  $a_i$  is the complex amplitude of the corresponding path. Hence, larger  $B$  results in fatter sinc functions. In particular, the resolution in separating multipath is the width of the sinc function main-lobe, given by [11]:

$$\text{Multipath Separability} = 1/B \quad (4)$$

Note that the final resolution in estimating each of the paths can be much finer, and is a function of the SNR and the sparsity of multipath [24]. In §6, we empirically evaluate the performance of RFind as a function of the sensing bandwidth.

### 4.2 Super-resolution Algorithm

In the previous section, we described how RFind can identify the LOS path and provide an initial distance estimate. However, this estimate is biased by noise and by leakage from other multipath components (due to the sinc effect described in §4.1). Below, we describe how RFind refines the initial estimate using a super-resolution algorithm.



**Figure 7: RFind’s super-resolution technique consists of a clustering algorithm that resolves phase cycle ambiguity. It constructs a cluster by choosing one potential distance from each frequency and then finds the cluster which has the minimum WCSS.**

To refine the distance estimate, we would like to leverage phase information. In particular, in the presence of a single LOS path, the phase  $\phi_k$  on the  $k$ -th carrier can be written as a direct function of the distance  $d$ :

$$\phi_k = \frac{2\pi}{\lambda_k} d \mod 2\pi \quad (5)$$

Leveraging this phase, however, is complicated by multiple factors. First, the phase can only be measured  $\text{mod } 2\pi$ ; this creates ambiguity in resolving the distance ( $\text{mod } \lambda$ ). Second, the above equation assumes a single LOS path and ignores both noise and multipath.

RFind’s solution to these challenges consists of two steps. First, it uses the initial distance estimate from §4.1 as a filter to mitigate the impact of multipath and recover a phase estimate that can be approximated by Eq. 5 at each of the frequencies  $f_k$ . Second, it performs an optimization algorithm across the approximate phases calculated at the different frequencies to mitigate the impact of residual noise and leakage. In what follows, we explain these steps in details.

#### 4.2.1 Obtaining Filtered Phase Estimates

Recall that by performing an inverse Fourier transform, RFind moves the channel estimates from the frequency domain to the time domain. In doing so, it loses the structure of the phase information at each of the subcarriers according to Eq. 5. In order to recover that information, it needs to transform the delay profile back to the frequency domain, while mitigating multipath.

To recover phase estimates at each of the frequencies  $f_k$  while mitigating multipath, RFind constructs a transform that exploits the LOS estimate of the distance  $\tilde{d}_0^c$  as a filter. Specifically, it projects the channels  $h_k$  at the different frequencies on its estimate of the channel caused by the LOS path. In Appendix B, we demonstrate that such projection can be realized through the following operation:

$$\theta_k = \mathcal{L} \sum_{i=1}^K h_k e^{j \frac{2\pi}{c} (f_i - f_k) \tilde{d}_0^c} \quad (6)$$

where  $c$  is the speed of propagation. Effectively, this operation reinforces the signal strength on line-of-sight path and suppresses the leakage from the multi-path reflections.

#### 4.2.2 Combining Estimates Across Frequencies

Now that we have the filtered phases at different frequencies, we formulate an optimization function that combines them in order to resolve phase ambiguity. RFind’s optimization algorithm is inspired by past work that leverage multi-frequency techniques for localization [37, 51, 56]. In contrast to past algorithms, RFind takes advantage of the fact that it already has an initial distance estimate.

Specifically, RFind’s search is bounded by Eq. 4 which limits the potential candidate distances to within a search range of  $c/B$ . Rather than searching over an infinite number of potential candidate distances due to the  $2\pi$  ambiguity of Eq. 5, the number of potential candidates from each  $\theta_k$  is:

$$\# \text{ candidates} = \frac{c}{B\lambda} \quad (7)$$

Since RFind can emulate a bandwidth over  $B = 220$  MHz on UHF RFIDs (whose  $\lambda$  is 33 cm), its search is limited to five candidate distances from each frequency  $k$ .

Fig. 7 demonstrates the intuition underlying RFind’s optimization function. The figure shows the five candidate distances at each  $\lambda_k$  (or frequency  $f_k$ ). The potential candidates can be clustered into different groups. Finally, the cluster that has the smallest width would correspond to the true location, since it is the one most robust to noise and leakage. Mathematically, RFind constructs different clusters  $\mathcal{C}$ , each of which consists of one distance estimate from each frequency. Then, it selects the cluster that has the minimum within-cluster sum of squares (WCSS). We can formulate the optimization function as follows:

$$\arg \min_{\mathcal{C}} \sum_{\hat{d} \in \mathcal{C}} |\hat{d} - \mu|^2 \quad (8)$$

where  $\mu$  is the cluster center. The optimization function can be solved in linear time by exploiting the fact that the unwrapping function is monotonic across  $\lambda_k$ . Said differently, a given candidate at  $\lambda_{k+1}$  can be mapped to the closest unwrapped candidate at  $\lambda_k$ .

After identifying the minimum WCSS cluster, the super-resolution distance estimate can be expressed as:

$$\tilde{d}_0^s = \mu \quad (9)$$

The above algorithm enables RFind to find the distance from a tag to a single receiving antenna, which determines a circle in 2D and a sphere with a fixed radius in 3D space. To obtain 2D or 3D locations, RFind leverages two or three antennas respectively and performs trilateration.

## 5 IMPLEMENTATION & EVALUATION

We implement a prototype of RFind using USRP N210 software radios [8] and test it with a variety of commercial RFID tags [12, 45, 49].

**Reader Implementation.** We adapt a USRP RFID reader developed in [30] and integrate RFind’s design into the EPC Gen2 protocol as described in §3.2.3. RFind’s transmit-side implementation of an RFID reader uses two USRPs with SBX daughterboards [7]: the first USRP transmits at 30 dBm at a frequency  $f_p$  for power delivery and communication<sup>10</sup> and the second USRP transmits a sensing frequency  $f_s$  at extremely low power (with an average radiation

<sup>10</sup> $f_p$  can be inside the ISM band (902-928 MHz) or in white spaces.

power at  $-15\text{dBm}$  and a peak power at  $-3\text{dBm}^{11}$ ) and sweeps it over  $220\text{MHz}$  bandwidth. These transmit powers are compliant to FCC regulations for consumer electronics [9]. The two USRPs are synchronized by an external clock [2].

To perform 3D localization, RFind’s receive-side implementation uses three USRP N210, each with a patch antenna [6], an external receive chain, and an LFRX daughterboard [5]. We design RFind’s external receive chain such that it performs coherent decoding similar to an off-the-shelf reader. The receive chain consists of a filter, a variable gain low noise amplifier (LNA), and an I/Q mixer. The filter eliminates strong leakage from the power delivery carrier  $f_p$ , and is essential to mitigate self-jamming and reduce the phase noise induced by the high-power self-leakage from  $f_p$ . After filtering, the received signal is amplified by an LNA and down-converted to baseband by mixing with the sensing frequency  $f_s$  through an I/Q mixer that feeds to an LFRX daughterboard of the USRP. The USRPs samples baseband I/Q signals which are postprocessed in MATLAB.


Our MATLAB implementation incorporates a Chebyshev-I digital bandpass filter that rejects residual low-frequency noise then performs matched filtering to recover the channel estimates. We perform a one-time calibration step to account for over-the-wire offsets and for changes in the reflection coefficient at different frequencies. In our evaluation, the estimated channels are divided by those acquired during the calibration step. The channels are then processed according to the algorithms described in §4 to obtain a tag’s 3D location.

**Latency.** Our current implementation requires around 6.4 seconds to output a location. This latency is primarily due to RFind’s frequency hopping pattern (i.e., the second step of §3.2.4) as it sweeps over  $220\text{MHz}$ . Specifically, we set  $\Delta f$  to  $10\text{MHz}$ , and the USRP-based reader requires around  $130\text{ms}$  to switch and lock to a new sensing frequency  $f_s$  in order to collect channels at that frequency. Note that this latency could be decreased by increasing the hopping step  $\Delta f$ ; however, increasing  $\Delta f$  may result in time-domain aliasing of the channel taps in multi-path rich environments. This is because increasing  $\Delta f$  results in sub-sampling the frequency domain, which leads to aliasing in the time domain. The super-resolution algorithm latency is sub-millisecond.

**Commercial RFID tags.** Unless otherwise noted, our experiments are performed with the most widely deployed type of passive RFIDs: the Alien Squiggle RFIDs [12]. Each of these tags costs 5-10 cents. To demonstrate the generality of RFind’s techniques, we also tested a variety of other commercial tags such as the Omni-ID Exo tag [45] and the Smartrac tag [49].

**Evaluation.** We evaluate RFind in multi-path rich indoor environments, and test it in both LOS and NLOS settings. Our experiments are performed in an office building with different types of indoor reflectors including tables, chairs, computers, ceilings, and walls. The experimental environment consists of an office lounge that spans an area of  $10\text{m} \times 12\text{m}$ , and we performed localization experiments in multiple sites against different multipath backgrounds. The area includes office cubicles that are separated by dividers consisting of  $20\text{cm}$  thick  $2\text{m}$ -tall separators made of two layers of wood. For

<sup>11</sup>Note that  $f_s$  is heavily duty-cycled due to the hopping pattern.



**Figure 8: CDF of 2D Localization Accuracy.** The figure shows 2D accuracy CDFs of RFind (blue), RFIDraw (red), and AoA (green).

NLOS experiments, these separators ensure that there is no LOS path between RFind’s antennas and the localized RFIDs. Finally, unless otherwise noted, RFind’s receive antennas are separated by a distance of  $20\text{cm}$  in the  $x/y/z$  dimensions from its transmit antenna.

**Baselines.** We compare RFind to state-of-the-art RFID localization schemes. Our baselines are: RFIDDraw [54] and the AoA [14]. Similar to RFind, these state-of-the-art schemes do not require reference tags or reader/tag motion.<sup>12</sup> We implement RFIDDraw and AoA with 8 antennas spanning an area of  $2\text{m} \times 2\text{m}$ , following the antenna patterns described in the respective papers. We implement RFind with only three receive antennas, placed within  $20\text{cm}$  of the transmit antenna.

**Ground Truth.** To obtain accurate ground truth RFID locations, we use the Bosch GLM 35 Laser Measure [4], which enables us to measure distances with sub-millimeter precision. Such precision is necessary since RFind achieves millimeter-scale ranging accuracy.

## 6 RESULTS

### 6.1 Comparison to State-of-the-Art

We start by comparing RFind’s localization accuracy to the state-of-the-art techniques described in §5: RFIDDraw and AoA. In fairness to these past techniques, we focus on 2D localization since both RFIDDraw and AoA are evaluated in 2D. In §6.2, we evaluate RFind’s 3D localization accuracy.


We performed 150 experimental trials. In each trial, an RFID tag was placed randomly within the evaluation environment described in §5. Since the read range of off-the-shelf RFIDs is constrained by the ability to power them up at distances larger than 5-6 m, we discard instances where the RFIDs do not respond. Fig. 8 plots the CDF of 2D localization error for RFind, RFIDDraw, and AoA. We observe the following:

- RFind achieves a median accuracy of  $0.91\text{cm}$  and a  $90^{\text{th}}$  percentile accuracy of  $1.92\text{cm}$  in 2D localization. In contrast, RFIDDraw’s median and  $90^{\text{th}}$  percentile accuracy is  $19\text{cm}$  and  $61.6\text{cm}$ ; these results are comparable to RFIDDraw’s implementation which also reports a median localization accuracy of  $19\text{cm}$ .<sup>13</sup> This demonstrates that RFind achieves  $21\times$  improvement for the median accuracy and more than  $32\times$  improvement for the  $90^{\text{th}}$  percentile accuracy over state-of-the-art localization systems that do not

<sup>12</sup>In contrast, past proposals that can achieve higher accuracy require moving the tag or the reader on a predefined trajectory at a predefined speed [40, 48, 57].

<sup>13</sup>[54] calls this the initial position accuracy.





**Figure 9: Sample output of the estimated RFID locations by RFind and RFIDraw.** RFind’s location estimations (indicated by crosses) enable pinpointing the location of the items they tag whereas RFIDraw’s estimations (indicated by x’s) cannot assign an RFID’s location to the item it tags.

require reference tags or motion. Moreover, RFind can achieve this improvement despite the fact that it operates with only three closely-spaced receive antennas in contrast to RFIDraw which needs 8 antennas that are separated by 2 m. Such improvement is expected since RFind incorporates mechanisms to directly estimate the time-of-flight on off-the-shelf tags.

- Both RFind and RFIDraw outperform AoA approaches, which achieve a median accuracy of 42.4cm and a 90<sup>th</sup> percentile error of 129cm.
- Finally, to demonstrate the significance of such accuracy, we show one of our experimental results in Fig. 9. The figure shows four RFID-tagged items placed on a mat and overlays the estimated locations by RFind and RFIDraw. The figure demonstrates that RFind can pin down the locations of these items while RFIDraw lacks the ability to assign the RFID to the location of the item it tags. These results demonstrate that RFind can be used for localization applications requiring high accuracy such as in supply chains, factories, and retail stores.

## 6.2 3D Localization Accuracy

Next, we evaluate RFind’s 3D localization accuracy in both LOS and NLOS settings. In our evaluation, we place two of RFind’s receive antennas on the ground, separated by 80 cm and elevate a third antenna by 60 cm above the ground. We evaluate RFind’s performance by placing an RFID at different 3D locations throughout the experimental environment described in §5. We also vary the location of the reader setup throughout the evaluation area. Since the read range of off-the-shelf RFIDs is constrained by the ability to power them up from a distance, we discard instances where the RFIDs do not respond. For NLOS experiments, the reader’s antennas are separated from the RFID by 20 cm-thick office dividers made of two layers of wood. Similar to past RFID literature [14, 42, 52, 54], we cannot do cross-room validation since a reader cannot power an RFID from another room.

We run 80 experimental trials and plot the CDF of accuracy along the  $x/y/z$  dimensions for the LOS and NLOS scenarios in Figs. 10(a) and (b) respectively. The figures reveal the following findings:

- In both LOS and NLOS settings, the median error is less than 1.1 cm along each of the  $x$ ,  $y$ , and  $z$  dimensions. Moreover, even

the 90<sup>th</sup> percentile error is less than 2 cm in  $x/y$  and less than 4 cm in the  $z$  dimension.

- The accuracy in LOS is higher than in NLOS settings. Such a result is expected since the SNR of the line-of-sight path degrades in NLOS, resulting in lower accuracy.
- The accuracy is higher in  $x/y$  than in the  $z$  dimensions. This is due to the fact that in our 3D evaluation experiments, the RFID was oriented vertically; hence, it is longer along the  $z$  dimension than along the  $x/y$ .
- Finally, we note that our discussion in §4.1 assumes that in NLOS settings, the direct path is not completely blocked (e.g., by metal) yet it may be significantly attenuated by an obstacle. To generalize to scenarios where the direct path is completely blocked, RFind can add additional antennas and perform clustering and outlier rejection (similar to past localization proposals, e.g., [56]).

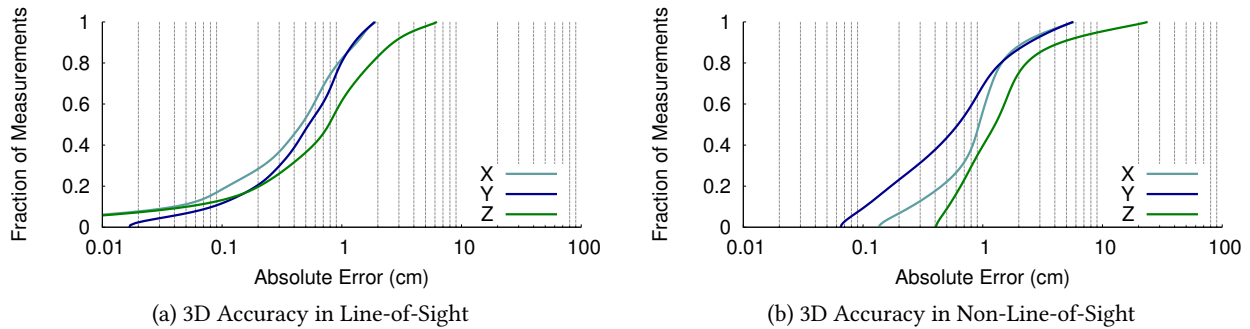
## 6.3 Accuracy vs. Different Parameters

Next, we would like to evaluate RFind’s localization accuracy as a function of different system parameters: distance, bandwidth, and antenna separation:

**Accuracy vs. Distance.** First, we evaluate RFind’s localization accuracy as a function of distance. For simplicity, we focus on 1D localization, and repeat the above experiments at 1 m increments. The ground truth is reported using the laser measure described in §5. Fig. 11(a) shows the median and 10<sup>th</sup> and 90<sup>th</sup> percentile errors as a function of distance from the receive antenna. The figure demonstrates that the median error increases gradually with distance, but remains about 1cm even at 6m. The 90<sup>th</sup> percentile at 6 m degrades to around 17 cm; this is due to the fact that at this distance, the SNR drops below the 3 dB threshold for a large number of frequencies, forgoing the benefits of a large bandwidth. Finally, note that beyond 6 m the RFID reader is unable to reliably power up an RFID tag similar to past reported work with USRP readers [53].

**Accuracy vs. Bandwidth.** Next, we evaluate RFind’s localization accuracy as a function of emulated bandwidth. Recall that RFind’s localization accuracy comes from its unique ability to emulate a large bandwidth on today’s off-the-shelf RFID tags. To understand the influence of bandwidth on localization accuracy, we vary the bandwidth of sensing frequencies provided to RFind’s localization algorithm described in §4. We use the same dataset from §6.2 but focus on 1D accuracy for simplicity.

Fig. 11(b) shows the impact of bandwidth on localization accuracy. The plot demonstrates that the localization accuracy monotonically improves with increased bandwidth. Specifically, if RFind only uses frequencies within the 26MHz ISM band, the median errors before and after super resolution are 73 cm and 33 cm respectively and the 90<sup>th</sup> percentile is over 3 m. The median error quickly decreases with bandwidth and drops below 1 cm for bandwidths larger than 120MHz. This is because increased bandwidth provides finer granularity in teasing apart the LOS path from the multipath. Note, however, due to the stochastic nature of indoor multi-paths, achieving sub-centimeter accuracy at the 90<sup>th</sup> is only realized when the bandwidth exceeds 200 MHz.



**Figure 10: 3D Localization Accuracy in LOS & NLOS** (a) shows RFind’s localization error in LOS in each of the x/y/z dimensions. (b) shows RFind’s localization error in NLOS in each dimension.



**Figure 11: Micro-benchmarks.** (a) shows that RFind’s accuracy gradually decreases with distance but the median error remains about 1 cm even at 6 m from the receiver. (b) shows that RFind’s accuracy increases with increased bandwidth both before and after super-resolution. (c) shows the independence of RFind’s accuracy from antenna separation. In each figure, error bars indicate  $10^{th}$  and  $90^{th}$  percentile accuracy.

**Accuracy vs. Antenna Span.** In past antenna array-based localization systems such as RFIDraw [54], a large antenna span (or aperture) is critical to achieve high accuracy since a larger aperture results in a narrower beamwidth. Hence, these systems require separating their antennas by more than 2 m. In contrast to these past systems, RFind can directly measure the time-of-flight; hence, we expect the antenna span to have less impact on its performance.

To evaluate the impact of antenna separation on RFind’s accuracy, we perform 2D localization experiments by gradually increasing the separation between RFind’s receive antennas. We run 50 experimental trials and plot the output in Fig. 11(c). The figure shows the median and  $90^{th}$  percentile accuracy as the antenna span is varied from 40 cm to 2 m. The figure demonstrates that the antenna span has minimal impact on 2D localization accuracy. In fact, RFind can achieve sub-centimeter localization accuracy even with 40cm antenna span. This is due to the fact that RFind can directly compute 1D estimates with sub-centimeter accuracy. Interestingly, the accuracy is slightly higher when the antenna span is smaller. This owes to a slight improvement in SNR when the antennas are closer to the RFID of interest with a smaller aperture. This ability to perform accurate localization using a small antenna span suggests that RFind’s design can enable compact and mobile RFID localization system.

## 7 RFIN IN APPLICATIONS

Finally, we would like to qualitatively test RFind in a number of real-world applications. Fig. 12 shows multiple instantiations of RFind for various applications. In each row of Fig. 12, the first column shows the test setup, the second column shows RFind’s output, and


the last two columns overlay the output over the real-world images. Across all these applications, RFind could achieve the same level of accuracy reported in the quantitative results above.

## 8 RELATED WORK


RFind is related to past work in three areas: localization of commercial RFID tags, ultra-wideband tag designs, and microscopic optical imaging. In contrast to all past work, RFind introduces the first system that emulates a large bandwidth on commercial tags, bringing the benefits of large bandwidth to billions of RFIDs already deployed. In doing so, it enables sub-centimeter 3D localization in multipath-rich environments. In what follows, we discuss how it relates to prior art.

**(a) RFID Localization.** The topic of RFID localization has gained much attention from the academic community over the past decade. Early proposals in this space relied on measuring the received signal strength (RSS) [16, 18, 41, 59], the angle of arrival (AoA) [14, 32, 38, 60], and the received signal phase [13, 35]. However, these proposals are based on line-of-sight assumptions and cannot deal with multipath. In particular, reflections off various objects in indoor environments (walls, ceilings, furniture) create constructive and destructive interference, making RSS and the phase of the received signal unpredictable. Moreover, since the strongest path between an RFID and reader may not always be the direct line-of-sight path but rather the one arriving off a side reflector, AoA can suffer from multi-meter errors in non-line-of-sight environments [53]. Moreover, in contrast to RFind, past multi-frequency techniques

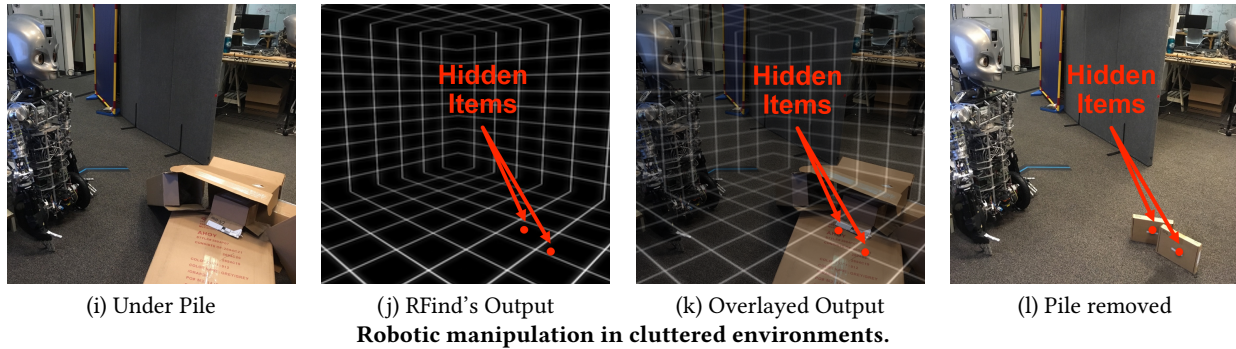




(a) Closed Box (b) RFind's Output (c) Overlaid Output (d) Opened box  
**Identifying and localizing items in closed boxes for packaging control and delivery services.**



(e) Behind Couch (f) RFind's Output (g) Overlaid Output (h) Couch removed  
**Tracking user gestures behind occlusions such as furniture.**



(i) Under Pile (j) RFind's Output (k) Overlaid Output (l) Pile removed  
**Robotic manipulation in cluttered environments.**

**Figure 12: RFind in Real-World Applications.** (a) shows closed boxes containing RFID-tagged objects. (e) shows a gamer whose fingers are tagged with RFIDs. (i) shows a pile of clutter under which lay two items the robot needs to find. (b), (f), and (j) show RFind's output in identifying and localizing the objects. (c), (g), and (k) overlay our output over the images. After removing the occlusions in (d), (h), and (l), we can see that RFind's localization is very accurate.

are restricted to the ISM band, and hence lack the ability to tease apart multipath from the line-of-sight path.

Recent proposals that can address multipath fall under two main categories: reference-based and motion-based techniques. Reference-based proposals require manually deploying a dense, surveyed grid of reference tags throughout the area of interest and localize a tag by matching it to the grid [17, 26, 53]. Motion-based proposals either require moving a reader's antenna over multiple wavelength (2-3 meters) [48, 52] or they require prior knowledge of the exact trajectory and exact speed of a tag and match it to a location on that trajectory, by leveraging inverse synthetic aperture radar (ISAR) or holographic techniques [40, 46, 57]. In contrast, RFind neither requires mechanical motion nor the deployment of reference tags in the environment. In particular, because RFind can directly measure the time-of-flight, it has intrinsic mechanisms that enable it to accurately localize RFIDs and eliminate multipath.

Finally, some recent proposals such as RFIDraw [36, 42, 54] demonstrate very high *tracking* accuracy but a *localization* accuracy of tens of centimeters. In particular, by measuring phase changes over time, these proposals can track small changes in the *relative*

location of a tag and recover its trajectory. Hence, they can recover a tag's trajectory but cannot know exactly where that trajectory was traced. In contrast, RFind enables sub-centimeter localization accuracy using time-of-flight measurements, demonstrating  $21\times$  localization improvement over these past proposals as we empirically demonstrated in §6.

**(b) Ultra-wideband RFID tags.** Past research has explored designing ultra-wideband tags. Due to their large bandwidth, these tags enable estimating the time-of-flight, and hence can deal with multi-path and achieve high accuracy localization [31, 37, 47, 58]. However, these proposals require building new hardware and, as a result, have multiple drawbacks: they are expensive, they are non-compliant with the RFID standard, and most of these proposals are not compliant with FCC regulations. More importantly, these proposals require replacing the billions of tags already produced and deployed annually.

In contrast to these past proposals, RFind leverages a realization that RFID backscatter communication is frequency agnostic,

enabling us to emulate ultrawide bandwidth on the billions off-the-self tags already deployed. In particular, RFind can sense an RFID tag's response even outside the ISM band by transmitting signals below FCC spurious emission limits, making it both EPC-Gen2 and FCC compliant. Moreover, RFind introduces multi-path filtering and super-resolution algorithms that enable it to zoom in on a tag's location with sub-centimeter precision, i.e., at a finer granularity than enabled by its sensing bandwidth.

**(c) Two photon microscopy.** RFind is inspired by a technique from optical physics called two-photon microscopy [21, 27], which is used for deep tissue imaging. The technique fires one photon to excite the layer of interest inside the tissue, and uses a second photon to perform imaging. These photons may be at the same or different visible-light frequencies. Similarly, RFind leverages a frequency inside the ISM band to excite or power up a tag and another frequency outside the ISM to perform sensing. However, RFind differs from two-photon microscopy in both techniques and problem domain. Specifically, in contrast to two-photon microscopy which was designed for improve fluorescence imaging in thick specimens, RFind is designed for RFID localization and makes a realization about RFID backscatter that enables it to accurately localize RFIDs.

## 9 CONCLUSION

We present RFind, a system that enables emulating a large virtual bandwidth on off-the-shelf passive RFIDs. In this paper, we use the large virtual bandwidth to estimate the time-of-flight and achieve centimeter-scale localization accuracy without reference tags or reader/tag motion. However, we believe that the implications of such large bandwidth extend beyond localization and pave way for exciting new avenues for exploration in RFID communication and security.

### Appendix A: Sensing Reflectivity Changes over a Wide Bandwidth.

The reflection of an RFID can be expressed as a function of the tag's electromagnetic (EM) reflection co-efficient  $\gamma$  [43]. Assume the incident electric field to the tag is  $E_{in}$ , we can express the reflected electric field  $E_{ref}$  as:

$$E_{ref} \propto E_{inc} \times \gamma \quad (10)$$

$\gamma$  is a complex number which a function of the frequency dependent antenna impedance  $Z_a(f)$  and the effective chip input impedance  $Z_c^{eff}(f)$ , and can be expressed as follows [43]:

$$\gamma = \frac{R_a(f)}{Z_a(f) + Z_c(f)^{eff}} \quad (11)$$

where  $Z_a(f) = R_a(f) + jX_a(f)$ .  $R_a(f)$  and  $X_a(f)$  are the frequency dependent real and imaginary parts of antenna impedance. The effective chip input impedance is affected by the switching transistor. Specifically, when the switch is open  $Z_c^{eff}(f) = Z_c(f)$ ; when it is closed,  $Z_c^{eff}(f) \approx 0$ . Hence,  $\gamma$  switches between:

$$Y_{open} = \frac{R_a(f)}{Z_a(f) + Z_c(f)} \quad \text{and} \quad Y_{closed} = \frac{R_a(f)}{Z_a(f)} \quad (12)$$

A RFind reader senses the reflected field difference in the complex domain which can be expressed as:

$$E_{diff} \propto E_{in}(Y_{closed} - Y_{open}) \propto E_{in} \frac{R_a(f)Z_c(f)}{Z_a(f)(Z_a(f) + Z_c(f))}$$

Typically, RFIDs are designed such that the field difference is maximized at the desired antenna center frequency. Outside the design bandwidth, the differential field gradually decreases due to impedance mismatch. Note that in RFID tag designs, matching is much more critical for power delivery on the downlink rather than it is for the backscatter response on the uplink. Since RFind powers up an RFID tag using one carrier inside ISM band, the tag would always power up and switch the impedance. Such complex impedance switching can be sensed outside the optimal frequency.

To illustrate this concept, consider the following simple practical model [50]: Assume  $Z_c = R$  which is purely resistive and an equivalent RLC circuit model for antenna,  $Z_a = R + j\omega L + \frac{1}{j\omega C}$ . The antenna center frequency is  $f_c = 900\text{MHz}$  and the quality factor of the antenna is  $Q = 10$ . So

$$Y_{closed} - Y_{open} = \frac{1}{(2 - jQ(1 - (\frac{f}{f_c})^2))(1 - jQ(1 - (\frac{f}{f_c})^2))} \quad (13)$$

The resultant SNR degradation from 900 MHz to 750 MHz is 17dB which is about the same amount of SNR degradation shown as measured from our experiment with a real tag in Fig. 4.

### Appendix B: Recovering Phase While Mitigating Multi-path.

We show that Eq. (6) reinforces the LOS path component and suppress leakages from multi-paths. Assume there are  $L$  multi-paths whose distances are  $d_1, d_2, \dots, d_L$ , and that  $d_0$  is the real LOS distance. The channel  $h_k$  can be expressed as:

$$h_k = a_0 e^{-j\frac{2\pi}{c} f_k d_0} + \sum_{l=1}^L a_l e^{-j\frac{2\pi}{c} f_k d_l} \quad (14)$$

where  $\tilde{d}_0^c$  is the coarse LOS distance estimate derived from IFRFT. By factoring in  $h_k$ , we rewrite Eq. (6) as:

$$\theta_k = \angle a_0 e^{-j\frac{2\pi}{c} f_k d_0} \sum_{i=1}^K e^{j\frac{2\pi}{c} (i-k)\Delta f (\tilde{d}_0^c - d_0)} + \angle \sum_{l=1}^L a_l e^{-j\frac{2\pi}{c} f_k d_l} \sum_{i=1}^K e^{j\frac{2\pi}{c} (i-k)\Delta f (\tilde{d}_0^c - d_l)} \quad (15)$$

Since  $(\tilde{d}_0^c - d_0)\Delta f/c \ll 1$ :

$$\sum_{i=1}^K e^{j\frac{2\pi}{c} (i-k)\Delta f (\tilde{d}_0^c - d_0)} \approx K \quad (16)$$

while  $\tilde{d}_0^c - d_l$  is large, resulting in

$$\left| \frac{\sum_{i=1}^K e^{j\frac{2\pi}{c} (i-k)\Delta f (\tilde{d}_0^c - d_l)}}{K} \right| \approx |\text{sinc}(B(\tilde{d}_0^c - d_l)/c)| \ll 1 \quad (17)$$

Thus, LOS is reinforced while multipath leakage is suppressed.

**Acknowledgments.** We thank Omid Abari, Dina Katabi, Zach Kabelac, our shepherd Shyam Gollakota, and the anonymous MobiCom reviewers for their insightful comments. This research is partially supported by NSF and the MIT Media Lab.

## REFERENCES

- [1] Calculating radiated power and field strength for conducted power measurements. <http://www.semtech.com>.
- [2] CDA2990. <http://www.ettus.com>. Ettus Inc.
- [3] EPC UHF Gen2 Air Interface Protocol. <http://www.gs1.org/epcrfid/epc-rfid-uhf-air-interface-protocol/2-0-1>.
- [4] Laser Measuring. <http://www.boschtools.com>. Bosch Inc.
- [5] LFRX daughterboard. <http://www.ettus.com>. ettus inc.
- [6] MTI RFID antenna. <http://www.mtiwe.com>. MTI Wireless Edge.
- [7] SBX daughterboard. <http://www.ettus.com>. ettus inc.
- [8] usrp n210. <http://www.ettus.com>. ettus inc.
- [9] *Understanding the Fcc Regulations for Low-power, Non-licensed Transmitters*. Office of Engineering and Technology Federal Communications Commission, 1993.
- [10] The State of RFID Implementation and Its Policy Implications: An IEEE-USA White Paper. IEEE USA, 2009.
- [11] F. Adib, Z. Kabelac, D. Katabi, and R. C. Miller. 3d tracking via body radio reflections. In *Usenix NSDI*, 2014.
- [12] Alien Technology Inc. ALN-9640 Squiggle Inlay. [www.alientechnology.com](http://www.alientechnology.com).
- [13] D. Armitz, K. Witrisal, and U. Muehlmann. Multifrequency continuous-wave radar approach to ranging in passive uhf rfid. *IEEE Transactions on Microwave Theory and Techniques*, 57(5):1398–1405, 2009.
- [14] S. Azzouzi, M. Cremer, U. Dettmar, R. Kronberger, and T. Knie. New measurement results for the localization of uhf rfid transponders using an angle of arrival (aoa) approach. In *RFID (RFID), 2011 IEEE International Conference on*, pages 91–97. IEEE, 2011.
- [15] D. H. Bailey and P. N. Swartztrauber. The fractional fourier transform and applications. *SIAM review*, 33(3):389–404, 1991.
- [16] M. Bouet and A. L. Dos Santos. Rfid tags: Positioning principles and localization techniques. In *Wireless Days, 2008. WD'08. 1st IFIP*, pages 1–5. IEEE, 2008.
- [17] M. Bouet and G. Pujolle. A range-free 3-d localization method for rfid tags based on virtual landmarks. In *Personal, Indoor and Mobile Radio Communications, 2008. PIMRC 2008. IEEE 19th International Symposium on*, pages 1–5. IEEE, 2008.
- [18] K. Chawla, C. McFarland, G. Robins, and C. Shope. Real-time rfid localization using rss. In *Localization and GNSS (ICL-GNSS), 2013 International Conference on*, pages 1–6. IEEE, 2013.
- [19] M. Chu, P. Jacob, J.-W. Kim, M. R. LeRoy, R. P. Kraft, and J. F. McDonald. A 40 gs/s time interleaved adc using sigle bimos technology. *IEEE Journal of Solid-State Circuits*, 45(2):380–390, 2010.
- [20] D. Dardari, R. D'Errico, C. Roblin, A. Sibille, and M. Z. Win. Ultrawide bandwidth rfid: The next generation? *Proceedings of the IEEE*, 98(9):1570–1582, 2010.
- [21] W. Denk, J. H. Strickler, W. W. Webb, et al. Two-photon laser scanning fluorescence microscopy. *Science*, 248(4951):73–76, 1990.
- [22] K. Finkelzeller. The rfid handbook, 2003.
- [23] L. Gang et al. Bandwidth dependence of cw ranging to uhf rfid tags in severe multipath environments. In *IEEE RFID 2011*.
- [24] S. Gezici, Z. Tian, G. B. Giannakis, H. Kobayashi, A. F. Molisch, H. V. Poor, and Z. Sahinoglu. Localization via ultra-wideband radios: a look at positioning aspects for future sensor networks. *IEEE signal processing magazine*, 22(4):70–84, 2005.
- [25] P. R. Gray and R. G. Meyer. *Analysis and design of analog integrated circuits*. John Wiley & Sons, Inc., 1990.
- [26] J. Han, C. Qian, X. Wang, D. Ma, J. Zhao, W. Xi, Z. Jiang, and Z. Wang. Twins: Device-free object tracking using passive tags. *IEEE/ACM Transactions on Networking*, 24(3):1605–1617, 2016.
- [27] F. Helmchen and W. Denk. Deep tissue two-photon microscopy. *Nature methods*, 2(12):932–940, 2005.
- [28] HTC. HTC Vive. <https://www.vive.com/us/>.
- [29] IEEE. *A new TX leakage-suppression technique for an RFID receiver using a dead-zone amplifier*, 2013.
- [30] N. Kargas, F. Mavromatis, and A. Bletsas. Fully-coherent reader with commodity sdr for gen2 fm0 and computational rfid. *IEEE Wireless Communications Letters*, 4(6):617–620, 2015.
- [31] N. C. Karmakar et al. Chipless rfid tag localization. *IEEE Transactions on Microwave Theory and Techniques*, 61(11):4008–4017, 2013.
- [32] R. Kronberger, T. Knie, R. Leonardi, U. Dettmar, M. Cremer, and S. Azzouzi. Uhf rfid localization system based on a phased array antenna. In *Antennas and Propagation (APSURSI), 2011 IEEE International Symposium on*, pages 525–528. IEEE, 2011.
- [33] I. Kwon, Y. Eo, H. Bang, K. Choi, S. Jeon, S. Jung, D. Lee, and H. Lee. A single-chip cmos transceiver for uhf mobile rfid reader. *IEEE Journal of Solid-State Circuits*, 43(3):729–738, 2008.
- [34] N. Levanon. Radar principles. *New York, Wiley-Interscience, 1988, 320 p.*, 1988.
- [35] X. Li, Y. Zhang, and M. G. Amin. Multifrequency-based range estimation of rfid tags. In *RFID, 2009 IEEE International Conference on*, pages 147–154. IEEE, 2009.
- [36] T. Liu, L. Yang, Q. Lin, Y. Guo, and Y. Liu. Anchor-free backscatter positioning for rfid tags with high accuracy. In *INFOCOM, 2014 Proceedings IEEE*, pages 379–387. IEEE, 2014.
- [37] Y. Ma, X. Hui, and E. C. Kan. 3d real-time indoor localization via broadband nonlinear backscatter in passive devices with centimeter precision. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking*, pages 216–229. ACM, 2016.
- [38] G. Mao, B. Fidan, and B. D. Anderson. Wireless sensor network localization techniques. *Computer networks*, 51(10):2529–2553, 2007.
- [39] A. McWilliams. RFID: Technology, Applications, and Global Markets. BCC Research, 2016.
- [40] R. Miesen, F. Kirsch, and M. Vossiek. Holographic localization of passive uhf rfid transponders. In *RFID (RFID), 2011 IEEE International Conference on*, pages 32–37. IEEE, 2011.
- [41] L. M. Ni, Y. Liu, Y. C. Lau, and A. P. Patil. Landmarc: indoor location sensing using active rfid. *Wireless networks*, 10(6):701–710, 2004.
- [42] P. V. Nikitin, R. Martinez, S. Ramamurthy, H. Leland, G. Spiess, and K. Rao. Phase based spatial identification of uhf rfid tags. In *RFID, 2010 IEEE International Conference on*, pages 102–109. IEEE, 2010.
- [43] P. V. Nikitin and K. S. Rao. Theory and measurement of backscattering from rfid tags. *IEEE Antennas and Propagation Magazine*, 48(6):212–218, 2006.
- [44] Oculus. Oculus Rift. <https://www.oculus.com/rift/>.
- [45] Omni-ID. Omni-ID Exo. [www.omni-id.com](http://www.omni-id.com).
- [46] A. Parr, R. Miesen, and M. Vossiek. Inverse sar approach for localization of moving rfid tags. In *RFID (RFID), 2013 IEEE International Conference on*, pages 104–109. IEEE, 2013.
- [47] M. Pelissier, J. Jantunen, B. Gomez, J. Arponen, G. Masson, S. Dia, J. Varteava, and M. Gary. A 112 mb/s full duplex remotely-powered impulse-ubw rfid transceiver for wireless nv-memory applications. *IEEE Journal of Solid-State Circuits*, 46(4):916–927, 2011.
- [48] L. Shangguan and K. Jamieson. The design and implementation of a mobile rfid tag sorting robot. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services*, pages 31–42. ACM, 2016.
- [49] Smartrac Group. Smartrac Shortdipole Inlay. [www.smartrac-group.com](http://www.smartrac-group.com).
- [50] W. L. Stutzman and G. A. Thiele. *Antenna theory and design*. John Wiley & Sons, 2012.
- [51] D. Vasisht, S. Kumar, and D. Katabi. Decimeter-level localization with a single wifi access point. In *Usenix NSDI*, 2016.
- [52] J. Wang, F. Adib, R. Knepper, D. Katabi, and D. Rus. RF-Compass: Robot Object Manipulation Using RFIDs. In *ACM MobiCom*, 2013.
- [53] J. Wang and D. Katabi. Dude, where's my card? rfid positioning that works with multipath and non-line of sight. In *ACM SIGCOMM*, 2013.
- [54] J. Wang, D. Vasisht, and D. Katabi. RF-idraw: virtual touch screen in the air using rf signals. In *ACM SIGCOMM*, 2015.
- [55] Y. Xie, Z. Li, and M. Li. Precise power delay profiling with commodity wifi. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*, pages 53–64. ACM, 2015.
- [56] J. Xiong, K. Sundaresan, and K. Jamieson. Tonetrack: Leveraging frequency-agile radios for time-based indoor wireless localization. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*, pages 537–549. ACM, 2015.
- [57] L. Yang, Y. Chen, X.-Y. Li, C. Xiao, M. Li, and Y. Liu. Tagoram: Real-time tracking of mobile rfid tags to high precision using cots devices. In *Proceedings of the 20th annual international conference on Mobile computing and networking*, pages 237–248. ACM, 2014.
- [58] C. Zhou and J. D. Griffin. Accurate phase-based ranging measurements for backscatter rfid tags. *IEEE Antennas and Wireless Propagation Letters*, 11:152–155, 2012.
- [59] J. Zhou and J. Shi. Rfid localization algorithms and applications—a review. *Journal of intelligent manufacturing*, 20(6):695–707, 2009.
- [60] J. Zhou, H. Zhang, and L. Mo. Two-dimension localization of passive rfid tags using aoa estimation. In *Instrumentation and Measurement Technology Conference (I2MTC), 2011 IEEE*, pages 1–5. IEEE, 2011.