

Dude, Where's My Card? RFID Positioning That Works with Multipath and Non-Line of Sight

Jue Wang and Dina Katabi
Massachusetts Institute of Technology
{jue_w,dk}@mit.edu

ABSTRACT

RFIDs are emerging as a vital component of the Internet of Things. In 2012, billions of RFIDs have been deployed to locate equipment, track drugs, tag retail goods, etc. Current RFID systems, however, can only identify whether a tagged object is within radio range (which could be up to tens of meters), but cannot pinpoint its exact location. Past proposals for addressing this limitation rely on a line-of-sight model and hence perform poorly when faced with multipath effects or non-line-of-sight, which are typical in real-world deployments.

This paper introduces the first fine-grained RFID positioning system that is robust to multipath and non-line-of-sight scenarios. Unlike past work, which considers multipath as detrimental, our design exploits multipath to accurately locate RFIDs. The intuition underlying our design is that nearby RFIDs experience a similar multipath environment (e.g., reflectors in the environment) and thus exhibit similar *multipath profiles*. We capture and extract these multipath profiles by using a synthetic aperture radar (SAR) created via antenna motion. We then adapt dynamic time warping (DTW) techniques to pinpoint a tag's location. We built a prototype of our design using USRP software radios. Results from a deployment of 200 commercial RFIDs in our university library demonstrate that the new design can locate misplaced books with a median accuracy of 11 cm.

Categories and Subject Descriptors C.2.2 [Computer Systems Organization]: Computer-Communications Networks

Keywords RFID, Localization, SAR, DTW

1. INTRODUCTION

How many times have you lost your car keys, driver's license, or credit card, and gone through an excruciating search to find the lost object? Locating misplaced objects is also a major problem in pharmacies, libraries, retail stores, and warehouses, where clients and staff move objects around and misplace them. Many industries are looking to address this problem and improve productivity with the deployment of RFIDs in the Ultra-High-Frequency band [16, 17, 7, 38]. For example, the department of veteran affairs has recently signed a \$500M contract to deploy RFIDs in their hospitals and health centers, highlighting their motivation as "Our driving factor was the frustrations in finding the right equipment at the right time" [7]. Current RFID systems, however, identify only whether

the object is within radio range, which can be tens of meters [6]. Many applications however would benefit from knowing the exact location of an RFID-tagged object. For example, in pharmacies, one would want to periodically check that sensitive medications are on certain shelves in a particular storage cabinet. Similarly, if you lose your credit card, instead of simply learning it is in the room as you probably already have guessed, you would like to know its exact location, e.g., whether it is under the couch, in the drawer, in the pocket of your coat, etc. Exact positioning information is also indispensable for some emerging RFID applications like automated customer checkout: pilot programs of RFID-automated checkout revealed that "a shopper could end up paying for the groceries of the person behind her", because today's systems confuse which basket the RFID-tagged goods belong to [4].

Fine-grained RFID localization has recently received much attention [32, 30, 20, 21, 29]. The two key approaches in this domain are based on received signal strength, i.e., RSSI, and the angle of arrival, i.e., AoA. Some researchers augment these methods by considering multiple frequencies, or multiple receiving antennas [32, 8, 41]. However, underlying both methods is an assumption that the signal is mainly propagating along a direct line-of-sight (LOS) path, which is unlikely in most practical RFID setups [42, 18]. Signals bounce off various reflectors in the environment, and multiple reflected copies combine at a wireless receiver, which undermines positioning metrics such as RSSI and AoA [43, 32]. This multipath phenomenon is inevitable in practice and is exacerbated in non-line-of-sight (NLOS) scenarios. Indeed, multipath is the natural mode of operation for UHF RFIDs, which were originally introduced to replace barcodes in non-line-of-sight settings [42].

This paper introduces PinIt, the first fine-grained RFID positioning system that works in the absence of a line-of-sight path and the presence of rich multipath. In line with common practice in locating RFIDs [30, 28, 36], PinIt employs reference tags, whose positions are known a priori. To locate a misplaced object, both the desired tag attached to it and a set of reference tags are queried, and the nearest reference tags are identified. The challenge however is to identify the nearest tags in a manner robust to multipath and NLOS.

Unlike past approaches [32, 36, 29], which consider multipath as detrimental, PinIt exploits a tag's *multipath profile* to locate it. Specifically, signals of nearby tags propagate along closer paths when being reflected off each surface. Hence, if we can obtain a description of all the paths along which a tag's signal propagates - i.e., its multipath profile - neighboring tags can be identified.

To illustrate PinIt's approach, Fig. 1 shows a toy example with three tags, where the blue and green tags are separated by 10 cm, whereas the blue and red tags are separated by 60 cm. As the figure shows, signals from a pair of spatially separated tags, like the blue and red tags, may arrive along exactly the same path from one perspective, but tend to have significantly different paths from a second perspective. In contrast, signals from nearby tags, like the blue and green tags, will always arrive along close paths. Hence, one needs to consider the full multipath profile; one path alone, like the path

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.
SIGCOMM'13, August 12–16, 2013, Hong Kong, China.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
Copyright 2013 ACM 978-1-4503-2056-6/13/08 ...\$15.00.

from Perspective 1, can be misleading. To see this more clearly, we can put an antenna array at the location marked as PinIt’s receiver in Fig. 1, and use its beamsteering capability to separate the signal power received along different paths. Fig. 2 shows the obtained multipath profiles of the three tags from a simulation using standard antenna array equations [33]. It is clear from this figure that looking at the dominant path alone –i.e., the strongest lobe– would wrongly indicate that the blue and red tags are close, while looking at the other path allows us to realize that they are far apart. Thus, a robust positioning scheme needs to compare the whole multipath profiles, by quantifying how the paths get shifted or warped across two profiles.

So how can we automatically quantify the shifts across the multipath profiles of two tags? To do so, we need to align the paths that share the same perspectives. In contrast to the illustrative example, however, real-world deployments have many potential paths, which may shift differently depending on the orientation of the reflectors with respect to the tags. Further, different RFID tags typically have different antenna gains, causing the power along a path to be scaled up or down independently of the location [19]. In designing a technique that finds the nearest neighbor despite these variations, we are inspired by the word matching problem in speech recognition. Even if the same person says the same word twice, she may intonate and speed up differently causing the same word to be stretched and scaled differently across clips. The speech recognition community matches such warped time signals using a technique called Dynamic Time Warping (DTW) [37], which we adopt for our problem. As opposed to warping a time series however, PinIt warps the spatial multipath profile to identify and quantify the similarity between neighboring tags. We present the design in §5, and demonstrate in §8 that the synergy between the multipath effect and DTW enables accurate positioning.

Finally, while one could obtain the multipath profiles using an antenna array, good beamsteering performance would require many antenna elements in the array, which is difficult and costly in practice [24]. Instead, PinIt emulates the antenna array with a single moving omni-directional antenna that slides back and forth. As it moves, at each point in time, the antenna emulates a different antenna element in an antenna array. PinIt collects the signals the antenna receives as it moves and combines these signals to acquire the multipath profiles in post-processing. This design owes its origin to synthetic aperture radar (SAR) imaging, where one antenna is mounted on a flying aircraft or satellite to emulate an antenna array [15].

Summary of Results: We built a prototype of PinIt using USRP software radios and evaluated it with a deployment of 200 commercial UHF RFIDs in our university library (Fig. 9). The tags are placed on 12 racks with a total of 60 shelves and separated by 15 cm. In every experiment, we tag a book and move it to a random shelf. We compare PinIt’s performance in locating the book with state-of-the-art RSSI and AoA based schemes, described in [30] and [32]. PinIt and the compared schemes all use the same set of reference tags. Our experiments lead to the following findings:

- Averaged across 100 different locations of the desired book, PinIt’s median error distance is 11.2 cm, $6\times$ lower than the AoA-based scheme and $10\times$ lower than the RSSI-based scheme. In terms of reliability, the 90th percentile error distances of the AoA and RSSI based schemes are 1.5 m and 2.3 m respectively, whereas PinIt’s 90th percentile is 16 cm. Given this performance, PinIt can address common complaints about unreliable RFID positioning in densely packed environments [3, 5].




Figure 1—Intuition underlying PinIt’s use of multipath in mapping tags: The figure shows a misplaced tag (in blue) and two reference tags (in red and green); the green tag is spatially closer to the blue tag. This physical proximity can be determined only if one considers the directions of both paths. Viewing the reflected paths off the left wall alone will mislead one to think that the desired tag is closer to the red reference tag.




Figure 2—Multipath profiles of tags in Fig. 1.

- In LOS experiments in a typical office lounge setup, PinIt’s performance holds and achieves $4\times$ and $2.4\times$ better accuracy than the RSSI and AoA based schemes.

Contributions: This paper makes the following contributions:

- It presents the first RFID positioning system that exploits multipath as a fingerprint of a tag’s spatial location. As a result, the design delivers centimeter-scale resolution even in rich multipath and non-line-of-sight environments.
- It is also the first to demonstrate the capability of dynamic time warping (DTW) to identify spatial similarity between the multipath profiles of nearby wireless devices, and to successfully use it to locate RFIDs. While our design and results are presented in the context of RFIDs, the basic idea can be extended to other wireless systems.
- It presents a low-cost solution for extracting multipath profiles by using a moving antenna that emulates an antenna array.
- It demonstrates a working system and evaluates it in a real-world large-scale deployment.

2. BACKGROUND

Ultra-High-Frequency (UHF) RFIDs communicate using backscatter. A device called the reader transmits a high power RF signal. Nearby RFID tags reply to the reader’s query by reflecting the high power RF signal using ON-OFF keying. The tags transmit a ‘1’ bit by changing the impedance on their antennas to reflect the reader’s signal and a ‘0’ bit by remaining in their initial silent state [14].

Two points about UHF RFID communication are particularly relevant to the positioning problem.

- There is no carrier frequency offset between the reader and the RFID tags, because the tags do not generate their own RF signal but rather reflect the reader’s signal [14]. Thus, the reader can perform coherent detection, fully recovering the complex channel values of the tags. These channels are available to positioning systems and are used by many past proposals [32, 8, 20, 29] as well as ours.
- The communication range of UHF RFIDs today is from a few meters up to tens of meters (157 feet) [6]. RFID manufacturers

compete to increase the range [6], and hence it is expected to continue growing over the coming years. Providing accurate position information while keeping the range large will be a necessary and valuable feature for RFID systems.

3. PINIT OVERVIEW

PinIt is a fine-grained UHF RFID localization system that can provide a resolution on the order of a few to tens of centimeters, much smaller than the read range of UHF RFIDs.

Following a common practice in RFID localization, PinIt leverages the ultra-low cost of UHF tags (5-10 cents each) to deploy a set of reference tags in the environment, e.g., reference tags in a library can be placed on the shelves of each rack. PinIt’s infrastructure also includes a database which stores the IDs and 3-D positions of the reference tags. Further, PinIt works under the assumption that the objects to be located are tagged with RFIDs, e.g., every book in a library can have a tag adhered to its spine. The association between each object and its RFID is registered in PinIt’s database.

To locate an object, at a high level PinIt goes through the following steps:

- PinIt queries the RFID tag attached to the object (i.e., desired tag), as well as the reference tags.
- PinIt acquires a multipath profile of each queried tag based on its replies, using the technique in §4.
- By comparing and matching the desired tag’s multipath profile with the reference tags’ as described in §5, PinIt identifies the nearest neighbors of the desired tag.
- Finally, based on the positions of these nearest neighbors, it estimates the desired tag’s position, locating the object.

In large-scale deployments like libraries and pharmacies, there can be a large number of reference tags. It is difficult and time-consuming to query them all. In §6, we describe PinIt’s hierarchical protocol to pinpoint the position of the desired tag by gradually zooming in to sub-regions.

The next few sections elaborate on the above steps, providing the technical details.

4. ACQUIRING A TAG’S MULTIPATH PROFILE

An RFID’s multipath profile is a collection of distinguishable paths along which the tag’s signal travels to reach PinIt’s receiver. In this section, we focus on how to differentiate various paths and acquire such a profile for each tag.

Copies of a signal arriving along different paths can be distinguished based on their different delays, like in radar systems. Such an approach however would require ultra-wideband communication, which low-cost RFIDs do not support. Alternatively, multiple paths can be distinguished based on their directions, which is the approach PinIt takes. In order to separate copies of the tag’s signal arriving from different directions, we need to measure the power coming from each direction alone. It is well-known that a large phased antenna array can achieve this goal by forming a very narrow beam and steering it around [40]. When steering its beam to a certain direction, the array essentially filters out power coming from all other directions, i.e., filtering out other paths.

To visually understand this idea, let us again consider the simple setup shown in Fig. 3(a) (same as in §1). Fig. 3(b) shows the power received by an antenna array as a function of its steering angle. The three tags are not replying at the same time, i.e., their signals do not overlap, so these curves are acquired individually. There are two peaks in each tag’s power curve, corresponding to the two reflected paths of its signal. When the array steers its beam to around 40°, it exclusively receives the tags’ signals reflected off the left wall;

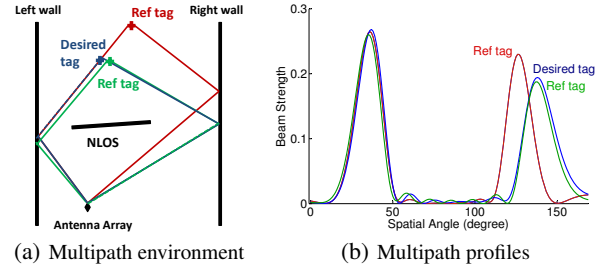


Figure 3—Illustration of multipath profiles: (b) shows the three tags’ multipath profiles captured by the antenna array in (a). The peaks on the left correspond to the reflected paths off the left wall, and the peaks on the right correspond to the reflected paths off the right wall. The spatial angle positions of the peaks reveal that the desired tag is closer to the green tag than the red one.

when the beam direction is steered to around 140°, the reflected paths off the right wall are observed alone. Fig. 3(b) successfully captures the fact that the desired tag (blue) is spatially closer to the green tag: positions of both peaks in the green and blue tags’ curves are close to each other, whereas the red tag’s peak on the right has a 12° shift from the blue tag’s. This example demonstrates the following observation: if two tags are close to each other in space, the whole set of paths in their profiles should closely match. In other words, they should appear to be close overall from various perspectives as revealed by their multipath profiles.

Now we formally define an RFID’s multipath profile as a vector $B(\theta)$, where each element is the power of the RFID’s signal received in the beam at a direction $\theta \in [0^\circ, 180^\circ]$. Next we start by explaining how to compute $B(\theta)$ in a simplified scenario – if we were allowed a uniform linear antenna array. Then in §4.2, we present PinIt’s acquisition of $B(\theta)$ using a single moving antenna.

4.1 Multipath Profile Acquisition Using Antenna Array

An array with K uniformly spaced antenna components is shown in Fig. 4. In order to obtain a narrow beam in the θ direction (i.e., steer the beam to the direction identified by θ), the array *projects* the received signals at all of its antennas on the θ direction. This projection is done by multiplying the received signal at each antenna by a complex number (i.e., weight) in post-processing. Then the received signals are linearly combined with these complex weights, which aligns their phases [33]. Such an operation focuses the beam in the θ direction while minimizing signals in other directions, creating a spatial filter. Below, we establish the standard mathematical formulation of linear array beamsteering.

Let s_k be the tag’s signal observed by the array’s k^{th} antenna, $k = 0, \dots, K - 1$. $w(k, \theta)$ is the complex weight assigned to s_k when steering the beam to θ . λ is the wavelength. D is the length of the array. $x_k = k \frac{D}{K-1}$ is the position of the k^{th} antenna. Then the power received in a beam in the θ direction is computed as:

$$B(\theta) = \left| \sum_{k=0}^{K-1} w(k, \theta) \cdot s_k \right|^2 \quad (1)$$

$$w(k, \theta) = e^{-j \frac{2\pi}{\lambda} x_k \cos \theta} \quad (2)$$

$w(k, \theta)$, the weight assigned to antenna k when steering to θ , has a phase proportional to the distance between antenna k and the origin projected along θ , as Fig. 4 shows. We refer the reader to [33] for the analysis and proof of this standard steering mechanism. At a high level, the weights in Eq. 2 project the received signals to θ by compensating their phases according to the antennas’ positions.

Steering its beam across 180°, the uniform linear array can acquire the multipath profile $B(\theta)$, $\theta \in [0^\circ, 180^\circ]$.




Figure 4—Standard uniform linear array beamsteering: When a linear antenna array steers its beam to direction θ , the received signal at its antenna k (position x_k) is projected to that direction with a phase compensation proportional to $x_k \cos \theta$.

4.2 Multipath Profile Acquisition via Antenna Motion

In practice, it is difficult and often infeasible to deploy a bulky and expensive antenna array to locate low-cost RFIDs [24]. PinIt aims to achieve the spatial filtering capability with a single off-the-shelf omni-directional antenna. Inspired by synthetic aperture radar (SAR), we use antenna motion to emulate an antenna array.

Say we want to emulate a K -antenna array using a single moving antenna. We slide the antenna back and forth along a line using a toy motor. For an array of length D , a naive approach would be to let the antenna move for $\frac{D}{K-1}$ every time, then remain static for a while to query all tags, and repeat this for K times. In this case, after the k^{th} movement, the antenna is at position $x_k = k \frac{D}{K-1}$. Thus the replies captured here are equivalent to the replies captured by antenna k in the uniform linear array in Fig. 4. Therefore, one could combine these replies using Eq. 1 and Eq. 2 to compute the multipath profile of each tag. However, using the moving antenna as such is problematic in practice, because the antenna cannot have an immediate stop and needs to decelerate. Frequent acceleration and deceleration consumes significantly more time and can cause uncertainty in deciding the actual positions where the antenna stops. It is much more manageable and preferable to let the antenna keep moving at a constant speed while receiving the signals. PinIt uses such a design and customizes its post-processing techniques as described below.

Design: PinIt moves the passively listening omni-directional antenna back and forth in a straight line at a constant speed v . We define one full motion of the antenna as the process where it moves from one end of the line to the other end.¹ PinIt’s reader repeatedly queries the desired tag and the reference tags in an interleaving manner. The tags reply with their IDs. PinIt then groups the replies received during one full motion based on the IDs they contain. For each tag, PinIt computes its multipath profile in post-processing.

Post-processing: Given the above design, we need to address several practical issues when post-processing the received signals. First, each tag does not necessarily reply in uniformly spaced time slots [12]. If we consider all the replies PinIt’s moving antenna captures of a particular tag, they do not translate into a uniform array as the one in §4.1. Further, the moving antenna receives the tags’ replies one at a time [12], resulting in slightly different emulated arrays for each tag.

So how can we account for these system constraints and acquire $B(\theta)$, the multipath profile of each tag? The idea is still to project and linearly combine the replies received from each tag over time, but with properly adjusted complex weights. Say the moving antenna has captured K replies from the tag at time t_0, t_1, \dots, t_{K-1} dur-

¹The antenna accelerates and decelerates close to the two ends of the line, and it moves at a constant speed in the middle. PinIt only uses the signals collected at the constant-speed part to acquire the tags’ multipath profiles.

ing one full motion.² Let $s(t_0), s(t_1), \dots, s(t_{K-1})$ be the complex values of the tag’s signal measured during the K replies. Without loss of generality, we use the position of the moving antenna at time 0 to define the origin. To measure the power of the tag’s signal received in a narrow beam at the θ direction, PinIt computes:

$$B(\theta) = \left| \sum_{k=0}^{K-1} w(k, \theta) \cdot s(t_k) \right|^2 \quad (3)$$

$$w(k, \theta) = e^{-j \frac{2\pi}{\lambda} t_k v \cos \theta} \quad (4)$$

As we can see, Eq. 4 has the same form as Eq. 2, only differing in that the phase compensation is now calculated based on timing information (i.e., t_k). This is because time and space are correlated due to the constant-speed antenna motion. At time instant t_k , the moving antenna is at position $t_k v$. Therefore, the reply received at t_k can be considered as one received by an antenna component at position $t_k v$ in a linear antenna array. Similar as in the linear array case, here Eq. 4 projects the received signals all to the direction of θ , and hence achieves the beamsteering effect.

PinIt uses a movement length of $D \approx 2\lambda$ and a speed of $v = 30$ cm/s, and performs the above post-processing to compute $B(\theta)$ for $\theta \in [0^\circ, 180^\circ]$.

Discussion: A few points are worth-noting:

- The above description refers to a linear antenna array, where the antenna elements are on a straight line. Alternatively, one may use a circular antenna array, where the antenna elements are along the circumference of a circle [13]. Circular arrays can produce a slightly sharper beam, but they are less popular because the gain is small and the circular array occupies a larger space. We have experimented with circular motion in acquiring the multipath profiles, moving the antenna using a rotating frame. We found that linear and circular arrays produce quantitatively similar results in terms of accuracy. Thus, we describe our system and also report the results in §9 and §10 using linear motion, which is simpler to deploy.
- Recall that RFIDs communicate by reflecting the reader’s RF signal (using different impedances for “1” and “0” bits). Hence, PinIt’s moving antenna receives both the reader’s signal and the tag’s signal together. Removing the reader’s signal to obtain the tag’s signal $s(t_k)$ is a standard procedure in RFID readers [19], which PinIt also performs.
- One may wonder: how does the number of replies affect the performance? As the moving antenna captures more replies, PinIt can more effectively distinguish the paths. This is because each reply emulates an antenna component in the array. Hence, the spatial filtering is more effective with more replies, as the side lobes are better suppressed in the beam pattern [33]. In our evaluation, PinIt captures about 20 replies from each tag. Not all tags send the same number of replies, since RFIDs may not harvest enough energy to answer every query. However, as long as the number of replies from a tag is above a threshold of 12, we keep that tag in our reference set. Note that PinIt works properly even when different tags have different numbers of replies. The difference in the number of replies merely makes some multipath profiles more precise than others. But the emulated arrays for all tags have the same reference angle θ and their multipath profiles can be properly compared against each other.

²Here, we describe the array formation using only one direction of movement. However, we note that we can use both directions of movement to build the array. When using the other direction, we reverse the order of the replies captured so as to keep the spatial angle axis consistent for all multipath profiles.

- In computing $B(\theta)$, we use uniform sampling of $\cos \theta = 1, 0.99, 0.98, \dots, -1$, instead of uniformly sampling θ . This is because in linear array beamsteering, not all spatial angles in $[0^\circ, 180^\circ]$ have equal resolution. If we uniformly sample θ , the acquired profile is often quite flat close to 0° and 180° but has narrower peaks (higher resolution) close to 90° , which is due to the $\cos \theta$ term in Eq. 4. To compensate for this uneven behavior, we compute $B(\theta)$ for $\theta = \arccos(1), \arccos(0.99), \dots, \arccos(-1)$.
- Finally, note that other movements in the environment can affect the wireless channels. However, their impact on the multipath profiles of neighboring tags is similar. Because PinIt’s positioning is based on reference tags, it is robust to such channel variations. Our experiments were conducted in the presence of people moving around in the deployments, i.e., library, office lounge.

5. COMPARING MULTIPATH PROFILES

After acquiring the multipath profiles, PinIt needs to compare the profile of the desired tag against the profiles of the reference tags in order to identify the nearest reference tags. Thus, we need a method to evaluate the similarity between two multipath profiles \mathbf{B}_i and \mathbf{B}_j of tag i and tag j . But what kind of similarity are we looking for? At an intuitive level, each profile shows peaks that typically refer to the various paths from the tag to PinIt’s receiver (see Fig. 3(b)). When comparing the desired tag’s profile against the profile of another tag, we would like an automated scheme that matches each path (i.e., peak) to one of the paths in the other tag’s profile and then computes the overall similarity.

The distance between two profiles manifests itself in the amount of *misalignment/shifting* between them. Specifically, the paths exhibited by the desired tag’s profile are shifted versions of the paths in its neighbor’s profile. And different paths in the profile are often shifted by different amounts. This is illustrated in Fig. 5, which shows a trace from our library experiment, where the multipath profiles of two tags spaced 30 cm apart on the same shelf are acquired by PinIt’s moving antenna.³ The three peaks reveal that there are three distinguishable dominant paths in the tags’ surroundings. It is clear that the two profiles bear a lot of similarity. Automatically measuring this similarity requires matching the corresponding peaks and quantifying their misalignment/shifts. Each pair of corresponding peaks in the two profiles are misaligned by different amounts, 17° , 14° and -12° respectively. Thus, simply correlating the profiles does not work ($\rho \approx 0.2$ in this case). Instead, what we need is a method that can capture an elastic misalignment of two sequences such as the ones in Fig. 5, in order to accommodate profiles that are intrinsically similar but out of sync at certain segments.

5.1 Dynamic Warping of Multipath Profiles

In designing PinIt’s method for comparing multipath profiles, we are inspired by the word matching problem in the speech recognition community. It is understandable that even if the same person says the same word/sentence twice, she may intonate and speed up differently. As a result, while an underlying pattern remains the same in her speech signal, the two time sequences captured are misaligned at different segments. Therefore, in matching the utterances, a key challenge is to automatically identify these shifts and align corresponding segments. Having recognized the similar nature of PinIt’s multipath profile matching and speech pattern match-

³Note that the SNRs of the two tags are below 8 dB in this trace. However, channel noise in time domain signals only translates into distortions in the *shapes* of multipath profiles, rather than noisy curves. The smooth curves in Fig. 5 do not indicate the experiment has low noise. In contrast to the profiles in Fig. 3(b), the wider bumps here are due to channel noise in real-world deployments.

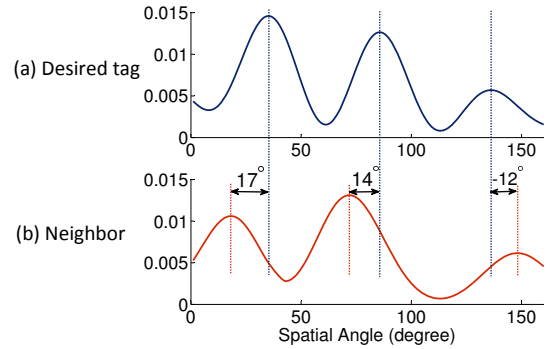


Figure 5—Multipath profiles of two tags 30 cm apart on same library shelf: The figure illustrates the need for a multipath matching algorithm that focuses on quantifying the shifts. The three peaks are shifted by 17° , 14° and -12° each. While the correlation of the two profiles is only $\rho \approx 0.2$, they are intrinsically similar. The question is how to automatically identify and quantify their similarity.

ing, we borrow a well-established technique from speech processing: Dynamic Time Warping (DTW) [37].

Design: Given two time series, and a cost metric, DTW finds an alignment that maps each point in the first series to one or more points in the second series, such that the cost of the mapping summed over all point pairs is minimized. In the context of PinIt, we give DTW two multipath profiles \mathbf{B}_i and \mathbf{B}_j . For any pair of points in the two profile, $B_i(\alpha)$ and $B_j(\beta)$, we define the cost of mapping these points to each other as the Euclidean distance between the two power values:

$$C_{\alpha,\beta} = |B_i(\alpha) - B_j(\beta)| \quad (5)$$

Given this input, DTW searches for the best alignment of the two profiles that minimizes the total cost, using standard dynamic programming [37]. To understand the solution, consider a cost matrix C in Fig. 6(a). The two profiles being compared \mathbf{B}_i and \mathbf{B}_j are shown along the two axes of C . Each cell (α, β) in C , i.e., $C_{\alpha,\beta}$, refers to the cost of aligning $B_i(\alpha)$ and $B_j(\beta)$. For example, the top left cell $C_{0^\circ,0^\circ}$ refers to aligning $B_i(0^\circ)$ and $B_j(0^\circ)$. The lighter color in Fig. 6(a) indicates larger values of $C_{\alpha,\beta}$.

To find the best alignment of the two profiles that minimizes the overall cost, PinIt looks for a route starting from the top left cell (beginnings of the two profiles) to the bottom right cell (ends of the two profiles). This route needs to satisfy the following requirements. First, it must be monotonic: at any cell, the route must next go to the right \rightarrow , or down \downarrow , or diagonal \searrow (right and down) adjacent cell. This means we cannot change the ordering of their elements when trying to align \mathbf{B}_i and \mathbf{B}_j . However, aligning one element in \mathbf{B}_i with multiple consecutive elements in \mathbf{B}_j is allowed and vice versa. Second, we are interested in the route with the lowest total cost. In other words, the sum of all cell values on this route should be the lowest among all possible routes (the green route in Fig. 6(a)). A dynamic programming algorithm is used to efficiently find this route, indicating the best way to align \mathbf{B}_i and \mathbf{B}_j , as visualized in Fig. 6(b). Then we use the total cost on this optimal route to represent the distance between \mathbf{B}_i and \mathbf{B}_j .

Algorithm: Next, we formulate the mathematical explanation of PinIt’s dynamic warping algorithm. Consider the two multipath profiles with equidistant sampling of $\cos \theta = 1, 0.99, 0.98, \dots, -1$:

$$\mathbf{B}_i = B_i(0^\circ), \dots, B_i(\alpha), \dots, B_i(180^\circ) \quad (6)$$

$$\mathbf{B}_j = B_j(0^\circ), \dots, B_j(\beta), \dots, B_j(180^\circ) \quad (7)$$

PinIt computes a cost matrix C using Eq. 5.

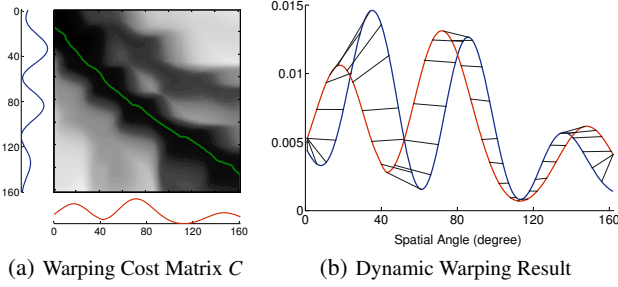


Figure 6—Dynamic warping of multipath profiles: (a) To compute the optimal alignment of the two profiles, DTW finds the route from the top left cell to the bottom right cell of C with the lowest total cost. (b) shows the automatically computed alignment.

The goal is to find a mapping M between \mathbf{B}_i and \mathbf{B}_j . A mapping M is a *contiguous* set of matrix elements in C : $M = m_1, m_2, \dots, m_L$ where $m_l = (\alpha_l, \beta_l)$. L is no less than the length of each profile by definition. The mapping of interest is the one that yields the minimum total cost.

$$\underset{M}{\text{minimize}} \quad \sum_{l=1}^L C_{m_l} = \sum_{l=1}^L C_{\alpha_l, \beta_l}, \quad (8)$$

under the following constraints:

(a) Boundary constraint:

$$\begin{aligned} m_1 &= (0^\circ, 0^\circ) \\ m_L &= (180^\circ, 180^\circ) \end{aligned}$$

(b) Monotonicity constraint:

$$\begin{aligned} \alpha_{l+1} &\geq \alpha_l \\ \beta_{l+1} &\geq \beta_l \\ \alpha_{l+1} + \beta_{l+1} &> \alpha_l + \beta_l \end{aligned}$$

(c) Window constraint:

$$|\alpha_l - \beta_l| \leq W, l = 1, \dots, L$$

The boundary and monotonicity constraints are by definition in DTW. The interpretation of PinIt’s window constraint is as follows: if a reference tag’s paths differ from the desired tag’s by an angle greater than a threshold (30° by default), it is highly unlikely to be a close neighbor. Such a constraint reduces the computational complexity of DTW, because only a diagonal band of width $2W$ in cost matrix C needs to be computed; the other cells can be directly set to $+\infty$. To solve the optimization problem in Eq. 8, PinIt uses a standard solver for DTW based on dynamic programming with linear complexity [37]. Then PinIt records the total cost of this alignment $\sum_{l=1}^L C_{\alpha_l, \beta_l}$ as the distance between the two profiles.

Finally, we note a difference between the algorithm described here and the original DTW. When DTW is used in speech recognition, the goal is to eliminate the effect of timing misalignment, instead of quantifying this misalignment. Hence they normalize the total cost incurred on the green route in Fig. 6(a) by the length of the route. In contrast, when PinIt compares two multipath profiles, the longer this route is, the more misaligned the two tags’ profiles are, indicating that they are spatially further apart. Thus we preserve this quality and compute the total cost without normalization.

5.2 Addressing Power Scaling Challenge

A practical challenge arises when we try to directly apply the above DTW algorithm to compare the multipath profiles: *scaling* of feature values (e.g., peak heights, valley depths). It is widely known that the backscatter efficiency of RFID tags – i.e., the amount of reader’s power they reflect – can differ significantly from one tag to another [30, 36, 20]. The effect of this variation can be removed

by normalizing each tag’s multipath profile by its total power. However, normalization alone is not enough. The RFID tag antenna is often orientation sensitive, i.e., the power it backscatters at each direction is different [19, 1]. Thus, even if two tags are very close, their signal power at a particular direction will not be the same due to their orientations [20, 19, 36]. This causes peaks and valleys to be scaled differently across multipath profiles, independent of the tags’ locations. As we can see from Fig. 5, the peaks in the desired tag’s profile are scaled versions (146%, 96%, 93%) of their counterparts in the neighbor’s profile. This scaling problem could degrade the performance of DTW. The algorithm may fail to find obvious and natural alignment in two profiles simply because a peak or valley in one profile is higher or lower than its corresponding feature in the other profile.

To address these potential variations in feature values, we leverage a simple modification [25] of DTW. Instead of performing DTW directly on the two multipath profiles \mathbf{B}_i and \mathbf{B}_j , PinIt first computes the derivatives of them: \mathbf{B}'_i and \mathbf{B}'_j . Next, each derivative sequence is normalized by its standard deviation. Then PinIt applies the algorithm in §5.1 to align the two normalized derivative sequences. The cost of this alignment is recorded as the distance between the two multipath profiles. It has been shown in [25] that such a design allows DTW to focus on the high level feature of "shape", rather than being bogged down by the absolute values of the sequences, meeting the needs of PinIt. In §8, we demonstrate that derivative DTW provides a robust metric to evaluate the similarity between multipath profiles.

6. HIERARCHICAL APPROACH FOR FINDING THE NEAREST NEIGHBORS

After comparing the multipath profile of the desired tag against the reference tags’, PinIt sorts the profile distances obtained and identifies a few nearest neighbors (the default is 2 nearest neighbors). Then PinIt estimates the desired tag’s position by calculating the weighted mean of its nearest neighbors’ 3-D positions. PinIt’s interface returns the estimated position to the user. In many applications, the user may prefer the location to be expressed with respect to physical objects in her environment, such as “the third shelf on the first rack”. In this case, the reference tags should also be annotated in the database with respect to physical objects in the environment when deployed.

However, trying to identify the nearest neighbors of the desired tag in one step is inefficient in large deployments. In a pharmacy or library, each reader may have hundreds of reference tags in its range. Acquiring the multipath profiles of all reference tags, comparing them with the desired tag’s profile, and sorting the computed distances could take minutes. Hence, in order to keep the communication cost and computational complexity low, PinIt searches for the nearest neighbors hierarchically using the protocol below.

Protocol: In stage 1, PinIt’s reader starts by querying and acquiring the profiles of a subset of reference tags spaced 1 m apart from each other as well as the desired tag. It is possible that not all of these reference tags will reply. PinIt identifies the n nearest neighbors of the desired tag among the ones that reply, where by default $n = 2$. If the nearest neighbors have significantly lower profile distances to the desired tag than other reference tags, PinIt proceeds to stage 2. Otherwise, PinIt checks whether any of the tags in stage 1 has not replied, and replaces that tag by a nearby reference tag. It queries the replacement tags until the similarity between the desired tag and its n nearest neighbors exceeds a threshold or PinIt already has replies from each 1 m^3 sub-region.



Figure 7—Antenna and commercial UHF RFID used in experiments: (a) We use the VERT900 6-inch vertical antenna as PinIt’s moving antenna. (b) the Alien Squiggle General Purpose UHF RFID Tags. Note that the sizes of (a) and (b) are not proportional to the real sizes of the devices.

In stage 2, PinIt zooms in to n sub-regions each of size 1 m^3 centered at the n nearest neighbors respectively. In each sub-region, PinIt queries a subset of reference tags spaced by 30 cm and identifies the n nearest neighbors among all the reference tags queried so far. In the vast majority of cases, the n nearest neighbors at this stage converge to the same sub-region, according to PinIt’s database. In the rare cases they do not converge, PinIt keeps zooming in on the candidate sub-regions until the n nearest neighbors converge.

In stage 3, PinIt zooms in to query a subset of reference tags spaced by 15 cm and identifies the nearest neighbors. In our field experiment in the library, the minimum spacing between reference tags is 15 cm. Hence, PinIt’s protocol terminates at this stage. However, if there are reference tags available at even finer granularity in the deployment, PinIt can keep zooming in until finding the nearest neighbors of the desired tag at the finest resolution supported.

Communication Cost and Computational Complexity: The complexity of PinIt comes in two parts: the time spent on hierarchically acquiring the multipath profiles, and the complexity of comparing the reference tags’ profiles with the desired tag’s. Let N be the total number of reference tags PinIt queries, K be the average number of replies captured from each tag, and L be the average length of the DTW route. The complexity of querying the reference tags and computing their multipath profiles is $O(KN)$. The complexity of computing derivative DTW between the desired tag’s profile and the queried reference tags’ is $O(LN)$. Sorting the reference tags has a complexity of $N \log(N)$. One should choose the parameters N , K and L properly based on system capabilities, constraints and requirements.

For example, the library section in our experiment is of size $6 \times 5 \times 2.2 \text{ m}^3$ and has 12 racks with 200 reference tags deployed. To pinpoint a UHF RFID to a median accuracy of 11.2 cm (error distance between estimated position and actual position), $N = 24$ reference tags (12 in stage 1, 8 in stage 2, and 4 in stage 3) are queried hierarchically. Acquiring their multipath profiles uses $K \approx 20$ replies from each of them. The total communication cost of getting these replies is 5.4 seconds using the SELECT and QUERY commands and the default rates in the EPC Gen-2 UHF RFID standard [12]. In computing the multipath profiles, we sample the spatial angle θ such that $\cos \theta = 1, 0.99, \dots, -1$. Thus $L \geq 201$. The time needed to perform derivative DTW of 24 pairs of multipath profiles of this length is less than 0.5 seconds using MATLAB on a 64-bit machine with Intel Core i7-2600 Quad-Core processor. Hence, the overall time complexity for locating a misplaced object in this environment is under 6 seconds.

7. IMPLEMENTATION

We build a prototype of PinIt using USRP software radios [22] to locate commercial EPC Gen-2 UHF RFIDs.

Reader: We adopt a USRP1 implementation of an EPC Gen-2 RFID reader developed in [10] and modify it to work with

USRP N210. UHF RFID systems in the U.S. operate in the 902-928 MHz band. We use RFX900 USRP daughterboards and Cushcraft 900 MHz RFID antennas [26] and run all experiments at a carrier frequency of 910 MHz.

Moving Antenna: We use an omni-directional VERT900 antenna (Fig. 7(a)) connected to a USRP N210 device to receive.⁴ To move the antenna, we place it on an iRobot Create robot [2] which is programmed to move back and forth along a line at a constant speed of 30 cm/s. The USRP samples the received signals using a rate of 4 MHz. Due to the overhead of user mode processing in software radios, we do not run PinIt’s algorithms in real time. Instead, we collect the traces from this USRP receiver and process them offline.

Tags: The Alien Squiggle General Purpose tag shown in Fig. 7(b) is a passive backscatter RFID in the UHF band. We use these RFIDs to tag objects and as reference tags.

8. MICROBENCHMARK

We start with a microbenchmark experiment to provide insights into the working of PinIt. In particular, to better understand how derivative DTW works with spatial multipath profiles, we construct a simple topology of one desired tag and three reference tags 15 cm, 30 cm and 50 cm away from it. All four tags are attached to objects on a movable shelf, facing PinIt’s moving antenna (4 m away) with no obstruction in the middle. This setup is in a typical office lounge.

Fig. 8(a) shows the spatial multipath profiles of the four tags PinIt has acquired using the moving antenna, Fig. 8(b) shows the derivatives of the multipath profiles, and Fig. 8(c) shows the warping PinIt performs on the derivative sequences to compare the desired tag against the three reference tags. These figures reveal the following:

- Even in strong line-of-sight, the multipath effect is still evident. Despite our effort to keep the shelf and moving antenna away from all reflecting surfaces except for the floor, the tags’ signals still experience multiple paths in the environment, as indicated by the multiple bumps in each profile. The reason for this is the reflection off the shelf and objects it contains. This is a representative case in RFID applications, where each tag is attached to an object. The object itself, as well as other objects nearby, naturally creates a multipath environment.
- Due to the power scaling problem, we face a trade-off between *scale* and *shift* in Fig. 8(a) when trying to decide which reference tag is closer to the desired tag. On one hand, it may seem that the red tag is closer to the desired tag (blue), since the two main bumps (paths) have similar heights (scales) in the red and blue profiles. In comparison, the green tag’s bump around 80° is much lower than the desired tag’s, although it is the closest spatially. On the other hand, from the "shift" perspective, the green tag’s profile aligns better with the desired tag’s, indicated by the positions of peaks and valleys. This dilemma demonstrates the power scaling challenge caused by the orientation-sensitive tag antennas (§5.2), and the need to focus on shifts rather than scales.
- Derivative DTW allows us to overcome the power scaling problem by focusing on shifts of the features (e.g. peaks and valleys). The similarity between the green tag and blue tag becomes clearer in their derivatives shown in Fig. 8(b), which justifies the use of derivative DTW.
- Derivative DTW can provide a robust metric for spatial similarity. Fig. 8(c) shows a gradual increase in the cost of derivative DTW as the distance to the desired tag increases. It costs the least to warp the derivative of the desired tag’s profile to align

⁴One could potentially incorporate the moving antenna into the reader. For ease of implementation, we keep them separate.

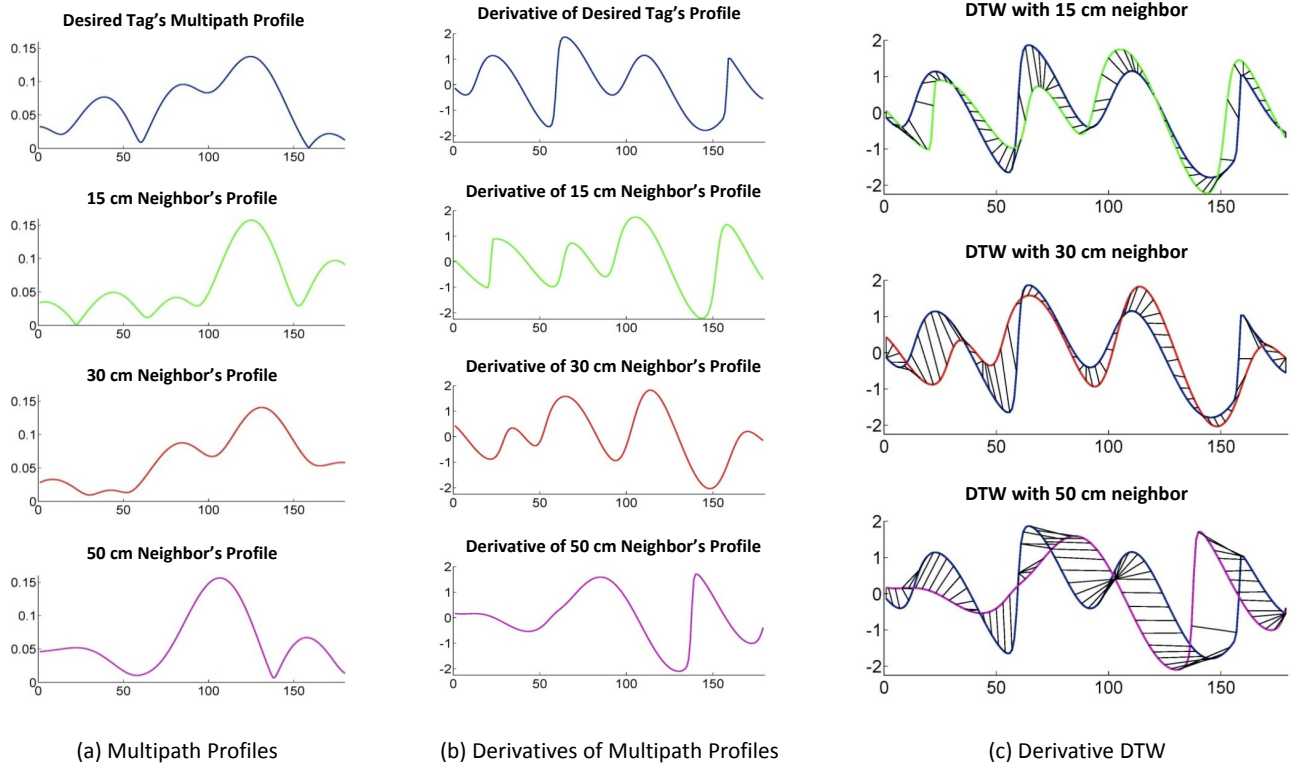


Figure 8—Identifying the nearest neighbor by dynamically warping the derivatives of multipath profiles: (a) shows 4 UHF tags’ multipath profiles acquired in a LOS experiment. As these profiles indicate, even with dominant LOS the tags’ signals still experience multiple paths. It is difficult to tell which reference tag is closer to the desired tag based on (a). In particular, due to the power scaling problem, the nearest neighbor (green) has a much smaller bump around 80° than the desired tag (blue). In order to overcome this problem and focus on "shifts", PinIt computes the derivatives of the profiles, shown in (b). Here the similarity between the desired tag and the nearest neighbor becomes more obvious. (c) shows how PinIt automatically measures this similarity by warping the derivatives. The cost of warping the blue curve to align with the green curve is the lowest. Thus, PinIt successfully identifies the green tag as the nearest neighbor of the desired tag. Further, as the spatial distance increases, the cost of derivative DTW gradually increases accordingly.

with the derivative of the green tag’s profile. Thus, PinIt identifies the green tag as the nearest neighbor. It costs $1.5\times$ more to align the blue curve with the red, and even more to align it with the purple curve.

- Finally, the reader may be wondering why we do not see the familiar white channel noise in these figures. Recall that the multipath profile curves do not show the time signals, but are rather the sum of the power in a beam along different spatial directions. Noise in time domain signals translates into distortions in the *shapes* of multipath profiles in spatial domain, e.g., peak and valley widening (as also clear in past work [45, 18]). In this experiment, the signals of the four tags have mean SNRs of 6.2 dB, 5.7 dB, 5.4 dB, and 7.9 dB each. The smooth curves in Fig. 8(a) do not suggest the experiment had low noise. Indeed, unlike the clear narrow beams in Fig. 2, the ambiguous profiles with wider bumps here are the result of channel noise.

In conclusion, this microbenchmark demonstrates how PinIt leverages the synergy between the multipath effect in wireless communication and the DTW technique to identify and quantify spatial similarity. It also underscores the need to use the derivative version of DTW in order to address the practical challenge of power scaling in RFID localization.

9. EVALUATION IN FIELD EXPERIMENTS

We conduct field experiments in the engineering library of our university (Fig. 9). 200 reference tags evenly spaced by 15 cm are deployed on 12 racks (60 shelves) in a library section of size

$6\times 5\times 2.2$ m³. Different rows of racks are separated by about 1m. Since there are no books to be located outside of the shelves, we do not put any tags on the floor or ceiling. The shelves are made of metal and full of books, resulting in a complex multipath environment. We adhere the desired tag to the spine of a book and move the book around to 100 different locations within the 12 racks.

We compare PinIt with two baseline schemes in locating the book, which also use the same set of reference tags.

- **RSSI-based:** The difference between the RSSIs of a pair of tags is used as a metric for their spatial distance as in past work [30, 36]. Based on the RSSI distances, the nearest neighbors of the desired tag are identified.
- **AoA-based:** The difference between the angles of arrival is used as a metric for the spatial distance between tags as in past work [28]. Based on the AoA distances, the nearest neighbors of the desired tag are identified.

There are two ways to obtain RSSIs and AoAs of the tags. First, we could provide both baselines with all the SAR antenna array measurements that PinIt uses. Alternatively, past work has proposed using multiple receive antennas at different locations to measure the RSSIs and AoAs of RFIDs [30, 20, 32]. To evaluate which of these two options works better for RSSI and AoA each, we first performed a set of calibration experiments. In particular, we allow the compared schemes to either use the same channel measurements as PinIt, or measure the RSSIs or AoAs at three locations. In the second option, one antenna at each location is used to measure the



Figure 9—Library deployment: We conducted field experiments with a deployment of 200 UHF RFIDs in our university library.

RSSI, while a pair of antennas at each location are used to compute the AoA as recommended in [30, 32].

The detailed calibration results for both LOS and NLOS are included in the appendix. The output of the calibration experiments shows that the AoA-based scheme works better using PinIt’s SAR measurements than three pairs of antennas at three locations. This is expected since a SAR antenna array enables a higher resolution in AoA calculation, which in the presence of reference tags is more valuable than three separate locations. On the other hand, the RSSI-based scheme does not benefit from the spatial angle information offered by SAR. Hence its performance is better when using three different locations. Based on these findings, in the rest of our evaluation, the RSSI-based scheme uses three antennas to receive at three different locations, and the AoA-based scheme and PinIt use channel measurements from PinIt’s SAR. The SAR moving antenna is placed in a corner of the library section, 3-7 meters away from the tags. One of the three RSSI receivers is co-located with the SAR receiver, and the other two are placed at two other corners of the same section, 4.5 and 5 meters away from the first antenna.

All three schemes follow the hierarchical approach: In stage 1, one reference tag at the center of each rack is queried.⁵ We identify the two nearest neighbors among these queried tags to decide which two racks to zoom in. In stage 2, for each of the two candidate racks, we query another four tags in it. Based on all the tags queried so far, we decide on two shelves to further zoom in. In stage 3, within each of the two candidate shelves, we query two more reference tags. Then we identify the two final nearest neighbors. Once having identified the nearest neighbors, all three schemes estimate the desired tag’s position as the weighted mean of the nearest neighbors’ positions as recommended in [30].

Finally, we use the error distance between the estimated position and the actual position of the desired tag as the performance metric in our comparison. The lower this error distance is, the more accurate a positioning system is. Note that this error distance could be lower than 15 cm (the reference tag spacing). Because the book is at a random location, its distances to its neighboring reference tags are less than the spacing between reference tag.

9.1 Accuracy in Locating the Mislaced Book

First we discuss the aggregated results over 100 different placements of the desired book. Fig. 10 shows the CDF of the error distances in PinIt, RSSI and AoA based schemes.

- The RSSI-based scheme has a median error distance of 112 cm and a 90th percentile of 229 cm. This agrees with previous results reported for RSSI-based RFID localization [30, 36]. The

⁵In the cases where the center tag does not reply, we continue to query the next tag closest to the center, and keep doing so until finding a tag in the rack that replies. The same applies to later stages.




Figure 10—CDF of localization error distance: PinIt locates the desired tag to a median accuracy of 11.2 cm, 10× and 6× better than the RSSI and AoA based schemes. The 90th percentile error distances of the RSSI and AoA based schemes are 2.3 m and 1.5 m respectively, while PinIt’s 90th percentile is 16 cm, significantly improving the reliability of RFID positioning.

shelves and books in the library create a complex multipath environment, undermining the correlation between location and signal power. In addition, RSSI also suffers from the variation in behavior across tags, such as the antenna gain and backscatter efficiency of individual tags. Note that the racks are spaced by approximately 1 m, which means in many cases the RSSI-based scheme found the wrong rack altogether.

- The AoA-based scheme has a median error distance of 68 cm, 39% reduction as compared to the RSSI-based scheme. Its 90th percentile is 149 cm. Past work has also noted that AoA is more robust than RSSI [32], because it focuses on the angle and is less sensitive to the different antenna gains across tags.
- PinIt has a median error distance of 11.2 cm and a standard deviation of 6.2 cm, outperforming RSSI and AoA by 10× and 6× respectively. Its 90th percentile is 16 cm, and 99th percentile is 25 cm. This significant improvement is due to PinIt’s ability to extract the rich spatial multipath profiles and its use of derivative DTW to identify intrinsic similarity between profiles. Note that in almost all traces, simple correlation did not work due to elastic path shifts and power scaling, emphasizing the need for derivative DTW. Recall the trace in Fig. 5 where the dominant peaks in the two profiles correspond to two different paths. Since AoA focuses on the dominant peak, it ends up confusing paths in these cases, while derivative DTW overcomes the issue by focusing on the shifts of features instead of scales.

In conclusion, bringing together the multipath profile and derivative DTW enables PinIt to achieve centimeter-scale accuracy and reliable positioning performance in a real-world deployment of commercial RFIDs.

9.2 Effect of the Distance from Desired Tag

Next we investigate the impact of the distance from the desired tag to PinIt’s moving antenna. Note that in the RSSI-based scheme, the definition of this distance is unclear since there are three antennas receiving at distributed locations. Thus, we exclude the RSSI-based scheme from this discussion. For PinIt and the AoA-based scheme, we group all the experiment locations of the tagged book into seven sets, 3.5 to 6.5 meters away from PinIt’s moving antenna.

Fig. 11 shows the error distance of each group. As we can see, PinIt consistently outperforms the AoA-based scheme at all distances. Further, no clear correlation is observed between the distance to the receiver and the performance of these two schemes. This is because the rich multipath in NLOS weakens the concept of "range" in the library setting.

9.3 Evaluation of Hierarchical Protocol

The hierarchical search has three stages. In stage 1, it operates at the rack level. In stage 2, it zooms in to the shelf level. In stage 3,

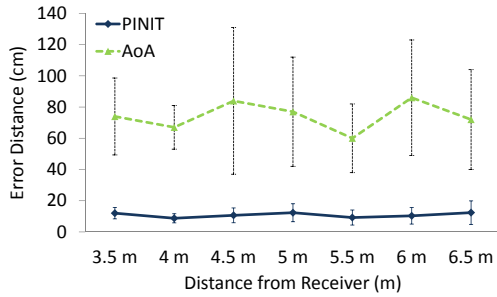


Figure 11—Impact of distance from receiver: Due to multipath, the performances of the AoA-based scheme and PinIt do not exhibit clear correlation with the distance between the desired tag and the SAR receiver. The RSSI-based scheme is excluded from this figure because its three receivers are distributed and hence it is unclear how to define the distance to tag.

it further zooms in to find the nearest neighbors among all tags. In this evaluation, we study the performance at the end of each stage.

A localization system should identify the correct nearest neighbors even when the reference tags are sparsely distributed in the environment, e.g., 1 m on the rack level, 30 cm on the shelf level. Further it is important to investigate the robustness of the hierarchical protocol. For example, what if the protocol identifies the wrong rack or shelf as the nearest neighbor in the early stages; would its performance improve in later stages by converging back to the correct position when more reference tags are queried?

Fig. 12 shows the localization error distances achieved by RSSI, AoA and PinIt at each stage of the protocol.

- PinIt performs better than RSSI and AoA based schemes in all three stages. Consider stage 1 as an example: PinIt achieves a median error of below 0.5 m, whereas the error distances for both RSSI and AoA are above 1 m. To understand this, we looked at in how many cases the three schemes successfully identified the reference tag in the correct rack as the nearest neighbor in stage 1. The success rates are shown in the first column of Table 1. As we can see, PinIt correctly identified the rack the desired tag belonged to in 98 out of 100 experiments. This explains why its median error distance is below 0.5 m. On the other hand, RSSI and AoA based schemes only found the correct rack in 19 and 43 experiments out of 100. Thus, their median error distances in stage 1 are considerably higher. Similarly, the rates that the RSSI and AoA based schemes successfully found the correct shelf or the nearest neighbor are also much lower than PinIt in stages 2 and 3, as the second and third columns of Table 1 show.
- All three schemes improve as the protocol proceeds and more reference tags are queried, indicating that the hierarchical approach is robust. This is true even if the wrong nearest neighbor is identified on the rack level. For example, RSSI and AoA identified the wrong rack as the nearest neighbor in 81% and 57% of the cases in stage 1, but managed to improve their accuracy in later stages as Fig. 12 shows. The reason for this is that, at the end of stages 1 and 2, the hierarchical protocol can keep two candidate racks/shelves to further zoom in next. This redundancy allows the schemes to potentially converge to the correct rack or shelf in the next stage, even if it is not identified as the top 1 nearest neighbor now.

10. EVALUATION IN LOS AND NLOS

In this section, we answer the following two questions: First, how do the localization systems compare in LOS and NLOS respectively? Second, what is the finest resolution each system can achieve given a large density of reference tags? The library environment, like most real-world RFID deployments, is largely dominated




Figure 12—Error distances at different stages of hierarchical search: All three schemes’ accuracy improves as the protocol proceeds and more reference tags are queried. PinIt consistently outperforms the baselines at each stage of the hierarchical search.

	Stage 1: Finding the rack	Stage 2: Finding the shelf	Stage 3: Finding nearest neighbor
PinIt	98/100	95/100	86/100
RSSI	19/100	6/100	4/100
AoA	43/100	8/100	0/100

Table 1—Identifying the nearest neighbor: This table shows the number of cases where the three schemes identified the reference tag in the correct sub-region (e.g. rack, shelf) as the top 1 nearest neighbor at each stage of the hierarchical protocol. Note that at the end of each stage, two nearest neighbors (sub-regions) are kept rather than one, which improves robustness of the protocol.

by NLOS with strong multipath [18], and hence does not allow us to understand the behavior of the localization schemes in LOS. Thus, to study these aspects of the systems, we use a separate setup in a typical office lounge. We place up to 30 tags on a few movable shelves and utility carts in the lounge and vary the minimum spacing between the tags in a range of [1.875 cm, 1 m]. To create a LOS setting, we remove all obstructions between the tags and the receivers, i.e., PinIt’s moving antenna and the three RSSI receive antennas. For NLOS, we block the direct paths using equipment on the shelves and carts, including metal structures. In different runs of the experiment, PinIt’s moving antenna is placed 3 to 5 meters away from the tags. One of the three receive antennas in the RSSI-based scheme is co-located with PinIt’s moving antenna, while the other two are 4 m away from the first antenna. The collection of the three antennas forms a 90° angle in the lounge. Since the number of tags is small in this experiment, we do not run the hierarchical protocol, but instead exhaustively search for the two nearest neighbors of each tag using the RSSIs, AoAs, and multipath profiles. We then estimate each tag’s position as the weighted mean of its two nearest neighbors identified.

Fig. 13 and 14 show the error distances for different tag spacings in LOS and NLOS experiments respectively.

- Given a set of reference tags, the achievable localization resolution is dependent on the minimum tag spacing. In both LOS and NLOS, the performance of all three schemes improves as the minimum tag spacing decreases.
- The RSSI and AoA based schemes perform significantly better in LOS than in NLOS. For example, with 15 cm spacing, their median errors in LOS are 38 cm and 24 cm, in line with past results for RFID localization in LOS setups [8]. In NLOS their performances degrade by about 3×, confirming the impact multipath and NLOS have on RSSI and AoA based RFID localization.
- Overall, PinIt achieves 4× and 2.4× better accuracy than RSSI and AoA in LOS (median of improvement across different spacings), and 10× and 6.4× in NLOS. The reason PinIt still outperforms the baselines in LOS is because multipath effects exist




Figure 13—Stressing the systems in line-of-sight: In LOS, PinIt’s overall accuracy is $4\times$ better than the RSSI-based scheme and $2.4\times$ better than the AoA-based scheme. The finest resolutions of PinIt, RSSI and AoA are 3.22 cm, 18 cm and 10 cm respectively.

even with an evident LOS path, as the microbenchmark experiments (§8) showed. PinIt’s ability to leverage the entire multipath profile improves localization accuracy. Note that even in an ideal free space channel model, the performance of PinIt will be equivalent to the AoA-based scheme.

- PinIt achieves a resolution of around 3 cm in both LOS and NLOS when the tags are spaced by 1.875 cm. The resolutions of RSSI and AoA based schemes do not improve beyond 18 cm and 10 cm in LOS, or 33 cm and 21 cm in NLOS. Reducing the spacing to under 7.5 cm does not linearly improve the resolution of PinIt. Its performances at 3.75 cm and 1.875 cm spacing are similar, as SAR reaches its limit in disambiguating neighboring tags and UHF RFIDs start to experience the near-field effect [14].

In conclusion, PinIt enables centimeter-scale RFID positioning in both line-of-sight and non-line-of-sight scenarios.

11. RELATED WORK

Early work on RFID positioning relied on RSSI and scene analysis [21, 30]. Typically, RSSIs of reference tags at known positions are measured and used to locate the desired tag [30]. RSSI-based techniques however cannot provide reliable location estimates in environments dominated by multipath or non-line-of-sight paths [18]. Further, RSSI is affected by the tag’s orientation and antenna gain [20, 19], which adds uncertainty to positioning.

There is a growing interest in using AoA information (Angle of Arrival) to locate RFIDs by measuring the phase difference between the received signals at different antennas [20, 8, 32, 41]. The feasibility of this approach has been demonstrated for both active [20] and passive [8] RFIDs in setups where an evident LOS path is present. However, it remains a major challenge for these techniques to mitigate the multipath effect and deal with NLOS [32, 41].

Researchers have also investigated the possibility of using FMCW (frequency modulated continuous wave) and UWB (ultra wideband) to estimate the time of flight of signals and locate RFIDs [18]. It has been reported that sub-meter accuracy can be achieved in strong multipath at a cost of a large bandwidth, e.g., 150MHz. However today’s UHF RFIDs communicate in 500 KHz narrow band channels [12]. Hence, applying these schemes would require modifying both RFID hardware and regulations.

Because of their directional beamsteering capabilities, antenna arrays and synthetic aperture (SAR) methods are commonly used in wireless localization [45, 15, 29]. Radar systems have traditionally used antenna arrays and SAR methods for both object localization and terrain imaging [15]. In the WiFi domain, the work in [45] presents a state-of-the-art WiFi localization system using antenna arrays at the APs. In the context of RFIDs, the approaches in [8, 29, 32, 34] use antenna arrays and SAR to infer a tag’s location based on a LOS model. While PinIt employs a SAR-style antenna array,

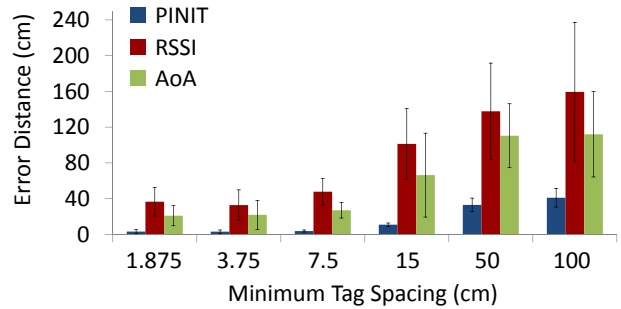


Figure 14—Stressing the systems in non-line-of-sight: In NLOS, PinIt’s overall accuracy is $10\times$ better than the RSSI-based scheme and $6.4\times$ better than the AoA-based scheme. The finest resolutions of PinIt, RSSI and AoA are 3.16 cm, 33 cm and 21 cm respectively.

it significantly differs from all past systems in how it uses the array. Past schemes use the array to extract the direct LOS path, as they struggle with the multipath effect and aim to eliminate it. In contrast, PinIt uses the array to obtain the whole multipath profile and is the first to establish a connection between the multiple paths in wireless propagation and the Dynamic Time Warping technique (DTW). The synergy between the multipath profile and DTW enables PinIt to map and compare the whole set of paths in one profile to the set of paths in another profile. This approach provides much richer information and enables RFID localization in scenarios with strong multipath effects and non-line-of-sight situations.

PinIt is related to a wide range of work in locating wireless user devices [9, 39, 31, 11, 35, 45, 23]. Most of these systems are based on RSSI measurements when communicating with multiple WiFi APs nearby. Some use the angle of arrival (AoA). In particular, the work in [31] uses a rotating directional antenna to measure the AoA and estimate ranges for triangulation and trilateration. While PinIt also uses antenna motion, it employs an omni-directional antenna and processes its signal using SAR to enable inexpensive, automatic beamsteering. Further, PinIt’s signal processing is different: it employs a variant of DTW to match multipath profiles for localization, which allows it to achieve a significantly higher accuracy. Also, recently a few WiFi localization proposals [45, 23] have used antenna arrays at multiple APs to mitigate multipath and distill the direct path for localization. Compared with these systems, PinIt differentiates itself in objectives, techniques, and operating environments. PinIt’s objective is to locate RFIDs, which by nature are deployed in rich multipath and often NLOS environments. To do so, PinIt exploits the whole multipath profile as opposed to the signal along the direct path. It also performs DTW to evaluate similarity between the desired tag’s profile and the reference tags’ profiles.

Finally, a few recent proposals apply DTW to accelerometer data to match the trajectories of pedestrians and identify situations where people walk along similar paths [44, 27]. PinIt’s use of DTW differs in that it is based on the spatial multipath profiles, which we extract using SAR processing.

12. CONCLUDING REMARKS

Industries such as retailing and healthcare are deploying billions of RFIDs to improve efficiency, safety, and availability. A common complaint they have about RFID systems is the lack of reliable, fine-grained position information [3, 5]. RSSI and AoA based solutions provide a 90th percentile accuracy of 2.3 m and 1.5 m, inadequate for many applications today. For example, it causes errors in RFID-automated customer checkout [4].⁶ Furthermore, the health-

⁶This is a major concern about replacing barcodes with RFIDs in customer-facing applications, and also a reason for inaccurate data collection in inventory management. While it could potentially be solved by reducing the range

care industry is expected to spend over \$1B on item-level tagging of drugs and medical disposables (e.g., chemotherapy) [38]. With potentially hundreds of doses within a small area, relying on RFIDs to track drugs and eliminate medical errors in such densely packed environments requires higher precision than what existing positioning schemes can offer.

By exploiting multipath, the natural mode of operation for RFIDs, PinIt presents a solution based on dynamic time warping that achieves 16 cm 90th percentile accuracy and 25 cm 99th percentile accuracy, enabling reliable centimeter-scale positioning in densely packed environments.

We also believe that while the operating assumptions of RFIDs are different (e.g., low-cost reference tags, low-power), the general techniques in this paper are transferable to other wireless systems.

Acknowledgments: We thank Haitham Hassanieh, Fadel Adib, Hariharan Rahul, Nate Kushman, Lixin Shi, Shuo Deng, Jonathan Perry, Arthur Berger, the reviewers and our shepherd, Kun Tan for their insightful comments. This research is funded by NSF. We thank the members of the MIT Center for Wireless Networks and Mobile Computing, including Amazon.com, Cisco, Google, Intel, MediaTek, Microsoft, ST Microelectronics, and Telefonica, for their interest and support.

13. REFERENCES

- [1] ALN-9640 Squiggle Inlay. www.aliantechnology.com.
- [2] iRobot Create Programmable Robots. www.irobot.com.
- [3] Physical and Psychological Barriers to RFID Adoption. *RFID Journal* 2010.
- [4] RFID, a Vision of the Future. rsa.com/rsalabs/node.asp?id=2117.
- [5] Rfid interrogators grow up. *RFID Journal* 2008.
- [6] Impinj revolutionizes rfid readers, 2011. www.rfid.net.
- [7] How rfid is transforming va hospital operations, 2012. www.rfidjournal.com/article/view/9819.
- [8] S. Azzouzi et al. New measurement results for the localization of uhf rfid transponders using an angle of arrival (aoa) approach. In *IEEE RFID 2011*.
- [9] P. Bahl and V. Padmanabhan. Radar: an in-building rf-based user location and tracking system. In *INFOCOM*, 2000.
- [10] M. Buettner and D. Wetherall. A software radio-based uhf rfid reader for phy/mac experimentation. *IEEE RFID*, 2011.
- [11] K. Chintalapudi, A. Padmanabha Iyer, and V. N. Padmanabhan. Indoor localization without the pain. *MobiCom '10*.
- [12] EPCglobal Inc. EPCglobal Class 1 Generation 2.
- [13] R. Fienberg. Limitations on directional patterns of phase-compensated circular arrays. 1965.
- [14] K. Finkenzeller. *RFID Handbook*. John Wiley & Sons, 2010.
- [15] J. Fitch. Synthetic aperture radar. 1988.
- [16] Frost & Sullivan. Global RFID healthcare and pharmaceutical market. Industry Report, 2011.
- [17] Frost & Sullivan. *RFID Market in Apparel Supply Chains*. 2011.
- [18] L. Gang et al. Bandwidth dependence of cw ranging to uhf rfid tags in severe multipath environments. In *IEEE RFID 2011*.
- [19] J. Griffin and G. Durgin. Complete link budgets for backscatter-radio and RFID systems. *IEEE Antennas and Propagation Magazine*, 2009.
- [20] C. Hekimian-Williams et al. Accurate localization of rfid tags using phase difference. In *IEEE RFID 2010*.
- [21] J. Hightower et al. Spoton: An indoor 3d location sensing technology based on rf signal strength, 2000.
- [22] E. Inc. Universal Software Radio Peripheral. <http://ettus.com>.
- [23] K. Joshi, S. Hong, and S. Katti. Pinpoint: Localizing interfering radios. *NSDI '13*.
- [24] N. Karmakar. Handbook of smart antennas for rfid systems, 2010.
- [25] E. Keogh et al. Derivative dynamic time warping. In *SDM 2001*.
- [26] Laird Technologies. Crushcraft S9028PCRW RFID antenna.
- [27] F. Li et al. A reliable and accurate indoor localization method using phone inertial sensors.
- [28] F. Manzoor and K. Menzel. Indoor localisation for complex building designs using passive rfid technology. 2011.

of RFIDs, such a design choice undermines the purpose of using RFIDs to improve efficiency.

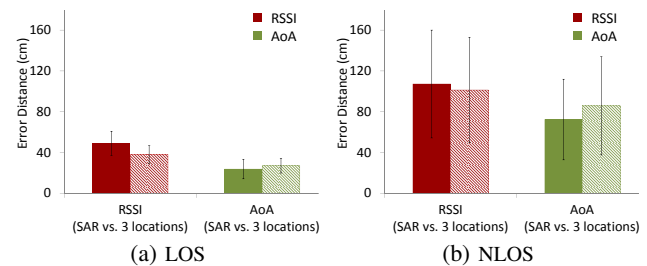


Figure 15—Choosing between SAR and multiple locations: (a) shows the comparison between the two options for both baselines in LOS. (b) shows the NLOS results. RSSI performs better when using three distributed receive antennas than using the SAR measurements, while AoA benefits more from the SAR measurements.

- [29] R. Miesen et al. Holographic localization of passive uhf rfid transponders. In *IEEE RFID 2011*.
- [30] L. M. Ni et al. Landmarc: indoor location sensing using active rfid. *Wirel. Netw.*, 2004.
- [31] D. Niculescu and B. Nath. Vor base stations for indoor 802.11 positioning. *MobiCom '04*.
- [32] P. Nikitin et al. Phase based spatial identification of uhf rfid tags. In *IEEE RFID 2010*.
- [33] S. J. Orfanidis. *Electromagnetic waves and antennas*, 2010.
- [34] A. Parr et al. A novel method for uhf rfid tag tracking based on acceleration data. In *IEEE RFID 2012*.
- [35] A. Rai, K. K. Chintalapudi, V. N. Padmanabhan, and R. Sen. Zee: zero-effort crowdsourcing for indoor localization. *Mobicom '12*.
- [36] S. N. Razavi and C. T. Haas. Using reference rfid tags for calibrating the estimated locations of construction materials. *Automation in Construction*, 2011.
- [37] S. Salvador and P. Chan. Toward accurate dynamic time warping in linear time and space. 2007.
- [38] SKYTRON. Making Sense of RFID/RTLS, 2011.
- [39] A. Smith et al. Tracking Moving Devices with the Cricket Location System. In *MobiSys*, 2004.
- [40] P. Stoica et al. *Spectral Analysis of Signals*. Prentice Hall, 2005.
- [41] C. Swedberg. MojiX announces the availability of its next-generation rfid system, 2012.
- [42] ThingMagic. Why Use RFID. www.thingmagic.com/.
- [43] D. Tse and P. Vishwanath. *Fundamentals of Wireless Communications*. Cambridge University Press, 2005.
- [44] Y. Wakuda et al. An adaptive map-matching based on dynamic time warping for pedestrian positioning using network map.
- [45] J. Xiong and K. Jamieson. Arraytrack: A fine-grained indoor location system. *NSDI '13*.

APPENDIX

Calibrating the Baselines: We calibrate the behaviors of the RSSI and AoA based schemes in both LOS and NLOS through a set of experiments using 30 reference tags spaced by 15 cm. The setup details are described in §10. For RSSI, we compare using the SAR measurements and using three distributed receive antennas. The positioning of the antennas is also described in §10. For AoA, we compare using SAR and using three pairs of antennas at three locations (same as in RSSI). When using a pair of antennas to calculate AoA, we adopt the method based on phase differences as in [32, 20]. When using SAR, we tried both the smoothed MUSIC algorithm [45] and standard beamsteering in estimating the AoA, and did not observe a clear advantage of either algorithm.

Fig. 15 shows the output of the calibration experiments across 150 traces with 5 different placements of the antennas. In both LOS and NLOS, the RSSI-based scheme performs better when using three locations. This is because the value of the SAR lies in the fine-grained spatial angle information, which the RSSI-based scheme does not exploit. In contrast, the AoA-based scheme benefits more from the SAR measurements than from measuring at three distributed locations. This is expected because the moving antenna's SAR emulates a large number of antennas (i.e., 20), which enables a higher accuracy in AoA estimation.