

The Design and Implementation of a Mobile RFID Tag Sorting Robot

Longfei Shangguan[†] and Kyle Jamieson^{†,‡}
[†]Princeton University, [‡]University College London
{longfeis, kylej}@cs.princeton.edu

ABSTRACT

Libraries, manufacturing lines, and offices of the future all stand to benefit from knowing the exact spatial order of RFID-tagged books, components, and folders, respectively. To this end, radio-based localization has demonstrated the potential for high accuracy. Key enabling ideas include motion-based synthetic aperture radar, multipath detection, and the use of different frequencies (channels). But indoors in real-world situations, current systems often fall short of the mark, mainly because of the prevalence and strength of “multipath” reflections of the radio signal off nearby objects. In this paper we describe the design and implementation of MobiTagbot, an autonomous wheeled robot reader that conducts a roving survey of the above such areas to achieve an exact spatial order of RFID-tagged objects in very close (1–6 cm) spacings. Our approach leverages a serendipitous correlation between the changes in multipath reflections that occur with motion and the effect of changing the carrier frequency (channel) of the RFID query. By carefully observing the relationship between channel and phase, MobiTagbot detects if multipath is likely prevalent at a given robot reader location. If so, MobiTagbot excludes phase readings from that reader location, and generates a final location estimate using phase readings from other locations as the robot reader moves in space. Experimentally, we demonstrate that cutting-edge localization algorithms including Tagoram are not accurate enough to exactly order items in very close proximity, but MobiTagbot is, achieving nearly 100% ordering accuracy for items at low (3–6 cm) spacings and 86% accuracy for items at very low (1–3 cm) spacings.

Keywords

Order tracking; RFID; Localization; Multipath propagation

1. INTRODUCTION

Radio-based positioning of objects equipped with RFID tags is poised to play a key role in the Internet of Things for libraries, retail business, and next-generation manufacturing:

1. A library patron is frustrated in her search for a misshelved book in her local library. The library has RFID-tagged its stock of books for inventory control and loan tracking and placed RFID




Figure 1— MobiTagbot roams a building floor, using its RFID antennas to scan objects such as these file folders on a bookshelf, checking that their spatial order is correct.

readers at the circulation desk, but not in the stacks. Could a RFID-based autonomous robot help locate the misshelved book, or continuously roam the library looking for misshelved books, alerting shelving clerks to the location of any that are found?

2. Supermarkets need to rotate stock as new stock arrives in order to maintain freshness for the consumer. This is highly labor-intensive, requiring the worker to visually inspect the sell-by date of each piece of stock, manually rotate the stock, and then visually check for any old stock out of place. Could an RFID-based autonomous robot roam the supermarket, scan the shelves for stock that needs rotating, and generate a report for the worker to action, saving worker time and reducing labor costs?
3. Manufacturing lines and supply chains already rely on passive RFID tags to track inventory as it moves through the manufacturing and supply chains to the consumer, benefiting from its non-line-of-sight functionality. This allows manufacturers to determine which parts recently passed by an RFID reader, but manufacturers must fall back on other mechanisms such as optical barcodes or physical separation to physically order parts in an assembly line, even if RFID tagged. Could the same technology used above in the autonomous robot be leveraged for a stationary RFID reader to order parts moving by on the assembly line, without the need for physical separation or optical barcode scanning?

Inspired by these and other applications of precision tracking and location for the Internet of Things, recent research efforts have leveraged many different techniques to improve the accuracy of RFID tag localization. To date, the fundamental challenge for the accuracy of indoor localization has been multipath reflections from

nearby objects: these reflections are especially strong and prevalent indoors. Pinit [34] addresses the multipath problem by deploying 200 manually-placed “reference tags,” at *a priori*-known locations, to locate an RFID-tagged book in a $6 \times 5 \times 2.2$ m³ space. RFIDraw [35] uses grating lobes to very accurately trace the trajectory of an RFID tag over small distances (tens of centimeters) at short range (2–5 meters), skirting the problem of multipath. Tagoram [41] uses synthetic aperture radar techniques for a similar result, suffering in the presence of multipath reflections. Older approaches, such as LibBot [11], rely on a “radio map” as first proposed by the classical RADAR [8] system, with the associated difficulties of radio map construction and maintenance.

From the motivating applications mentioned above, we observe and articulate a subtly-new problem: to precisely ascertain the spatial order of closely-spaced RFID tags spread out over a large area (*e.g.*, an entire library floor filled with book stacks), rather than locate them in absolute space with precision, or trace their fine-grained motion. This motivates an approach where an autonomous robot, whose location is only imprecisely known (down to decimeters of precision), scans an area of interest with a mobile RFID reader.

In this paper we present the design and implementation of MobiTagbot, a robot-based RFID ordering system that roams a large area (*i.e.* an entire building floor), ascertaining the spatial order of nearby RFID-tagged objects such as the file folders shown in Figure 1 without the need for a tediously-constructed radio map. Like prior systems such as Tagoram [41], our robot moves its RFID reader over small distances in a straight line, taking readings every few centimeters and forming a *synthetic aperture* radar system, essentially mimicking a radar system with multiple antennas separated by the distance between each reading. Through Wi-Fi, MobiTagbot maintains contact with a backend server cataloguing the expected order of RFID tags. When MobiTagbot detects a mis-ordered pair of tags, it flags the location of the tags and their identities to the backend server, which alerts the user.

MobiTagbot’s contributions are inter-related. First, we observe a problem endemic to phased array processing: the appearance of sidelobes in the generated *holograms* (location likelihood heatmaps) when a single frequency is used. Sidelobes appear because of ambiguities in the response of the antenna array to a tag at a certain distance. To address this problem, we leverage multiple carrier frequencies (§4.1) and combine them into a single hologram. In the absence or near-absence of strong multipath reflections, the differing wavelengths resulting from the use of multiple carrier frequencies reinforce at the true tag location and cancel the sidelobes as shown in Figure 5.

Next, we observe that Tagoram and similar systems run into practical problems indoors where multipath reflections are prevalent and strong: in these situations, their accuracy suffers. To mitigate this problem, MobiTagbot introduces a multipath suppression technique made possible by its use of multiple frequencies. At each location in the synthetic aperture, MobiTagbot sweeps through frequencies, measuring the RFID tag’s phase response as it moves. This forms the basis of a technique described in §4.2 that tests for the presence of multipath reflections. If it detects the presence of multipath in a particular location, MobiTagbot eliminates that location from the synthetic aperture forming a location estimate of the tag using the remaining locations. Together these two techniques significantly improves tag localization: qualitative results are shown in Figure 2.

As a third, systems-building contribution, we implement MobiTagbot using a combination of off-the-shelf hardware: an Imp-




Figure 2— Qualitative comparison of holograms generated by Tagoram (*left*) and MobiTagbot (*right*).

inJ RFID reader equipped with a UHF-band directional antenna sitting atop an iRobot Create 2 robot. Backend custom software written in Java collects phase readings from the RFID reader, location displacement estimates from the iRobot, and processes the data to yield our experimental results. Further details are available in §5.

Our experimental evaluation (§6) comprehensively measures MobiTagbot’s end-to-end performance at ordering RFID tags in two different locations of a real-world office environment: one with significant nearby clutter, resulting in a multipath propagation environment (§6.2), and the other with less nearby clutter, resulting in less multipath (§6.3). Further experiments drill down into these results, measuring basic localization accuracy, the effect of varying the reader-to-tag distance and the tag read rate, and the effect of distance between MobiTagbot and the RFID tags. Microbenchmarks (§6.4) then further explain our performance improvements, measuring the effect of varying the number of frequencies at which MobiTagbot interrogates the RFID tags, the distance MobiTagbot moves, and the spacing between RFID tags. Results show that in multipath-prevalent environments, MobiTagbot achieves nearly 100% ordering accuracy for items at low (3–6 cm) spacings and 86% accuracy for items at very low (1–3 cm) spacings. Performance degrades only slightly up to 1.2 meters away from the tags. Our multipath suppression algorithm is responsible for much of the performance improvements over prior work, almost halving the uncertainty in simple localization.

We next (§2) present the related work and a holography primer (§3) to provide a context for our design (§4). Our implementation (§5) and experimental evaluation follow (§6). We discuss future work in §7, and conclude in §8.

2. RELATED WORK

MobiTagbot draws on prior works in multiple areas, chiefly indoor localization, robot navigation, and Simultaneous Localization and Mapping (SLAM).

2.1 RSSI-based localization

The received Signal Strength Indicator (RSSI) is an indication of the power level received by the radio receiver. Many RSSI-based localization techniques have been proposed during the past, which generally fall into two categories.

RF Fingerprinting. This idea, pioneered by Radar [8], employs RSSI based fingerprinting matching against a database to determine the location of interest indoors. Horus [44] introduces a probabilistic fingerprint matching scheme and achieves a much higher localization accuracy. LANDMARC [22] introduces this technique to RFID localization. SurroundSense [7] further extends this idea by introducing a multi-source (*e.g.*, ambient sound, light, color, *etc.*) based fingerprinting matching scheme for human localization.

Later on several other improvements over RSSI fingerprinting have been proposed, such as incorporating inertial sensor hints [24, 33], embracing fine-grained channel state information [26, 36], and extending to outdoor scenarios [21]. But RF fingerprinting based techniques suffer from ambient dynamics, hence failing to provide fine-grained location estimates.

Modeling instead of fingerprinting. An alternative to fingerprinting based method is to use RF propagation model to estimate the distance from the transmitter to multiple receivers, and estimating the target’s location. In [18], the authors develop a zero-configuration localization system, where a Line-of-sight (LOS) path is assumed and a signal propagation model is used for ranging. More recent work makes an effort to use modeling with minimal assumptions. EZ [9] requires site surveys at only a few user locations. Whitehouse *et al.* [37] extend the RF propagation model to node localization in wireless sensor networks and achieve a localization error of 4.1 m in a 48-node network. But RF signal modeling based techniques suffer from prevalent multipath indoors due to the superposition of signals that propagate along all paths, making highly accurate localization difficult.

2.2 Phase-based localization

Phase reflects the distance that a wireless signal traverses in the physical world. There is a growing interest in using phase measurement for localization:

Distance ranging. The first category of works maps the distance between transmitter and receiver based on the received phase measurements. Liu *et al.* [19] introduces an RFID localization scheme by using multiple antennas to acquire phase measurements from tags. The hyperbolic positioning method is then employed to correlate the phase reading and the tag location. Li *et al.* [17] propose a multi-frequency based ranging method for passive RFID tag localization. Similarly, Zhou *et al.* [45] design a hybrid ranging system by leveraging dual-frequency continuous-wave (DFCW) and continuous-wave (CW) for passive RFID localization. It achieves cm-level localization accuracy using dedicated 5.8 GHz RFID tag. Using phase measurement for distance ranging, theoretically, could achieve high localization accuracy. However, due to the multipath, the phase measurement of an incoming signal is not corresponding to the direct path, and hence will lead to high ranging error.

Holography imaging. Holography is used for capturing the tomography of an object by the vision community. Later on it is introduced to both the radar and acoustic community for target localization [12, 31, 43]. Miesen *et al.* [20] employ holography to locate a moving tag on a transponder. It achieves an overall accuracy of 7 cm. Parr *et al.* [23] exploit tag mobility and adopt Inverse Synthetic Apertures Radar (ISAR) to generate hologram for tag localization and tracking. Tagoram [41] improves the hologram-based localization accuracy to around one centimeter by considering the impact of RFID hardware thermal noise on phase readings. However, it fails to address the multipath issue, hence will likely experience practical problems indoors where multipath reflections are prevalent and strong. MobiTagbot is inspired by above works in phase-based tag localization, but advances them in two ways. First, we propose a more robust holography method by leveraging multiple channels and combining them into a single hologram. Second, we design a novel multipath suppression algorithm, which can detect and find those multipath-prevalent reader locations, and exclude phase readings from these locations for hologram generation.

Angle-of-Arrival (AoA). Wong *et al.* [38] present measurements at high SNR and use AoA and channel impulse response measurements for localization. ArrayTrack [39] adds a novel multipath suppression algorithm, the first system to achieve sub-meter accuracy




Figure 3— The geometry of synthetic-aperture array formed by the movement of MobiTagbot.

in a multipath-rich environment. PinPoint [15] introduces a novel AoA estimation algorithm to localize interfering radios and simultaneously identify the type of interference indoors. SpotFi [16] provides accurate indoor localization services using COTS WiFi NICs with three antennas. It achieves a median accuracy of 40 cm in multipath rich environment. Similarly, WiDraw [30] implements an AoA-based hand gesture recognition system using COTS WiFi NICs. These schemes, however, only provide decimeter localization accuracy for human localization, hence failing to correctly ordering tags to high precision.

AoA information is also employed specifically for RFID localization. PinIt [34] uses a moving antenna to measure the multipath profile of the targeting tag, and compare this profile with reference tags for localization. The overall localization error of PinIt is above 10 cm, hence still not satisfying MobiTagbot’s precise tracking and localization requirement. Rf-IDraw [35] uses grating lobes to very accurately trace the trajectory of an RFID tag over small distances (tens of centimeters) at short range, skirting the problem of multipath.


2.3 Simultaneous Localization and Mapping

Simultaneous Localization and Mapping (SLAM) jointly localizes a robot in an unknown environment, building a consistent map of the environment using sensors such as lidar and cameras [10, 14, 29]. With the proliferation of inertial sensors on smart phones, this problem is extended to human indoor localization. Walkie-Markie builds an indoor floorplan using multiple users’ data [28]. WiFiSLAM [13] relates RSS fingerprints and human movement (*e.g.*, distance, direction) pattern by a Gaussian process latent variable model. When a small part of an indoor space is tagged with RSS measurements, WiFiSLAM performs a semi-supervised localization technique to estimate another user’s location based on RSSI dissimilarity. LiFS [42] further extends SLAM to indoor radio map construction.


MobiTagbot differs from SLAM-related works in two aspects. First, MobiTagbot targets fine-grained object ordering by leveraging antenna mobility, whereas the above works focus on localizing the robot itself and drawing the indoor floorplan. Second, MobiTagbot’s core technology is its multipath suppression algorithm, which allows it to achieve a localization accuracy of about 3 cm in a real indoor environment, whereas the above works significantly suffer in the presence of multipath indoors, achieving an overall localization accuracy of above one meter.

3. HOLOGRAPHY PRIMER

In this section, we review the basic mechanism of the *holography* technique [31, 32] in preparation for our discussion of MobiTagbot design in the next section.



(a)— Traditional holography generation uses a large number of positions to interrogate tags on a single frequency.



(b)— Multi-frequency based holography generation uses a smaller number of positions to interrogate tags on multiple frequencies.

Figure 4— Holography generation technique comparison.

The basic idea behind holography is to correlate an echo profile measured by multiple radar units with a reference echo profile generated at every possible location. The correlation reaches a maximum value if the target is actually located in the tested location. Conversely, the correlation will be low for a false assumption of the target position.

There are two steps involving in using holography for tag localization. The first step is to measure and store the backscattered signals at each possible location. The second step comprises hologram generation and tag localization. We derive the mathematical expression of the hologram by referring to the geometry shown in Figure 3. In this figure, a 2D image grid of points is used to characterize the possible location of a tag. The signal that bounces off grid location $g_{i,j}$ and then arrives at antenna location A_k can be written as:

$$s_r(A_k, g_{i,j}) = a_r(A_k, g_{i,j}) e^{j\theta_r(A_k, g_{i,j})} \quad (1)$$

where $a_r(A_k, g_{i,j})$ and $\theta_r(A_k, g_{i,j})$ are the amplitude and the phase of the backscatter signal, respectively. If the reader antenna moves to M known locations, we can compute M signals for grid location $g_{i,j}$. These M signals together comprise a *reference profile* for the grid location $g_{i,j}$ and are stored.

We use an *inverse correlation filter* to correlate two signals for hologram generation. Let $s_t(A_k)$ be the signal backscattered from the tag and received by the antenna A_k . The inverse correlation filter correlates $s_t(A_k)$ with the reference signal $s_r(A_k, g_{i,j})$ using the following formula:

$$C(s_t(A_k), s_r(A_k, g_{i,j})) = e^{j\theta_t(A_k)} \cdot e^{-j\theta_r(A_k, g_{i,j})} \quad (2)$$

In this equation,¹ $\theta_t(A_k)$ is the received phase reading by antenna A_k , which can be directly acquired from the COTS RFID reader. For each grid location, above correlation is then done for all M antennas and finally summed up. The pixel value $I(g_{i,j})$ of grid location $g_{i,j}$ in the hologram is then computed as follows:

$$I(g_{i,j}) = F \left| \sum_{k=1}^M C(s_t(A_k), s_r(A_k, g_{i,j})) \right| \quad (3)$$

¹The amplitude of the signal $s_t(A_k)$ and $s_r(A_k, g_{i,j})$ is simply set to one, since the COTS RFID reader reports the amplitude as a superposition of radio signals off nearby objects, which fails to correctly reflect the amplitude changing of LOS path.

where F denotes an appropriate normalization factor. The interpretation of Equation 3 is straightforward. If the pixel position is the actual tag position and there is no multipath, the phase measurement of the received signals should equal to that of the reference signals and the resultant correlation coefficient (pixel value) is maximized. Conversely, if the tag is not located in grid $g_{i,j}$, the complex signals will be summed up with random phase, which causes destructive interference and, thus, results in a small value of $I(g_{i,j})$. To locate the tag, the hypothesis is tested for every grid on the hologram. In the absence of multipath (which we address in §4), the grid yielding highest $I(g_{i,j})$ should be the true location of the tag.

The advantage of holography lies in its probabilistic approach using signal correlation to compute the likelihood of each location to be the tag's true location. However, the granularity of the localization result is highly sensitive to the number of antennas used in hologram generation. The more signals captured at different antenna locations, the higher the accuracy of the hologram. Hence to capture the precise location of a tag, bulky antenna arrays would have to be used. This is due to the phase ambiguity problem: phase changes periodically from zero to 2π . As a result, a signal bouncing at two different grid locations may yield equivalent phase readings. To uniquely label each grid location with phase readings, it is thus necessary to collect readings from many spatially-separated antennas. Also, since holography is based on an assumption that multipath is weak, accuracy degrades significantly when multipath is prevalent.

4. DESIGN

In this section, we present the design of MobiTagbot's multi-frequency holography and multipath suppression techniques.

4.1 Multi-frequency based holography

In practice, it is difficult or even infeasible to deploy expensive antenna arrays for low-cost tag ordering. MobiTagbot moves, creating a number of spatially-separated *virtual antennas*, which potentially improves its localization accuracy. However, the challenge for a wheeled robot lies in knowing its precise antenna position as it moves. If these robot position estimations are inaccurate, adding up acquired signals will cause destructive interference at the true location, further degrading tag localization accuracy.

In order to minimize this adverse effect, we propose a multi-frequency based holography. The basic idea is to fully exploit the channel hopping capability of COTS RFID devices by changing the carrier frequency (channel) of the RFID query. Our new inverse correlation filter becomes:

$$I(g_{i,j}) = F' \left| \sum_{k=1}^M \sum_{l=1}^N C(s_t(A_k, f_l), s_r(A_k, f_l, g_{i,j})) \right| \quad (4)$$

where F' denotes an appropriate normalization factor. In this equation, N is the number of channels we use for RFID query; $s_t(A_k, f_l)$ represents the signal on l^{th} channel received by the antenna A_k . Figure 4 shows the basic principle of traditional holography and our multi-frequency based holography. Compared with traditional holography, the advantage of MobiTagbot's multi-frequency based holography is twofold. First, we reduce the large number of positions needed, hence minimizing the adverse effect of location uncertainty associated with a moving antenna. Second, we create many more virtual antennas via channel hopping, so the combined result of signals received from these virtual antennas reinforces at the true tag location and cancels the side lobes (shown qualitatively in Figure 5).




Figure 5— The hologram generated by using different numbers of carrier frequencies. The ground truth location of this tag is (25, 40). Spurious side lobes are weakened when phase readings from multiple frequencies are incorporated into the hologram.

Figure 5 shows five pictures of holograms generated by using signals received on different numbers of carrier frequencies. As the figure shows, multiple side lobes appear in the hologram generated by the phase reading on a single carrier frequency, which renders the true location of the tag unclear. The likelihood of the tag being falsely located on these side lobes decreases significantly when using phase data from four, eight, and 12 frequencies. This is as expected since transmitting at a single location but over N different frequencies has a similar effect as transmitting a signal at N different locations but over a single frequency. The granularity improves as more channels are used. We also observe that the magnitude of the location accuracy improvement decreases as we further add channels.

4.2 Multipath suppression

A wireless signal returning from the tag emanates in all radial directions, reflecting off nearby walls, tables and other objects, leading to multipath propagation. The phase measurement reported by the RFID reader is the superposition of signals on all propagation paths, not necessarily just the direct path. Hence the localization accuracy achieved by holography will degrade significantly in prevalent multipath scenarios. While the multi-frequency based holography described above reduces moving antenna-induced distortion of phase measurements, it alone does not handle the adverse effect of multipath, hence multipath reflections are free to reduce system accuracy.

But the strength of multipath changes in space [25, 40], and so the basic idea of MobiTagbot’s multipath suppression algorithm is to detect and find multipath-prevalent reader locations, excluding phase readings from these locations for hologram generation. This algorithm is based on the following observation—In a low-multipath reader location, the RFID reader’s phase measurement corresponds to the signal that propagates along just the direct path. Hence the phase measurement change is constant when MobiTagbot changes the carrier frequency. In contrast, if the wireless signal suffers from severe multipath at a given reader location, the phase measurement should be the superposition of phase values corresponding to many propagation paths, and hence change irregularly (due to nonlinear superposition of phase values) when we change the carrier frequency.

We conduct a micro-benchmark to validate this observation. We program an RFID reader to hop over 16 carrier frequencies. The reader is placed one meter away from the tags. We conduct this micro-benchmark in both a low and prevalent multipath scenario (shown in Figure 6). We choose two tags and plot their phase series in Figure 6. In the low multipath scenario, for both these tags, the phase series stay around a stable value within each carrier frequency, and change linearly when hopping from one carrier frequency to another. In contrast, in prevalent multipath scenario, the phase series of these two tags change irregularly. Specifically, the phase value jumps abruptly due to human motion-induced dynamic




Figure 6— Comparison of phase profiles in low and prevalent multipath scenario.

multipath, and changes non-linearly during the channel hopping due to the superposition of phase from all paths.

Algorithm. Suppose MobiTagbot is programmed to hop over N channels. On each channel, the reader captures multiple phase measurements. Let θ_l be the average phase measurement on l^{th} channel. The phase difference between two consecutive channels is: $\Delta\theta_l = \theta_{l+1} - \theta_l$ (there are $N - 1$ phase difference values in total). We use entropy to measure the similarity of average phase differences across locations: we define $p_l = \frac{\Delta\theta_l}{\sum_{l=1}^{N-1} \Delta\theta_l}$, so we can write the entropy of $N - 1$ phase differences at the k^{th} reader location as:

$$e_k = -\frac{1}{\ln(N-1)} \cdot \sum_{l=1}^{N-1} p_l \lg p_l \quad (5)$$

In this equation, if all phase difference values are similar, e_k will be close to one. Otherwise, e_k will be close to zero. We further define the following weighting function and our new inverse correlation filter as follows:

$$w_k = \frac{e_k}{M - \sum_{i=1}^M e_i} \quad (6)$$

$$I(g_{i,j}) = F' \left| \sum_{k=1}^M w_k \sum_{l=1}^N C(s_t(A_k, f_l), s_r(A_k, f_l, g_{i,j})) \right|$$

where M is the number of reader locations (motion-induced virtual antennas). The interpretation of this weighting function is straightforward: a higher entropy of phase differences at location k indicates a lower likelihood of prevalent multipath, and hence phase measurements reported at this location should have a higher weight for the final hologram generation. A lower entropy of phase differences at location k indicates a higher likelihood of prevalent multipath, hence phase measurements reported at this location should have a lower weight for final hologram generation.

Figure 7 shows the hologram generated without and with MobiTagbot’s multipath suppression. The side lobes are weakened significantly when the multipath suppression algorithm is used. Also, the beam width of the main lobe is narrowed down to 3 cm after multipath suppression.




Figure 7— Hologram generated without (*left*) and with (*right*) multipath suppression.




Figure 8— Experimental setups: MobiTagbot hardware (*left*); multipath-prevalent scenario (*center*); multipath-low scenario (*right*).

5. IMPLEMENTATION

So far we have provided the algorithmic basis of how MobiTagbot works. This section describes the specific implementation of MobiTagbot.

5.1 Frontend hardware

MobiTagbot’s hardware consists of two parts: an RF communication module and a motion controller. All these hardware components are off-the-shelf devices.

An ImpinJ Speedway R420 RFID reader [2], two RFMAX S902-8PCR directional antennas [5], and a set of Avery Dennison AD-227m5 UHF passive RFID tags [6] comprise the RF communication module. The R420 reader works at the UHF frequency band 920.625–924.375 MHz and is programmed to hop over 16 channels. The S9028PCR antenna has a 9dBi gain, and 70° elevation and azimuth beam width.

The motion controller is based on an iRobot-Creat 2 robot [3]. We hack it using the API provided by [1]. As Figure 8 shows, the robot serves as the chassis of MobiTagbot. The body of MobiTagbot is composed of lightweight support frames, which mount the RFID reader and antenna. In our experiment, we program MobiTagbot to move at the speed of 5 cm/s for tag ordering.

5.2 Backend Software

We now briefly discuss the back-end processing of MobiTagbot. The three main tasks of the back-end are to control the motion of MobiTagbot; manipulate the RFID reader to query tags; run the algorithms that are described in §4. The software of MobiTagbot is implemented in Java, and runs on a Lenovo ThinkCentre PC equipped with Intel(R) i5-4590 CPU and 4 GB RAM. It collects phase readings via the Low-level Reader Protocol (LLRP) [4].

Speed calibration. Figure 9 shows the CDF of the moving distance error when MobiTagbot moves on different kinds of floors. As the result indicates, error varies significantly when MobiTagbot moves on different kinds of floors, because the friction force of different material varies significantly. Hence the moving speed of MobiTagbot may deviate from its speed setting due to the change of floor




Figure 9— CDF of the moving distance error of MobiTagbot when it moves on two kinds of floors. The groundtruth distance is 100 cm.

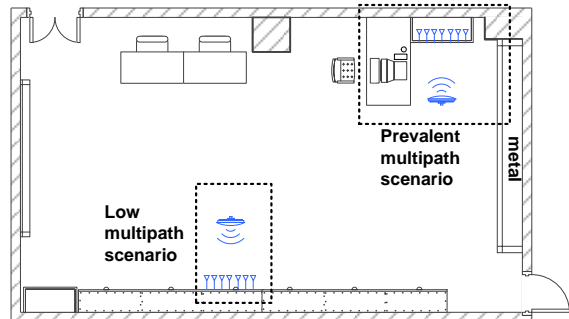


Figure 10— Floorplan of our experimental lab locations.

material. We also find that the span of each CDF is small, indicating a relatively stable distance error for each kind of floors. Based on this finding, we conduct a quick speed calibration before each experiment. Let v be the robot’s speed set in program, and v' be the real speed of the robot. We have $v = v' + \Delta v$, where Δv is the speed drift caused by the friction force. Before the experiment, we let the robot move t s at the speed v . So the expected moving distance d is $t \cdot v$. We measure the robot’s real moving distance (represent as d'), and compute the speed drift as: $\Delta v = \frac{d' - d}{t}$. The moving speed of MobiTagbot is finally calibrated by the equation $v - \frac{d' - d}{t}$.

6. EVALUATION


In this section, we first conduct field experiments to compare the ordering performance of MobiTagbot and other two state-of-the-arts. After that, we present micro-benchmark experiments to provide insights into MobiTagbot, particularly to explain which factors impact MobiTagbot’s performance. Since the spacing between neighboring bookshelves is large, it can be more easily distinguished by prior efforts in time of arrival localization. Hence, in our experiments we focus on the more challenging part: ordering accuracy along one dimension where books closely stand by each other.

6.1 Experimental methodology

Unless the floor material is specified, experiments are conducted on wooden surface. We conduct field experiments to measure the end-to-end performance of MobiTagbot in two different locations of a real-world office environment. Figure 10 shows the floorplan of these two locations. In the first location, multiple objects are placed around RFID tagged objects, in order to generate a multipath-prevalent environment. In the second location, we move the RFID tagged objects to a relatively open space where only walls and tables generate a low amount of multipath. Figure 8 gives a snapshot of these two scenarios. In each scenario, we attach 20 passive RFID tags on the books’ spines. The CDF of the spacing between adjacent tags is shown in Figure 11.

Table 1— Review of state-of-the-art RFID localization schemes.

Localization scheme	Antennas required	Method	Hardware	Applicability
PinIt [SIGICOMM’13]	One moving antenna	Proximity-based	SDR	Prevalent multipath
RF-compass [MobiCom’13]	One moving antenna	Geometry-based	SDR	Low multipath
RF-IDraw [SIGICOMM’14]	Eight antennas	Ranging-based	COTS RFID	Low multipath
Tagoram [MobiCom’14]	Four antennas	Hologram-based	COTS RFID	Low multipath
STPP [NSDI’15]	One moving antenna	Proximity-based	COTS RFID	Low multipath
MobiTagbot (our proposal)	Two moving antennas	Hologram-based	COTS RFID	Prevalent/low multipath

**Figure 11**— CDF of the distance between two adjacent tags attached on books in our experiment.**Table 2**— Fraction of correct book order identification in a multipath-prevalent environment.

System	Group 1	Group 2	Group 3	Group 4
STPP	46%	52%	38%	40%
Tagoram	74	66	74	72
MobiTagbot	86	84	86	88

Metrics. In the evaluation of MobiTagbot, we mainly focus on the following two metrics: *ordering accuracy*, and *localization error*. The ordering accuracy is defined as the number of tags ordered correctly out of total tag number. The localization error is computed as the distance between the estimated location and the true location of this tag. In each experimental scenario, we first show how MobiTagbot performs tag ordering, and then we drill down into these results, measuring basic localization accuracy, the effect of varying the reader-to-tag distance and the tag read rate.

Comparisons. There are many notable research proposals in RFID based indoor localization. Table 1 summarizes the key features of several latest cutting-edge works. We highlight MobiTagbot in the last row of this table. As the table shows, PinIt and RF-Compass both rely on software defined radio for tag localization, while RF-IDraw, Tagoram, STPP and MobiTagbot are built upon COTS RFID devices. RF-IDraw is designed for tracking human finger movement in a fine-grained manner. It achieves cm-level accuracy based on eight spatially-separated antennas. Similarly, Tagoram achieves very high tracking accuracy by using multiple antennas and synthetic aperture radar techniques. STPP adopts a moving antenna to directly capture the tag order in 2D space, with an error margin of 8 cm. These three schemes, however, all require the antenna to communicate with tags over a short distance, so that the reported phase readings correspond to the signals propagate along the direct path.

We implement Tagoram [41] (DAH algorithm) and STPP [27], two state-of-the-art works that use the same hardware as MobiTagbot, for comparison. Tagoram employs multiple pre-calibrated an-

Table 3— Fraction of correct book order identification (plus or minus one book) in a multipath-prevalent environment.

System	Group 1	Group 2	Group 3	Group 4
STPP	50%	56%	42%	44%
Tagoram	88	90	92	90
MobiTagbot	100	98	100	100

tennas to generate a hologram for real-time tag tracking. It minimizes the impact of thermal noise on received phase readings and relaxes the computational overhead for hologram generation. STPP employs a moving antenna to generate a phase profile for each tag, exploring the spatial-temporal correlation of phase profiles to directly estimate the tag order.

6.2 Field experiment: Prevalent multipath

This section presents results from the “multipath-prevalent” scenario in Figure 8.

6.2.1 Overall ordering accuracy

We first look at how each algorithm performs in book ordering. We put the reader antenna one meter away from the bookshelf and conduct four groups of experiments to measure the variance across experimental trials. In each group, we repeat the experiment five times and get 100 orders of books in total. The success rates are shown in Table 2. As we can see, MobiTagbot identified the right order of books with a minimum success rate of 84%. In contrast, STPP and Tagoram achieve a success rate of at most 52% and 74%, respectively. This clearly demonstrates that MobiTagbot achieves significantly higher ordering accuracy than these schemes.

We then look at how each algorithm works in ordering the books with an imprecision of up to one book. As Table 3 shows, there is no significant ordering accuracy improvement in STPP. Tagoram’s performance increases a slightly from 72% on average to 90% on average, while MobiTagbot achieves almost 100% accuracy for all the four test groups.

6.2.2 Overall localization accuracy

To better understand why MobiTagbot outperforms other two algorithms in book ordering, we then analyze the localization accuracy of these three algorithms. Figure 12 shows the Cumulative Distributed Function (CDF) of the error distances in MobiTagbot, Tagoram and STPP. STPP achieves a median error distance of 8.6 cm and a 90th percentile of 11.2 cm. The maximum error margin of STPP reaches 14.2 cm. This large error margin comes from both thermal noise and severe multipath effect indoors. Tagoram outperforms STPP. It achieves a median error margin of 5.6 cm and a 90th percentile of 6.8 cm in a severe multipath environment. The maximum error distance is 7.6 cm. Tagoram achieves higher accuracy than STPP since it is based on a more robust localization scheme and it carefully handles the thermal noises introduced by both the reader antenna and the tag’s hardware. However, due to




Figure 12— CDF of localization error margin achieved by three algorithms in a multipath-prevalent environment.




Figure 13— Impact of the distance from the reader. The error bar specifies the maximum, minimum and average localization error under different distance settings.

the following reason, it fails to achieve comparable accuracy as the previous result reported in [41]. Reflectors such as tables and books in the experimental environment create a complex multipath environment, undermining the correlation between location and phase readings. Tagoram did not explicitly handle the phase drifts introduced by the multipath effect. Hence it runs into practical problems indoors where multipath reflections are prevalent and strong.

MobiTagbot achieves a median error margin of 3.4 cm, outperforming both STPP and Tagoram by 2.5 \times and 1.6 \times , respectively. The 90th percentile and the maximum error margin is 4.8 cm and 5.6 cm, respectively. The improvement is due to MobiTagbot’s ability to handle the distortion introduced by the multipath effect. Compared with Tagoram, MobiTagbot creates more virtual antennas by using multiple frequencies. The hologram combination of these virtual antennas reinforces the true location and removes false positive locations. Combined with the tag ordering accuracy result, we can see that 2.2 cm localization accuracy gain results in a significant ordering accuracy gain when applied to the application of book ordering.

6.2.3 Varying the reader-to-tag distance

We now show how the three localization algorithms perform under different tag-to-antenna distance settings. In these experiments, we vary the reader-to-tag distance from 0.4 m to 1.2 m, with a step of 0.2 m. In each distance setting, we collect phase readings at 10 locations. Figure 13 shows the error distance of each localization scheme. As the result indicates, the localization error of all these three schemes decreases steadily as we move the reader closer to the tags. This is because the short tag-to-reader distance results in relatively weaker reflections.

Although Tagoram shows similar error margin with MobiTagbot under short tag-to-reader distance setting (*e.g.*, 0.4 m), its perfor-

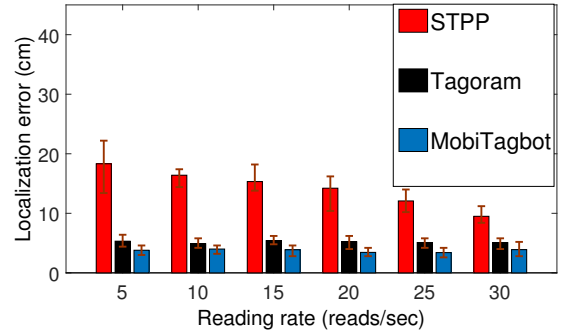


Figure 14— Impact of the tag reading rate. The error bar specifies the maximum, minimum and average localization error under different reading rate settings.

Table 4— Fraction of correct book order identification in a low-multipath environment.

System	Group 1	Group 2	Group 3	Group 4
STPP	52%	48%	54%	42%
Tagoram	82	84	84	86
MobiTagbot	88	88	90	86

mance degrades significantly when we gradually move the antenna away from tags. In contrast, MobiTagbot performs consistently robust under the five distance settings. This clearly demonstrates that MobiTagbot can successfully minimize the impact of multipath effect.

6.2.4 Varying tag reading rate

We now show how these three schemes perform under different tag reading rate settings. In this series of experiments, we put the RFID reader 1 m away from tags. Since the reader uses the ALOHA protocol to interrogate the tag, each tag will be interrogated a similar number of times. We thus fix the reader’s sampling rate, and manually increase or decrease the number of tags within the reader’s interrogation range to change the tag reading rate. Figure 14 shows the error margin achieved by different algorithms with respect to different reading rates.

As the result indicates, MobiTagbot consistently outperforms other algorithms in all of the six reading rate settings. From this figure we can see that STPP’s performance is highly sensitive to the tag’s reading rate. Specifically, STPP achieves a localization error of 9.8 cm on average, when the tag reading rate is set to around 30 reading/second. It then increases significantly and reaches over 18 cm on average when we set the reading rate to 5 reading/second. Because the temporal phase profile becomes extremely sparse when the tag reading rate is low. Hence, it fails to correctly reflect the spatial relationship between tags.

6.3 Field experiment: Low multipath

This section presents results from the “Low multipath” scenario in Figure 8.

6.3.1 Overall ordering accuracy

We start by presenting the overall ordering accuracy of the three algorithms. Similar to the setup in multipath-prevalent scenario, we put the reader antenna 1 m away from books and conduct four groups of experiments. Table 4 summarizes the success rate of each algorithm in book ordering. As the result shows, STPP still achieves a success rate around 50%, much similar to the perfor-

Table 5— Fraction of correct book order identification (plus or minus one book) in a low-multipath environment.

System	Group 1	Group 2	Group 3	Group 4
STPP	54%	56%	60%	54%
Tagoram	94	94	90	92
MobiTagbot	100	100	100	100




Figure 15— CDF of localization error margin achieved by three algorithms in a low-multipath environment.

mance in multipath-prevalent scenario. In contrast, Tagoram’s performance increases remarkably and reaches to a level that is comparable to MobiTagbot. This is expected as the backscatter signal suffers from slight multipath effect in this scenario.

Similarly, we also look at how each algorithm works in ordering the books with one book in-between. As Table 5 shows, there is a small improvement of STPP. Tagoram’s performance improves from 84% on average to above 92% on average, while MobiTagbot correctly identifies the order of all books in these four groups of experiments.

6.3.2 Overall localization accuracy

To understand why there is no significant ordering accuracy improvement in STPP, we present the overall localization results of these three localization schemes. The experiment setups are the same as the settings in prevalent multipath scenario. Figure 15 shows the CDF of the error distances in MobiTagbot, STPP and Tagoram.

STPP has a median error distance of 7.4 cm and a 90th percentile of 8.8 cm, much lower than it achieves in prevalent multipath environment. This is expected as the phase profile of each tag experiences light multipath effect. Hence it can reliably reflect the spatial-temporal correlation of tags. However, we can see that the median error margin achieved by STPP is still larger than the width of most books. Hence STPP fails to distinguish the correct order of most books. This explains why STPP fails to achieve a significant improvement in book ordering accuracy.

Tagoram outperforms STPP in a low-multipath environment. It achieves a median error distance of 3.8 cm and a 90th percentile of 4.2 cm. This error margin is mainly due to the small phase drift introduced by the multipath effect.

MobiTagbot achieves a median error margin of 2.8 cm, outperforming both STPP and Tagoram by $2.6\times$ and $1.3\times$, respectively. The 90th percentile of MobiTagbot is 3.3cm, which outperforms STPP and Tagoram by $2.6\times$ and $1.2\times$, respectively. These results clearly demonstrate that MobiTagbot can successfully detect and handle the multipath propagation.

6.3.3 Varying the reader-to-tag distance

We next show how these three localization algorithms perform

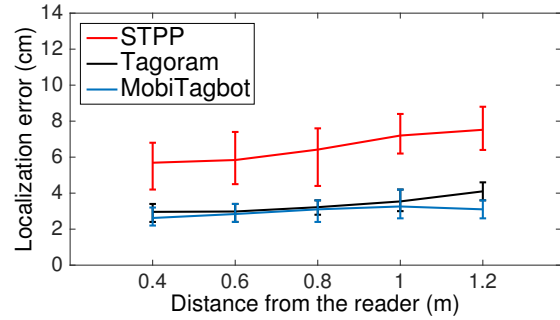


Figure 16— Impact of the distance from the reader. The error bar specifies the maximum, minimum and average localization error under each distance settings.




Figure 17— Impact of the tag reading rate. The error bar here specifies the maximum, minimum and average localization error achieved under different reading rate settings.

under different tag-to-antenna distance settings. As in prevalent multipath scenario, in this trail of experiments we also vary the tag-to-reader distance from 0.4 m to 1.2 m, with a step of 0.2 m. Figure 16 shows the localization error achieved by each localization scheme. We make the following three observations. First, similar to the variation of localization accuracy in prevalent multipath scenario, the localization error of all three schemes decreases steadily with the decrease of the distance between the reader and the tag. Second, for each localization scheme, its performance degradation is not as significant as in the prevalent multipath scenario. This is because the light multipath effect does not change so much on different tag-to-reader distances. Third, as expected, when there is low multipath, Tagoram achieves relatively similar performance to MobiTagbot. Nevertheless, the average localization error achieved by MobiTagbot is still below the performance achieved by Tagoram.

6.3.4 Varying the tag reading rate

We now show how these algorithms perform under different reading rate settings. Similarly to the experiment settings in multipath-prevalent scenario, we put the RFID reader 1 m away from tags in these experiments. Figure 17 shows the error margin achieved by different algorithms with respect to different reading rates. As the result indicates, STPP’s performance degrades significantly when the reading rate drops. On the other hand, as expected, the multipath propagation of signals introduces faint distortion to phase readings in the low multipath scenario. Accordingly, Tagoram achieves a comparable performance to MobiTagbot in this setting.

6.4 Microbenchmarks

We conduct micro benchmarks to give further insight into MobiTagbot’s performance. In particular, to understand which factors




Figure 18— Illustration of micro benchmark scenario.

Table 6— Fraction of correct book order identification (plus or minus one book) with (Case 02) and without (Case 01) the multipath suppression algorithm.

Case	Group 1	Group 2	Group 3	Group 4
Case 01	92%	88%	90%	92%
Case 02	98	100	100	98

will most affect the localization performance of MobiTagbot. Figure 18 shows a view of the benchmark setup. In these experiments, we attach 20 UHF passive tags on the paperboard. The spacing between adjacent tags is 3 cm. The reader is placed 1 m away from the paperboard. Multiple reflectors are then put near the tag to generate a multipath-prevalent environment.

6.4.1 Effect of multipath suppression

We first look at how MobiTagbot works in ordering the books with an imprecision of up to one book. We put the reader antenna 1 m away from the books and conduct four groups of experiment. In each group, we repeat the experiment five times. Table 6 summarizes the ordering accuracy of MobiTagbot before and after multipath suppression. As the result shows, when MobiTagbot does not perform multipath suppression and blindly uses phase readings reported at all reader locations, the accuracy is at most 92%. In contrast, with multipath suppression, the accuracy boosts to almost 100%. This result clearly demonstrates the effectiveness of our multipath suppression algorithm.

We further explore the benefit of our multipath suppression algorithm on the tag localization accuracy. The result is shown in Figure 19. As this figure shows, when MobiTagbot does not execute the multipath suppression algorithm, the median and 90th percentile localization errors are 4.3 cm and 6.4 cm, respectively. With multipath suppression algorithm, the localization error decreases significantly. The median and 90th percentile localization errors are becoming 2.9 cm and 4.1 cm, respectively. These results clearly demonstrate the effectiveness of our multipath suppression algorithm.

6.4.2 Varying the number of channels

Next, we show how MobiTagbot works while using different number of channels. Specifically, we randomly choose 4, 6, 8, 10, 12, 14 channels out of 16 frequencies and use the phase readings on these channels for multipath suppression and hologram generation. The CDF of localization accuracy under each channel setting is shown in Figure 20. As this figure shows, the median accuracy increases significantly with the increased number of channels we use. The reason behind is two-fold. On one hand, MobiTagbot creates more virtual antennas when hopping in more channels. The holo-

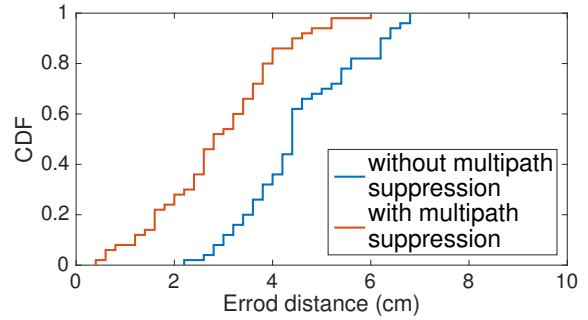


Figure 19— Localization accuracy with/without multipath suppression.




Figure 20— Impact of the number of channel hopping.

gram combination of these virtual antennas naturally reinforces the true location and removes false positive locations. On the other hand, it is more reliable to detect the multipath effect based on a larger amount of channels, which helps to minimize its impact on the localization result further. From the result we also find that when the number of channel we use is large enough (*i.e.*, $n = 10$), and it is unlikely to see further improvement on localization performance as the dominate factor will be correct alignment of moving antennas and even the ranging error for ground-truth.

6.4.3 Varying the number of robot read locations

We then show how MobiTagbot works while using different number of virtual antennas (*i.e.*, number of positions MobiTagbot reads from). In these experiments, we make the reader hop over 16 channels and then capture the tag reading on each channel, respectively. Figure 21 shows the snapshot of the hologram generated by MobiTagbot under different number of virtual antenna settings. In this figure, we divide the X-Z plane into equal size blocks. The block size is set to 1 cm. The value on each block represents the likelihood of tag existing. The ground-truth of the tag is (100,100). As the result indicates, the localization accuracy improves significantly with the increasing number of virtual antennas we use. Specifically, when the number of virtual antennas is large such as five or six, the hologram combination inherently reinforces the true location and removes the false positives. In contrast, when the number of virtual antennas is small, *e.g.*, 2 or 3, the reinforcement is not strong hence MobiTagbot fails to distinguish the true location from a set of location candidates.

6.4.4 Varying the spacing between adjacent tags

We next examine how MobiTagbot performs with different adjacent tag spacings. In these experiments, we vary the tag spacing from 0.5 cm to 5 cm, and plot the localization accuracy in Figure 22. As this figure shows, when the spacing is large, *i.e.*, > 3 cm,




Figure 21— Hologram showing the location likelihood of a tag with differing numbers of virtual antennas computing its location. The ground-truth location is (100,100).




Figure 22— Impact of distance between adjacent tags. The error bar specifies the maximum, minimum and average localization error under different tag spacing settings.

there is no significant performance variation. Besides, the maximum and minimum localization errors retain in a relatively stable and small region. However, as we reduce the spacing between two tags from 2 cm to 1 cm and further to 0.5 cm, we can see a clear performance degradation. This is because when two tags are within extremely short distance, they form two identical circular loops, which will generate inductive coupling effect. Consequently, each tag will be shadowed by another, resulting in reduced signal power at the tag side.

7. DISCUSSION AND FUTURE WORK

We discuss practical deployment issues in this section.

Speed calibration. Calibration of the robot’s speed is essential to ensure accuracy in environments with changing floor surfaces. We have obtained accurate results in our lab environment but plan next to investigate wireless, real-time wheel speed measurement using a small number of calibration tags with known spacings, placed near floor boundaries.

Dynamic multipath suppression. Although our experiments are conducted in the presence of moving humans and are robust to nearby motion, we believe schemes that actively reject dynamic multipath may lead to even better performance. This would involve examining the variance and other statistics of multiple phase readings from a tag taken at the same location and on the same channel.

NLOS tag readings. MobiTagbot requires a clear LOS between the tagged-book and the reader antenna. However, in libraries the RFID reader may interrogate the tagged-book behind the current bookshelf, leading to NLOS tag readings. In future work, we plan to address this issue by measuring the time-of-flight (ToF) of each incoming signal, filtering out tags with significantly different response delays in order to isolate a particular bookshelf.

Impact of tag orientation. The readability of a tag is affected by the relative orientation of this tag and the reader antenna. An orientation mismatch will result in a lower tag reading rate, hence

degrading system performance. We plan to put antennas in different angles such that tags with diverse orientations will be covered.

Extending to tag moving case. MobiTagbot conducts a roving survey of the interested area to achieve an exact spatial order of RFID-tagged objects. Its core technique can be applied to other applications such as robot manipulation, assembly line construction, and baggage sorting in an airport.

8. CONCLUSION

In this paper, we present the design, implementation and evaluation of MobiTagbot, an autonomous wheeled robot reader that conducts a roving survey of libraries, manufacturing lines and offices to achieve an exact spatial order of RFID-tagged objects in very close (1–6 cm) spacings. MobiTagbot advances the cutting-edge RFID localization systems in its ability to detect and minimize the adversary effect of multipath. MobiTagbot is purely built upon COTS RFID devices and manipulated on backend custom software. The real-world based experiments demonstrate that MobiTagbot could achieve nearly 100% ordering accuracy for items at low (3–6 cm) spacings and 86% accuracy for items at very low (1–3 cm) spacings in multipath-prevalent scenario.

Acknowledgements

We sincerely thank our shepherd Prof. Marco Gruteser and the anonymous reviewers for their insightful comments. The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP/2007–2013) / ERC Grant Agreement No. 279976.

References

- [1] Hacking Roomba. Web page.
- [2] ImpinJ Speedway R420 reader. Web page.
- [3] iRobot Create 2. Web page.
- [4] LLRP Toolkit. Web page.
- [5] RFMAX S9028PCR antenna. Web page.
- [6] UHF RFID Inlays AD-227m5. Web page.
- [7] M. Azizyan, I. Constandache, and R. Roy Choudhury. SurroundSense: Mobile phone localization via ambience fingerprinting. In *MobiCom*, 2009.
- [8] P. Bahl and V. N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In *Infocom*, 2000.
- [9] K. Chintalapudi, A. Padmanabha Iyer, and V. N. Padmanabhan. Indoor localization without the pain. In *MobiCom*, 2010.
- [10] E. Eade and T. Drummond. Scalable monocular SLAM. In *CVPR*, 2006.

- [11] I. Ehrenberg, C. Floerkemeier, and S. Sarma. Inventory management with an RFID-equipped mobile robot. In *CASE*, 2007.
- [12] H. Ermert and R. Karg. Multi-frequency acoustical holography. *TOSU*, 1979.
- [13] B. Ferris, D. Fox, and N. D. Lawrence. WiFi-SLAM using Gaussian Process latent variable models. In *IJCAI*, 2007.
- [14] A. Flint, C. Mei, I. Reid, and D. Murray. Growing semantically meaningful models for visual SLAM. In *CVPR*, 2010.
- [15] K. R. Joshi, S. S. Hong, and S. Katti. PinPoint: Localizing interfering radios. In *NSDI*, 2013.
- [16] M. Kotaru, K. Joshi, D. Bharadia, and S. Katti. SpotFi: Decimeter level localization using WiFi. In *SIGCOMM*, 2015.
- [17] X. Li, Y. Zhang, and M. G. Amin. Multifrequency-based range estimation of RFID tags. In *RFID*, 2009.
- [18] H. Lim, L.-C. Kung, J. C. Hou, and H. Luo. Zero-configuration, robust indoor localization: Theory and experimentation. In *Infocom*, 2005.
- [19] T. Liu, L. Yang, Q. Lin, Y. Guo, and Y. Liu. Anchor-free backscatter positioning for RFID tags with high accuracy. In *infocom*, 2014.
- [20] R. Miesen, F. Kirsch, and M. Vossiek. Holographic localization of passive UHF RFID transponders. In *RFID*, 2011.
- [21] A. Musa and J. Eriksson. Tracking unmodified smartphones using Wi-Fi monitors. In *SenSys*, 2012.
- [22] L. M. Ni, Y. Liu, Y. C. Lau, and A. P. Patil. LANDMARC: indoor location sensing using active RFID. *Wireless networks*, 2004.
- [23] A. Parr, R. Miesen, and M. Vossiek. Inverse SAR approach for localization of moving RFID tags. In *RFID*, 2013.
- [24] A. Rai, K. K. Chintalapudi, V. N. Padmanabhan, and R. Sen. Zee: zero-effort crowdsourcing for indoor localization. In *MobiCom*, 2012.
- [25] S. Sen, J. Lee, K.-H. Kim, and P. Congdon. Avoiding multipath to revive inbuilding WiFi localization. In *MobiSys*, 2013.
- [26] S. Sen, B. Radunovic, R. R. Choudhury, and T. Minka. You are facing the Mona Lisa: Spot localization using PHY layer information. In *MobiSys*, 2012.
- [27] L. Shangguan, Z. Yang, A. X. Liu, Z. Zhou, and Y. Liu. Relative localization of RFID tags using spatial-temporal phase profiling. In *NSDI*, 2015.
- [28] G. Shen, Z. Chen, P. Zhang, T. Moscibroda, and Y. Zhang. Walkie-Markie: Indoor pathway mapping made easy. In *NSDI*, 2013.
- [29] H. Strasdat, A. J. Davison, J. Montiel, and K. Konolig. Double window optimisation for constant time visual SLAM. In *ICCV*, 2011.
- [30] L. Sun, S. Sen, D. Koutsonikolas, and K.-H. Kim. WiDraw: Enabling hands-free drawing in the air on commodity WiFi devices. In *MobiCom*, 2015.
- [31] M. Vossiek, V. Magori, and H. Ermert. An ultrasonic multi-element sensor system for position invariant object identification. In *US*, 1994.
- [32] M. Vossiek, A. Urban, S. Max, and P. Gulden. Inverse synthetic aperture secondary radar concept for precise wireless positioning. *TMTT*, 2007.
- [33] H. Wang, S. Sen, A. Elgohary, M. Farid, M. Youssef, and R. R. Choudhury. No need to war-drive: unsupervised indoor localization. In *MobiSys*, 2012.
- [34] J. Wang and D. Katabi. Dude, where's my card?: RFID positioning that works with multipath and non-line of sight. In *SIGCOMM*, 2013.
- [35] J. Wang, D. Vasisht, and D. Katabi. RF-IDraw: Virtual touch screen in the air using RF signals. In *SIGCOMM*, 2014.
- [36] Y. Wang, J. Liu, Y. Chen, M. Gruteser, J. Yang, and H. Liu. E-eyes: device-free location-oriented activity identification using fine-grained WiFi signatures. In *MobiCom*, 2014.
- [37] K. Whitehouse, C. Karlof, and D. Culler. A practical evaluation of radio signal strength for ranging-based localization. *MCCR*, 2007.
- [38] C. Wong, R. Klukas, and G. Messier. Using WLAN infrastructure for Angle-of-Arrival indoor user location. In *VTC*, 2008.
- [39] J. Xiong and K. Jamieson. ArrayTrack: A fine-grained indoor location system. In *NSDI*, 2013.
- [40] J. Xiong, K. Sundaresan, and K. Jamieson. ToneTrack: Leveraging frequency-agile radios for time-based indoor wireless localization. In *MobiCom*, 2015.
- [41] 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 *MobiCom*, 2014.
- [42] Z. Yang, C. Wu, and Y. Liu. Locating in fingerprint space: wireless indoor localization with little human intervention. In *MobiCom*, 2012.
- [43] M. Younis, C. Fischer, and W. Wiesbeck. Digital beamforming in SAR systems. *TOGRS*, 2003.
- [44] M. Youssef and A. Agrawala. The Horus WLAN location determination system. In *MobiSys*, 2005.
- [45] C. Zhou and J. D. Griffin. Accurate phase-based ranging measurements for backscatter RFID tags. *Antennas and Wireless Propagation Letters*, 2012.