RF-IDraw: Virtual Touch Screen in the Air Using RF Signals

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

ABSTRACT

Prior work in RF-based positioning has mainly focused on discovering the absolute location of an RF source, where state-of-the-art systems can achieve an accuracy on the order of tens of centimeters using a large number of antennas. However, many applications in gaming and gesture based interface see more benefits in knowing the detailed shape of a motion. Such trajectory tracing requires a resolution several fold higher than what existing RF-based positioning systems can offer.

This paper shows that one can provide a dramatic increase in trajectory tracing accuracy, even with a small number of antennas. The key enabler for our design is a multi-resolution positioning technique that exploits an intrinsic tradeoff between improving the resolution and resolving ambiguity in the location of the RF source. The unique property of this design is its ability to precisely reconstruct the minute details in the trajectory shape, even when the absolute position might have an offset. We built a prototype of our design with commercial off-the-shelf RFID readers and tags and used it to enable a virtual touch screen, which allows a user to interact with a desired computing device by gesturing or writing her commands in the air, where each letter is only a few centimeters wide.

1. INTRODUCTION

RF-based positioning has become the next frontier for innovation in mobile computing, business analytics, and human-computer interaction [16, 29, 17]. The topic has attracted much interest from both the industry [30, 14, 23, 3] and the research community, which translated to many advanced RF-localization systems [41, 19, 39]. So far, however, the literature has mainly focused on the problem of discovering the exact location of an RF source on a building floor [41, 19, 13, 28]. Many applications, on the other hand, care more about the detailed trajectory of a target as opposed to its exact location; for example, in gesture-based user interfaces, it is more important to be able to precisely track the shape of a gesture with high resolution and fidelity, while the actual position may still have an offset. For such applications, existing schemes cannot reproduce an accurate version of the trajectory shape to satisfy the applications' needs.

This paper introduces RF-IDraw, a system that can accurately trace the trajectory of an RF source, particularly an RFID. RF-IDraw's trajectory tracing is so accurate that it enables a *virtual touch screen* based on RF signals. Today, a user can write, scroll, or swipe on a touch screen of a smart phone or tablet. Taking this a step further, RF-IDraw allows a user to input her commands by
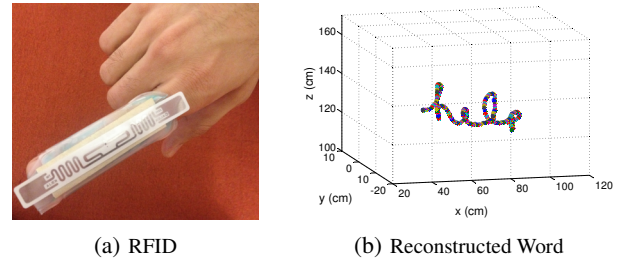
(a) RFID

(b) Reconstructed Word

Figure 1—Enabling Virtual Touch Screen: By accurately tracing the trajectory shape of an RFID on the user's finger, RF-IDraw can transform any plane or surface into a virtual touch screen, allowing the user to input her commands in the air using RF signals.
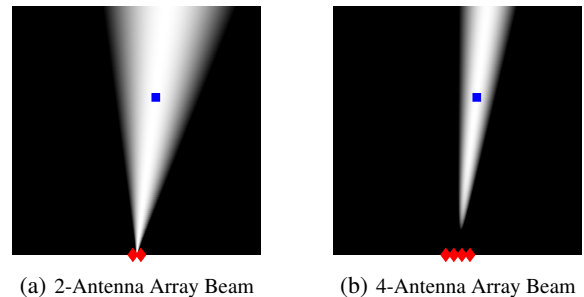

(a) 2-Antenna Array Beam

(b) 4-Antenna Array Beam

Figure 2—Antenna Array Beam Resolution: The figure shows an RF source in blue, and the beams of two antenna arrays with 2 and 4 antenna elements each. The more antennas in the array, the narrower its beam, and the tighter it can bound the source direction.

writing, scrolling, swiping, etc., without being in physical touch with a screen – the user can write in the air any word or command using an RFID attached to a pen or a finger splint, as shown in Fig. 1(a). RF-IDraw would reconstruct the RFID's trajectory and interpret the user's writing and gestures as input to the desired computing device. Essentially, RF-IDraw can transform any plane or surface in space into a virtual touch screen. Fig. 1(b) shows an example of RF-IDraw's output which was entered by writing in the air using an RFID on the user's finger. Such a virtual touch screen can be used to realize a variety of applications. For example, it can be used to interact with a remote screen, to send commands to a cell phone without touching it, or to interface with small devices (e.g., sensors) that do not have space for a keyboard.

RF-IDraw's technology is based on the realization that a different design principle for leveraging multiple antennas can lead to a significant improvement in both tracing and localization accuracy. Specifically, state-of-the-art RF-based positioning systems typically use an antenna array, and leverage its beamsteering capability to detect the direction of the source [41, 39, 12, 21]. The location of the source can then be computed by intersecting the beams of multiple such arrays. Hence, to obtain a high accuracy, they need an array with a narrow beam, which requires a large number of antennas. For example, Fig. 2 compares the beam width of 2-antenna and 4-antenna arrays; both have the default antenna spacing of $\frac{\lambda}{2}$,
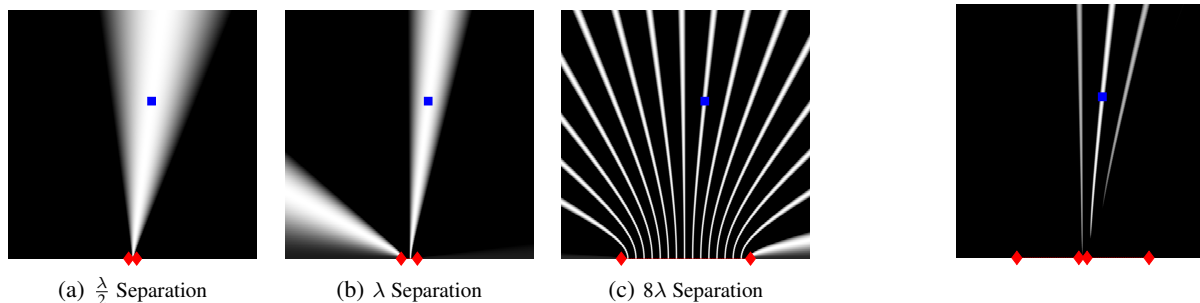
(a) $\frac{\lambda}{2}$ Separation   (b) λ Separation   (c) 8λ Separation

Figure 3—Tradeoff Between Improving Resolution and Removing Ambiguity: As the separation of the antenna pair (marked in red) increases, the number of beams increases accordingly, causing ambiguity in localizing the source (marked in blue). On the other hand, each beam gets narrower, leading to a higher resolution.

Figure 4—Multi-Resolution Design: The wide beam in Fig. 3(a) acts as a filter on Fig. 3(c), removing ambiguity while maintaining the high resolution.

where λ is the wavelength. Clearly, the more antennas on the array, the narrower its beam and the higher the resolution it provides in identifying the direction of the source. Thus, a standard approach in RF positioning systems is to increase the number of antennas in the array in order to achieve higher accuracy [41, 39].

In contrast, in designing RF-IDraw, we realize that for any number of antennas, there is always a tradeoff between *resolution* and *unambiguity*. Specifically, let us fix the number of antennas to be 2, and consider the beam patterns of this pair of antennas when they are separated by $\frac{\lambda}{2}$, λ, and 8λ respectively, as shown in Fig. 3(a), (b) and (c). As we can see, when the antenna separation is $\frac{\lambda}{2}$, the array has a single beam; yet the beam is wide and does not pinpoint the direction of the source, i.e., low resolution. As the antenna separation increases, the beam becomes narrower and more tightly bounds the direction of the source. However, instead of a single beam, we start to see multiple beams, despite the fact that there is a single direction along which the actual signal arrives, which causes ambiguity. Thus, there is a tradeoff in antenna spacing, between localizing with high resolution and localizing with no ambiguity.

Past work uses arrays with a maximum antenna spacing of $\lambda/2$ [41, 39, 19], hence avoiding ambiguity, but also sacrificing the potential of higher resolution. In contrast, in RF-IDraw, we leverage our understanding of the structure of the ambiguity to increase localization accuracy. In particular, RF-IDraw uses a few antenna pairs with different separations. The pairs with larger separation have very narrow beams and hence define the resolution of the positioning system. The pairs with small separation operate as filters to eliminate the ambiguity introduced by the widely separated antenna pairs, while maintaining the high resolution. Fig. 4 shows the result after applying the wide beam in Fig. 3(a) as a filter on Fig. 3(c). As we can see, most of the unintended beams have been filtered out and there is one distinctive narrow beam (whiter means higher likelihood) that bounds the correct direction of the source. Note that, both Fig. 2(b) and Fig. 4 are produced using a total of 4 antennas, yet the latter offers significantly higher resolution (i.e., much narrower beam) than the standard 4-antenna array, and thus demonstrates a more effective way of arranging the antennas.

A key characteristic of RF-IDraw's design is its ability to trace the detailed shape of an RF source's trajectory, even when the actual position has some offset. The tracing accuracy is due to the use of the fine beams in Fig. 3(c). Recall that tracking with an antenna array means that the beam of the array rotates to follow the source direction. When we have multiple beams as in Fig. 3(c), the correct beam rotates to track the source, but the wrong beams also rotate with it. Thus, even if one mistook the correct beam, the shape of the trajectory would continue to match the performed gesture, as further explained in §4.

We have built a prototype of RF-IDraw using commercial RFID readers and used it to track off-the-shelf UHF RFIDs. Further, we integrated RF-IDraw with the handwriting recognition functionality in the MyScript Stylus [36] Android app to evaluate RF-IDraw's function as a virtual touch screen. We run our experiments with five users in both line-of-sight and non-line-of-sight, 2–5 meters away from the reader antennas.[1] Our experiments lead to the following findings:

- RF-IDraw can track the detailed trajectory shape of the users' writing with a median accuracy of 3.7 cm, 11× more accurate than the state-of-the-art antenna array based technique using the same number of antennas. Note that, this accuracy number does not fully capture the power of RF-IDraw's trajectory tracing. In particular, even with a median 3.7 cm error on points along the trajectory, the shape of the overall trajectory is still preserved, because this error only reflects the coherent stretching or squeezing of the trajectory. Therefore, the handwritten characters reconstructed by RF-IDraw are correctly recognized by MyScript Stylus in 97.5% of the cases, and the words reconstructed by RF-IDraw are correctly recognized in 92% of the cases. In contrast, for the trajectories reproduced by the antenna array based technique using the same number of antennas as RF-IDraw, the character recognition success rate is less than 4%, which is equivalent to a random guess, and the word recognition success rate is 0%.
- RF-IDraw's trajectory tracing accuracy holds in non-line-of-sight scenarios, where the median accuracy is 4.9 cm, and the character recognition success rate is 96.8%.
- Even when identifying the exact location, as opposed to tracing, RF-IDraw's accuracy is 2.2× higher than the baseline that uses antenna arrays with the same number of antennas as RF-IDraw.

To our knowledge, RF-IDraw is the first RF-based system that can transform any plane or surface into a virtual touch screen, which opens up an entirely new range of applications in user interaction interfaces.

2. RELATED WORK

Prior work on RF-based positioning has primarily relied on RSSI (Received Signal Strength) [13, 28] or AoA (Angle of Arrival) information [41, 19, 22]. State-of-the-art systems use antenna arrays or synthetic aperture radar (SAR) to extract AoA of an RF signal and can achieve a positioning accuracy on the order of tens of centimeters [41, 19]. Using a dense, surveyed grid of reference sources,

---

[1]Unlike WiFi based localization work that aims to locate users on an office floor (e.g., [41]), RF-IDraw operates at a shorter range since it is based on RFID. §9.3 discusses the prospect of applying RF-IDraw's techniques to WiFi systems.

several schemes can achieve a higher accuracy of a few centimeters by identifying the nearest references [39, 37]. For example, the work in [37] leverages the motion of a robot equipped with reference RFIDs to enable centimeter-scale accuracy in grasping an object tagged with RFIDs. RF-IDraw differs from these past schemes in both techniques and capabilities. By effectively exploiting the high resolution of antenna pairs with large separations, RF-IDraw offers unmatched accuracy in tracking an RF device's detailed trajectory without the use of references. As such, it enables a whole new class of applications which are previously infeasible using RF signals, such as virtual touch screen in the air.

The conception and design of RF-IDraw are inspired by astronomical interferometry, where telescopes are used to image the sky and search for planets [20, 34, 2]. In particular, in astronomy, pairs of telescopes with large separation are used to produce high-resolution fringes. One can consider these telescopes as antennas. The narrow beams produced by RF-IDraw's antenna pairs with large separation and the fringes in interferometry are similar in nature, i.e., both offer high resolution at the cost of ambiguity. Ambiguity in interferometry is resolved using delay lines which effectively orient each telescope pair towards a particular part of the sky, whereas in RF-IDraw, we use a pair of antennas with small separation to focus on a particular region in the area of interest. Furthermore, in astronomy, the rotation of the Earth/sky is exploited to facilitate better coverage and accuracy. In RF-IDraw, although we do not have a known motion like the rotation of the Earth to leverage, we integrate information gathered throughout the trajectory of the RF device to improve the elimination of ambiguity, which has a similar flavor. However, while the underlying intuition is transferable between the two, RF-IDraw's algorithms, signal processing techniques, and applications significantly differ from astronomical interferometry.

In the context of RF antenna arrays, recently there is a growing interest in exploring the use of sparse arrays to estimate the angle of arrival [31, 25, 18]. For example, [31, 25] propose the use of co-prime sampling in a large uniform linear array to reduce the number of antennas needed; [18] evaluates the use of compressive sensing in sampling the antenna positions. By virtue of emulating very large arrays, these schemes show asymptotic improvement over the naive AoA approach through theoretical analyses. However, for a small number of antennas available to our application of interest (i.e., 8 in total), the asymptotic analysis does not lead to any meaningful gain, and hence these proposals cannot achieve the resolution enabled by RF-IDraw. Furthermore, they focus on estimating a single measurement of angle of arrival, as opposed to tracing the trajectory shape as RF-IDraw does.

Systems such as [7] and [32] make use of depth sensors (e.g., Kinect) and infrared cameras (e.g., Wii) to turn a projector screen or wall into a touch screen and allow a user to interact with the display using a specially designed pen. RF-IDraw is the first RF-based virtual touch screen; unlike solutions based on depth imaging or infrared, it does not require line-of-sight to work. Further, since RF sources have unique IDs (e.g., RFID EPC ID [15]), it is easy to scale to a larger number of users simultaneously interacting through the virtual touch screen without causing confusion.

Finally, RF-IDraw's application is inspired by recent work on motion tracking [27, 10, 9] which uses RF signals to enable a user to interact with the environment. Differing from these systems, RF-IDraw is the first RF-based solution that can accurately reconstruct the detailed trajectory of a user's writing or gesturing in the air, where each letter or gesture is only a few centimeters wide. Such capability is not supported by prior work in RF-based gesture recognition. For example, [27] presents a state-of-the-art WiFi-based in-
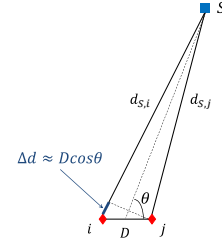

Figure 5—Angle of Arrival at Antenna Pair: Based on the signal phase difference measured between a pair of antennas, one can estimate the spatial direction along which the source's signal arrives.

terface, yet it only supports the detection and classification of a predefined set of nine gestures. Further, past work on RF-based gesture recognition requires modified hardware [27, 10, 9]. In contrast, RF-IDraw is a software patch that can be added to today's standard RFID readers, and does not require any hardware modification.

3. MULTI-RESOLUTION POSITIONING

3.1 Primer

In RF-based positioning, the Angle of Arrival (AoA) of an RF source is computed by comparing the phases of the received signals at multiple antennas [24]. Underlying the AoA computation is the widely known principle that the phase of an RF signal rotates by 2π for every λ distance the signal travels [35], where λ is the wavelength. Specifically, let us consider a signal source at position S, and a pair of receive antennas i and j separated by D in Fig. 5. $d_{S,i}$ and $d_{S,j}$ denote the distances from S to the two antennas respectively, and ϕ_i and ϕ_j are the phases of the received signal that we measure at the two antennas. $\phi_i, \phi_j \in [0, 2\pi)$. The distances and the received signal phases have the following known relations due to the phase rotation [35]:

$$\phi_i = -mod(\frac{2\pi}{\lambda}d_{S,i}, 2\pi)$$
$$\phi_j = -mod(\frac{2\pi}{\lambda}d_{S,j}, 2\pi) \tag{1}$$

Hence, the phase difference between the received signals at the two antennas, $\Delta\phi_{j,i} = \phi_j - \phi_i$, relates to the difference in their distances from the source, $\Delta d_{i,j} = d_{S,i} - d_{S,j}$, as follows:

$$\frac{\Delta d_{i,j}}{\lambda} = \frac{\Delta\phi_{j,i}}{2\pi} + k, \tag{2}$$

where k can be any integer in $[-\frac{D}{\lambda} - \frac{\Delta\phi_{j,i}}{2\pi}, \frac{D}{\lambda} - \frac{\Delta\phi_{j,i}}{2\pi}]$.

When S is relatively far from the antenna pair, as Fig. 5 shows, $\frac{\Delta d_{i,j}}{\lambda}$ can be approximated as $\frac{D\cos\theta}{\lambda}$ where θ is the angle of arrival. Thus, in this case we can rewrite Eq. 2 as

$$\frac{D\cos\theta}{\lambda} = \frac{\Delta\phi_{j,i}}{2\pi} + k \tag{3}$$

We note that Equations 2 and 3 are the same except for the approximation $\frac{\Delta d_{i,j}}{\lambda} \approx \frac{D\cos\theta}{\lambda}$, which is satisfied when the receive antennas are at some distance from the source. Eq. 3 is more intuitive since it directly articulates the spatial angle to the source, θ. Hence, we use this form in the rest of the section for our explanation. However, Eq. 2 is more accurate at close distances from the source and could be used directly in the implementation to avoid unnecessary approximation noise. Indeed, the precise formulation in Eq. 2 is a standard equation representing one or a set of hyperbolas. When the distance is large, a hyperbola reduces to a ray pointing towards a specific spatial direction.

3.2 Grating Lobes

There is a tradeoff between increasing the resolution and removing the ambiguity in detecting a signal's angle of arrival. As Fig. 3 shows, when the separation between the two antennas is greater than $\lambda/2$, the beam pattern exhibits multiple beams, although there is only one direction along which the signal arrives. The additional beams are often referred to as *grating lobes* [24].

Here we explain the cause and characteristics of grating lobes in the context of angle of arrival. For a particular phase difference measured $\Delta\phi_{j,i}$, there are one or a set of spatial angles θ that satisfy Eq. 3. Specifically, since $\cos\theta \in [-1, 1]$, when $D \leq \lambda/2$, k can take only one value which is 0. Thus, in this case, we can derive a unique angle of arrival $\theta = \arccos(\frac{\lambda}{D}\frac{\Delta\phi_{j,i}}{2\pi})$, which means we will observe a single beam in the antenna pair's beam pattern.

As D increases, the number of possible k values increases. For $D = K\frac{\lambda}{2}$, the number of possible values k can take is K. Hence, there are K different values of θ that will satisfy Eq. 3: $\theta = \arccos(\frac{\lambda}{D}\frac{\Delta\phi_{j,i}}{2\pi} + \frac{k\lambda}{D})$, only one of which refers to the actual angle of arrival. This is where ambiguity arises. We will see multiple beams (i.e., grating lobes) in the antenna pair's beam pattern corresponding to the multiple spatial angles. In summary, each grating lobe corresponds to one spatial direction θ, and the number of grating lobes increases linearly with the antenna separation D.

3.3 Resolution and Robustness to Noise

Since grating lobes cause ambiguity in deciding the actual angle of arrival, they are typically considered detrimental and past RF localization systems try to avoid them by imposing the constraint that adjacent antennas must be separated by no more than $\frac{\lambda}{2}$ [41, 39]. However, while grating lobes introduce ambiguity, there are two properties about grating lobes that are highly desirable in the context of localization: high resolution and robustness to noise.

Resolution: Let us first look at the resolution of the spatial angle of arrival as computed by an antenna pair with D separation. We can rewrite Eq. 3 as

$$\cos\theta = \frac{\lambda}{D}\frac{\Delta\phi_{j,i}}{2\pi} + \frac{k\lambda}{D}. \tag{4}$$

Recall that the value we can measure is the signal's phase difference, $\Delta\phi_{j,i}$, and the value we intend to compute is the signal's spatial angle of arrival θ. Any hardware has some resolution, δ, for how it expresses $\phi_{j,i}$ and hence $\Delta\phi_{j,i}$. The value of δ defines the minimum measurable change in $\Delta\phi_{j,i}$, and hence the finest quantization of $\cos\theta$, i.e., its resolution. In particular, the finest quantization we have in computing $\cos\theta$ is $\frac{\lambda}{D}\frac{\delta}{2\pi}$. As the antenna separation D increases (i.e., more grating lobes), the minimum quantization level for expressing $\cos\theta$ decreases, leading to a finer resolution in estimating the spatial angle θ.

Robustness to Noise: In addition to providing high resolution, grating lobes of antenna pairs with a large separation also prove to be more robust to wireless noise. Eq. 4 represents the noise-free case. Now let us take into consideration a phase noise term ϕ_n in the measured phase difference:

$$\cos\theta = \frac{\lambda}{D}\frac{\Delta\phi_{j,i} + \phi_n}{2\pi} + \frac{k\lambda}{D} \tag{5}$$

Eq. 5 shows that as D increases, $\cos\theta$'s sensitivity to phase noise ϕ_n decreases linearly. To put this into perspective, when the phase difference measurement has a $\phi_n = \frac{\pi}{5}$ noise, that translates into 0.2 additive error in $\cos\theta$ if $D = \frac{\lambda}{2}$. But if we increase the separation to $D = 8\lambda$, the same signal phase noise of $\phi_n = \frac{\pi}{5}$ would only lead to 0.0125 additive error in $\cos\theta$, which is minimal. Thus, the larger

the antenna pair separation is, the less effect wireless noise has on the spatial angle of arrival, i.e., the more resilient the system is.

3.4 Intersecting Grating Lobes for Positioning

The discussion above shows that grating lobes of a widely separated antenna pair offer high resolution and noise resilience in detecting the spatial angle of a signal source. To actually localize an RF source, we need to translate knowledge about the spatial angle of the source into information about the source position. To do so, we can intersect the grating lobes of multiple antenna pairs.

Consider the example in Fig. 6(a) with antennas 1, 2, 3, and 4 placed at the four corners of a square of size $8\lambda \times 8\lambda$.[2] The goal here is to localize the RF source marked in blue based on the phases measured at the antennas. The 4 antennas form a total of 6 pairs as indicated by the red lines in Fig. 6(a). Based on the phase difference measured at each antenna pair, we can obtain a set of grating lobes for each antenna pair. We can then intersect these grating lobes to identify candidate point locations for the source as indicated by the white dots in Fig. 6(a). One of these candidate locations matches the exact location of the source (the white dot is covered by the blue dot in the figure).

3.5 Resolving Ambiguity

Although each dot in Fig. 6(a) is very fine (i.e., high resolution), there are a number of them and we cannot tell which one includes the actual position of the source. The challenge is to resolve the ambiguity in positioning while maintaining the high resolution.

Multi-Resolution Filtering: One interesting observation is that in Fig. 6(a), the intersections (i.e., white dots) are fairly sparse, i.e., they are far away from each other. This means that if we can apply a *coarse* spatial filter to them, we will be able to eliminate the undesired intersections and resolve ambiguity.

Thus, our design of RF-IDraw adopts a multi-resolution approach – we create coarse spatial filters by leveraging the unique, low-resolution beams of antenna pairs with small separation to resolve the position ambiguity caused by the high-resolution grating lobes of the widely-separated antenna pairs. Specifically, in addition to the widely spaced antennas, we use a second set of antenna pairs that are spaced by $\lambda/2$. Our implementation uses four additional antennas organized in two pairs <5,6> and <7,8>. Since the spacing is $\lambda/2$, each pair produces one unique wide beam. Intersecting the two wide beams leads to a coarse spatial filter indicated as the white region in Fig. 6(b). We can further refine this filter to Fig. 6(c) by also utilizing all other antenna pairs, namely <5,7>, <5,8>, <6,7>, and <6,8>. Finally, when we overlay Fig. 6(c) as a filter on Fig. 6(a), ambiguity is resolved and the correct position of the RF source is uncovered, as shown in Fig. 6(d) (again, the white dot is covered by the blue mark representing the source).

This demonstrates how RF-IDraw's multi-resolution design leverages the unambiguous coarse filter produced by antenna pairs with small separation to effectively reduce the ambiguity caused by the high-resolution grating lobes.

We note that our implementation of RF-IDraw uses two RFID readers each equipped with 4 antennas. The first reader is used for the widely spaced antenna pairs, while the second reader is used for the tightly spaced antenna pairs. We only measure the AoA using antenna pairs that belong to the same reader. While using antenna pairs across the two readers could further improve the resolution of the system, in practice that would add complexity and potential

---

[2]Today's commercial RFID reader typically has four antenna ports and there is no phase offset between the received signals on these ports. Hence, one can properly compare the phases measured at any pair of antennas on the same reader.
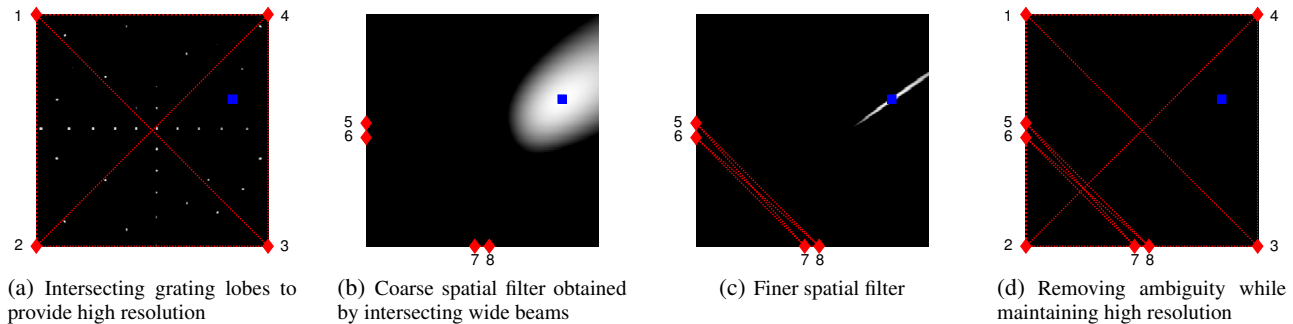
(a) Intersecting grating lobes to provide high resolution

(b) Coarse spatial filter obtained by intersecting wide beams

(c) Finer spatial filter

(d) Removing ambiguity while maintaining high resolution

Figure 6—Multi-Resolution Positioning: This figure shows an example of using RF-IDraw's multi-resolution positioning to localize a signal source marked in blue. The red dotted lines represent which antenna pairs are used in each subplot. (a) shows the common intersections of different antenna pairs' grating lobes, which offer high resolution in positioning yet causing ambiguity. (b) shows a coarse spatial filter formed by intersecting the wide beams of two tightly spaced antenna pairs. (c) shows the finer filter obtained when four more antenna pairs with larger separation are used to refine (b). Applying the filter in (c) on (a) eliminates ambiguity and uncovers the correct position as (d) shows.

errors because the phase offset between the two readers will need to be calibrated or removed. Hence, in our implementation, we only use antenna pairs within the same reader.

4. THE POWER OF GRATING LOBES FOR TRAJECTORY TRACING

Instead of just localizing a static device, many applications in gaming, smart homes, and healthcare are more interested in tracking the trajectory of an RF source as it moves. The unique property of RF-IDraw's grating lobe based approach is its ability to accurately detect the shape of a trajectory. In fact, this is true even in the case where there are errors in the absolute positioning along the trajectory. For example, one may get the initial point of where the trajectory starts wrong by a small offset, yet the shape of RF-IDraw's reconstructed trajectory will still match the actual shape of the trajectory with very high fidelity. This property is desirable for many applications, such as gesture recognition and virtual touch screen. In these applications, it is important to recognize the gesture or writing of the user, while a small offset in the exact location of the gesture is tolerable.

To provide an intuition for why RF-IDraw's design has this property, let us start with the basic case of a standard antenna array. Consider the scenario where the RF source is moving along a continuous trajectory. In the case of an antenna array, tracing the trajectory of this source means that the antenna array's beam will rotate accordingly with the movement of the source such that it keeps pointing towards the source. Now consider what happens when we have grating lobes. The actual source will be on one of these grating lobes. As the source moves, the correct grating lobe will rotate such that it continues to track the source direction. However, since all grating lobes rotate together, they will all follow the movement of the source. Hence, even if one makes the wrong assumption that the source lies on a particular grating lobe which differs from the correct one, this wrong grating lobe's motion (i.e., rotation in particular) does not fundamentally differ from the motion of the correct grating lobe as it tracks the source direction.

To illustrate this property, Fig. 7(a) shows what happens when we pick the correct grating lobe and also for the cases where we pick wrong grating lobes adjacent to the correct one. Specifically, this figure shows a trajectory that follows the handwriting of the letter 'q', the ground truth of which is shown in black. Fig. 7(a) shows the reconstructed trajectories for nine different sets of grating lobes. The trajectory in the center is reconstructed by tracking the motion of the correct set of grating lobes. The other eight trajectories are reconstructed by using wrong grating lobes adjacent to the correct

ones. As we can see, although the absolute positions have an offset, the shapes of the trajectories remain accurate and the letter 'q' can still be easily recognized. This demonstrates the point that even if we end up tracking the motion of a wrong grating lobe, which is not too far away from the correct one, the shape of the reconstructed trajectory will still remain similar to the ground truth.

It is important to note that, the further the correct grating lobe is from the one we pick, the more distortion there is in the shape of the reconstructed trajectory. The reason is as follows. While all grating lobes rotate together, (i.e., if the correct one turns with the source, the others will turn accordingly), they do not have exactly the same forms. Let us refer to the grating lobes in Fig. 3(c) again to better understand this point. There we can see that the grating lobes close to each other tend to have similar forms, while further apart grating lobes' forms are more different. Thus, although all grating lobes rotate together, their form differences will introduce distortion to the shape of the reconstructed trajectory, and this distortion increases as we pick a grating lobe further away from the correct one. For instance, Fig. 7(b) shows the reconstructed trajectory of the handwritten 'q' by using a set of grating lobes far away from the correct ones. As we can see, the shape of this reconstructed trajectory is less similar to the ground truth than the trajectories in Fig. 7(a).

This emphasizes the importance of using the multi-resolution system in §3. Although it may not completely eliminate the ambiguity among grating lobes in certain cases, it helps confine us to a region very close to the correct set of grating lobes, which ensures good shape fidelity.

5. ALGORITHM

In this section, we describe the algorithm behind RF-IDraw's multi-resolution positioning and trajectory tracing.

5.1 Multi-Resolution Positioning Algorithm

RF-IDraw's multi-resolution positioning design is implemented as a two-stage voting algorithm to identify the position of an RF source. In stage 1, each antenna pair with small separation casts one vote on each point in space, based on its belief that the RF source is at that position. Then, the points with high total votes from all antenna pairs with small separation will form the candidate region to be used in stage 2. In stage 2, each antenna pair with large separation casts one vote on each point within the candidate region, based on its belief that the RF source is at this position. Then the points with higher total votes from all antenna pairs will be considered as the more likely positions for the RF source.

Now the only question is how to define the vote. First, let us consider the vote cast by each antenna pair with $\frac{\lambda}{2}$ separation. Recall
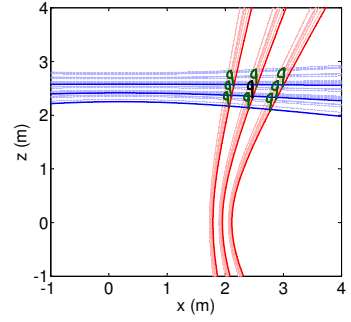
(a) Reconstructed trajectories using grating lobes adjacent to the correct ones
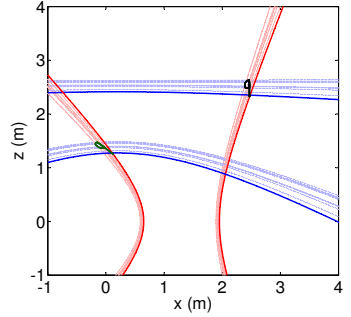
(b) Reconstructed trajectories using grating lobes far away from the correct ones

Figure 7—Effect of Choosing Wrong Grating Lobes in Trajectory Tracing: The ground truth trajectory is a handwritten 'q' as shown in black. When we pick a wrong grating lobe adjacent to the correct one to track, the reconstructed trajectory has a very similar shape as the actual shape, despite the absolute position offset, as (a) shows. However, if we end up tracking a wrong set of grating lobes far away from the correct ones, the distortion in the shape will be large, rendering the letter difficult to recognize, as (b) shows.

that an antenna pair with $\frac{\lambda}{2}$ separation has one beam. We let the antenna pair vote on a point in space based on how far the point is from the antenna pair's beam center. In particular, the vote is designed to be less or equal to 0. When the point lies along the direction of the center of the beam, the vote is 0; as it gets further away from the beam center, the vote becomes lower (i.e., more negative).

For example, consider the beam of an antenna pair in Fig. 8, whose center is shown as the red solid line. This antenna pair's votes on P_1 and P_2 will both be 0 because they lie along the beam center, while P_3 and P_4 will have a lower vote because they are further away from the beam center.

Formally, the direction of antenna pair $<i,j>$'s beam is defined by Eq. 3 with $k = 0$. In the more precise formulation, the center of the beam is defined by Eq. 2 also for $k = 0$. Hence, antenna pair $<i,j>$'s vote on point P is calculated as:

$$V_{i,j}(P) = -\|\frac{\Delta d_{i,j}}{\lambda} - \frac{\Delta \phi_{j,i}}{2\pi}\|^2, \quad (6)$$

when the two antennas are separated by $\frac{\lambda}{2}$.

Now let us extend this vote to the case of antenna pairs with large separation. For an antenna pair with large separation, the only difference is that instead of one beam, it will have multiple grating lobes. Thus, instead of voting based on how far the point is from one beam center, a widely spaced antenna pair votes based on how far the point is from the grating lobe closest to it. Specifically, the vote will be:

$$V_{i,j}(P) = -\min_k \|\frac{\Delta d_{i,j}}{\lambda} - \frac{\Delta \phi_{j,i}}{2\pi} - k\|^2, \quad (7)$$

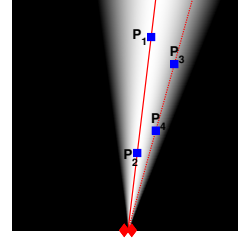**Figure 8—Interpretation of Vote:** Points along the same spatial direction have the same vote from the antenna pair, e.g., P_1 and P_2, P_3 and P_4 have the same votes. Points closer to the beam center have higher votes than the ones further away from the center, e.g., P_1, P_2 have higher votes than P_3 and P_4.

where the minimization over k captures the distance to the closest grating lobe among all grating lobes.

Now we have the votes from all antenna pairs, we combine them to identify the intersections of the antenna pairs' beams or grating lobes. The higher the total vote $V(P) = \sum_{i,j} V_{i,j}(P)$ is, the more likely P is the actual position of the RF source.

5.2 Trajectory Tracing Algorithm

For the purpose of trajectory tracing, recall the discussion in §4 which shows that even if we start from a wrong initial position and track a wrong grating lobe (not too far away from the correct one), the shape of the reconstructed trajectory will still be quite accurate. Aside from this insight, from a slightly different perspective, Fig. 7 also reveals another important point: to achieve shape fidelity, we are better off tracking the continuous rotation of a grating lobe, as opposed to switching grating lobes in the middle of a trajectory.

Based on this idea, we develop RF-IDraw's trajectory tracing algorithm exploring the continuity of the RF source's motion and the continuous rotation of the grating lobes, as described below:

- RF-IDraw selects a few candidate initial positions with the highest total votes as calculated in §5.1 using the initial phase measurements.
- For each candidate initial position, RF-IDraw identifies the grating lobe of each antenna pair that is closest to this position, and keeps tracking the continuous rotation of this grating lobe.
- To estimate the next position, RF-IDraw evaluates the votes for all points within the vicinity of the current position. In computing each antenna pair's vote, RF-IDraw also enforces the rule to rotate with the same grating lobe, which is done by fixing k and unwrapping $\Delta \phi_{i,j}$ in Eq. 7, i.e., ensuring continuity when angles wrap around $[0, 2\pi]$.
- RF-IDraw iteratively estimates the consecutive positions to reconstruct the entire trajectory.
- After having reconstructed one trajectory for each candidate initial position, RF-IDraw picks the one whose sum of votes across all points on the trajectory is the highest as its final estimate.

The reason for choosing the trajectory with the highest overall vote is as follows. Recall that we have an over-constrained system, i.e., there are more antenna pairs than needed to specify a point in space (using their grating lobes). In this over-constrained system, if we start from the correct initial position, the set of grating lobes (one from each widely spaced antenna pair) we track will move coherently and hence always intersect, because they all agree on the actual trajectory of the RF source. This means their total vote will be high always. In contrast, for a wrong initial position, the associated wrong grating lobes will not move coherently, and hence do not consistently intersect throughout the trajectory. Therefore, their total vote will drop later on. In other words, with an over-constrained system, we can detect incohesive rotations of

Figure 9—Commercial UHF RFID Used in Experiments: Alien Squiggle General Purpose UHF RFID

these grating lobes. Thus, by choosing the reconstructed trajectory with the highest overall trajectory vote, RF-IDraw further refines the initial position estimate, as demonstrated in §7.

6. IMPLEMENTATION

We built a prototype of RF-IDraw using commercial off-the-shelf UHF RFID readers to locate and trace the trajectories of EPC Gen-2 [15] UHF RFIDs.

Readers and Software Implementation: We use two ThingMagic M6e 4-port UHF RFID readers [33], each connected with four AN-900LH 900 MHz omni-directional antennas [1]. We program the readers to continuously query the RFIDs using a carrier frequency of 922 MHz and return the signal phase for every RFID reply. We implemented RF-IDraw's multi-resolution positioning and trajectory tracing algorithms in MATLAB and ran them in real-time.

RFIDs: We use the Alien Squiggle tag [11] shown in Fig. 9, which is a commercial off-the-shelf passive UHF RFID widely used in supply chain and asset tracking applications. Each of them costs 5-10 cents. We also experimented with the Omni-ID Exo 800 tag [5] to verify RF-IDraw's design across different types of RFIDs.

Ground Truth: We use the VICON motion capture system [8] to measure the ground truth trajectory of the RFID. In a 5×6 m^2 room instrumented with a grid of infrared cameras, the VICON system can provide sub-centimeter accuracy in tracking an object tagged with infrared reflective markers. For experiments in the VICON room, we measure the ground truth by putting infrared reflective markers on the user's hand, around the RFID. For experiments outside the VICON room in an office lounge, we let the user write along a specific set of trajectories for which the ground truth has been measured manually in advance.

Compared Schemes: We compare RF-IDraw with the state-of-the-art antenna array AoA-based approach [12] using the same number of antennas as RF-IDraw. In particular, both RF-IDraw and the antenna array AoA-based systems employ 8 antennas connected to two RFID readers. RF-IDraw's 8 antennas are deployed as shown in Fig. 6(d). Since RFIDs communicate by backscattering the reader signal [38], the signal phase reading returned by the reader reflects the round trip distance instead of the one-way distance. Hence, each tightly spaced antenna pair (i.e., <5, 6>, <7, 8>) has a separation of $\frac{\lambda}{4}$ (as opposed to $\frac{\lambda}{2}$) to ensure there is only one beam.[3] The widely spaced antenna pairs (i.e., each edge) have separation of 8λ (i.e., 2.6 m). The antenna array AoA-based scheme uses two antenna arrays, each with 4 antennas spaced by $\frac{\lambda}{4}$ (to account for the backscattering round trip also). One antenna array is placed along the left edge of the square in Fig. 6(d), the other is placed along the bottom edge.

Application Evaluation: We interface RF-IDraw with the MyScript Stylus app [36] on an Android phone using the MonkeyRunner Android API [4]. This API allows developers to send specified sequences of events to an Android device. We use this API to convert the reconstructed trajectory of the RFID to touch screen input sequences on an Android phone. The MyScript Stylus app, then, interprets the input as text. We let the users write words

---

[3]The equations in this paper apply to the general case of an RF transmitter. To account for the special case of RFID backscattering, one needs to add a 2 factor to all the $\Delta d_{j,i}$ and $\cos\theta$ in the equations.

randomly sampled from the top 5000 most commonly used words in the Corpus of Contemporary American English [6]. Some examples include: "*play*", "*clear*", "*import*", etc. Then, we evaluate the character and word recognition success rates of the reconstructed trajectories as recognized by the MyScript Stylus app.

7. MICROBENCHMARK

We start with a microbenchmark experiment to provide insights into the working of the system. In particular, to better understand the capabilities and properties of RF-IDraw's trajectory tracing, we let a user write the word "*clear*" using the RFID on his hand, in the 5×6 m^2 VICON room. RF-IDraw uses a total of 8 antennas connected to two RFID readers. All reader antennas are deployed on a wall at one side of the VICON room. The user stands 2 m away. Fig. 6 presents the side view of the room, facing the wall of antenna deployment. Antenna separations and placement are chosen based on the application's needs and room size. Fig. 10(a) shows the ground truth trajectory of the user's handwriting measured by the VICON motion capture system. Now let us see how RF-IDraw's tracing works in practice.

7.1 Granularity of Tracing

In this experiment, RF-IDraw's multi-resolution positioning algorithm (described in §5.1) returned two candidate initial positions. Fig. 10(b) and Fig. 10(c) show the trajectories reconstructed from these positions using RF-IDraw's trajectory tracing algorithm, described in §5.2. As we can see, RF-IDraw is able to trace every minute change in the RFID position during the user's writing. For example, the letter 'e' in the middle is only about 5 cm wide, yet RF-IDraw is able to reproduce the details in its shape by tracking the rotation of the high-resolution grating lobes.

7.2 Choosing the Correct Initial Position

In order to decide which trajectory to choose out of the two (as the final output), as §5.2 describes, RF-IDraw looks at the total vote from all antenna pairs for each position on the trajectory and picks the trajectory whose overall vote (summing up total votes from all antenna pairs for all positions along the trajectory) is the highest. Fig. 10(f) shows the evolution of the total vote from all antenna pairs as the blue and red reconstructed trajectories progress. As we can see, initially, both the blue and the red have high total votes (i.e., close to 0). That is exactly why their starting points have been identified by the multi-resolution positioning algorithm as the candidate initial positions. However, as the two reconstructed trajectories progress, the red one's total vote quickly drops, significantly deviating from 0, while the blue one's vote stays fairly stable and close to 0, having only occasional deviation. This is because the grating lobes of the antenna pairs closest to the red initial position do not consistently agree throughout the trajectory, leading to no common intersection later in the trace, which results in their poor total vote later on. The detection of this is made possible by the over-constrained system where we have more antenna pairs than needed to specify a single point in space, as explained in §5.2.

Comparing the overall votes throughout the trajectories, RF-IDraw picks the blue trajectory as its final output. Fig. 10(d) shows the ground truth, RF-IDraw's final output, and the red trajectory on the same plot. As we can see, the blue trajectory which RF-IDraw picks indeed better matches the ground truth than the red one, and the blue initial position is also closer to the ground truth than the red candidate. This demonstrates RF-IDraw's power in leveraging its trajectory tracing to further refine its positioning.
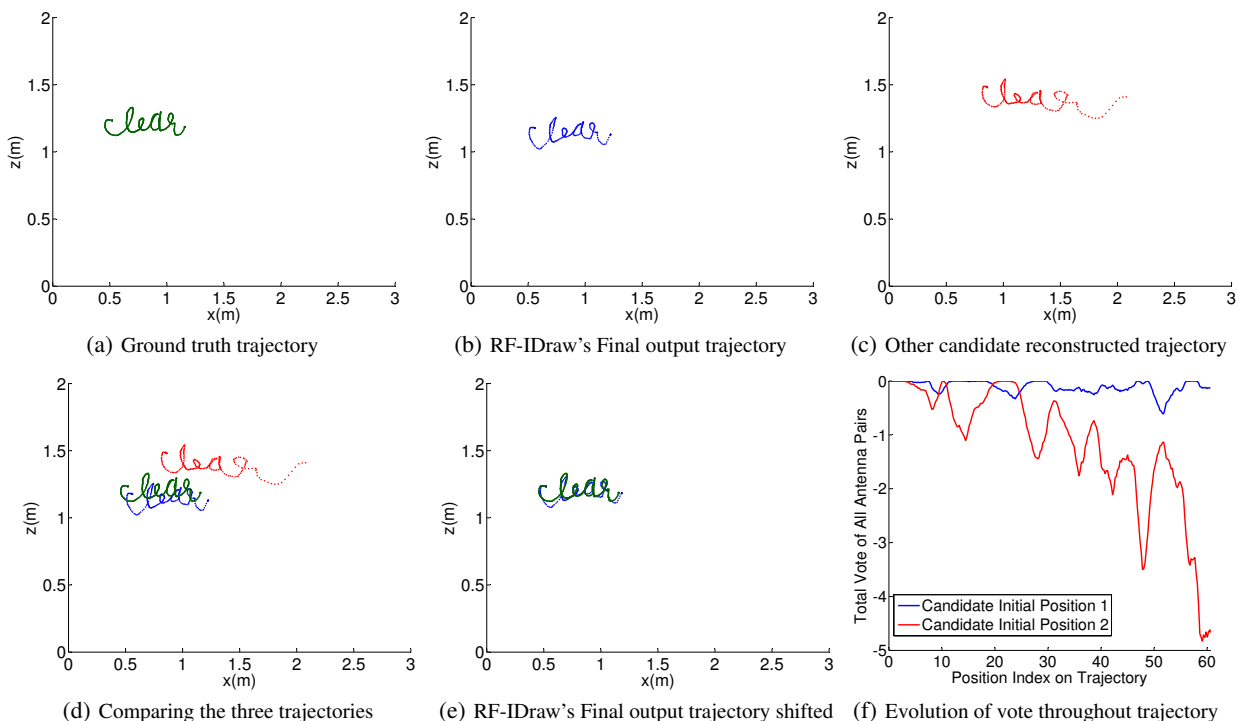
(a) Ground truth trajectory   (b) RF-IDraw's Final output trajectory   (c) Other candidate reconstructed trajectory

(d) Comparing the three trajectories   (e) RF-IDraw's Final output trajectory shifted   (f) Evolution of vote throughout trajectory

Figure 10—Tracing a User's Writing in the Air: (a) shows the actual trajectory of the RFID as the user writes the word "*clear*" in the air, as measured by the infrared VICON motion capture system. (b) and (c) show two trajectories reconstructed by RF-IDraw, starting from two candidate initial positions provided by the multi-resolution positioning algorithm. Both reproduce all the detailed turns in the trajectory shape and can be recognized as the correct word. Based on the overall trajectory votes in (f), RF-IDraw chooses the blue one as its final output, which is indeed closer to the ground truth position, as shown in (d), demonstrating RF-IDraw's capability to refine positioning through trajectory tracing. Finally, (e) shows the comparison between the actual trajectory and RF-IDraw's output after removing the initial offset. The similarity between the two curves proves the unique shape resilience property of RF-IDraw's trajectory tracing.

7.3 Shape Resilience

Finally, these figures also demonstrate the unique property of RF-IDraw's trajectory tracing design, which is that, the *shape* will be well preserved in the reconstructed trajectory, even if the initial position estimate has an offset.

Let us first look at the comparison between the blue reconstructed trajectory and the ground truth. The blue initial position is about 7 cm away from the ground truth initial position.[4] In Fig. 10(e), we show the blue reconstructed trajectory after removing this 7 cm initial position offset from it. As we can see, this shifted version of the reconstructed trajectory closely matches the ground truth, accurately preserving the shape of the writing. Furthermore, even in the red reconstructed trajectory, whose initial position is off by 41 cm, the shape of the word is still roughly preserved.

As §4 explains, shape fidelity can be achieved despite initial position offset because all grating lobes of an antenna pair rotate together, and hence picking a wrong grating lobe's rotation to track only results in a transform of the trajectory. Furthermore, if the grating lobe picked is fairly close to the correct one, the reconstructed shape will be very similar to the correct shape, and the distortion in shape increases as we pick a grating lobe further away from the correct one. Indeed, as compared with the blue reconstructed trajectory, the red one is less similar to the green ground truth's shape, because its initial position is off by more.

This demonstrates the shape resilience property of RF-IDraw's trajectory tracing, while also underscoring the value of finding an initial position reasonably close to the correct one. This property can be particularly useful in applications where recognizing the trajectory of a motion offers high value. In this trace, in fact all three trajectories in Fig. 10(d) have been correctly identified as the word "*clear*" by the handwriting recognition app RF-IDraw interfaces with. In §9, we validate the robustness of RF-IDraw's trajectory tracing through the evaluation of character and word recognition success rates.

8. EVALUATION

We evaluate our prototype of RF-IDraw by having five users write 150 words randomly sampled from the 5000 most frequently used words in the Corpus of Contemporary American English [6]. The users write in the air with a UHF RFID on their hands as shown in Fig. 1(a). The average width of each letter written is around 10 cm. We conduct experiments in both LOS and NLOS settings, in the 5×6 m^2 VICON room as well as a large office lounge of size 8×12 m^2 divided into cubicles.

We compare the performance of RF-IDraw with state-of-the-art antenna array based approach [12] using the same number of antennas. Both RF-IDraw and the antenna array based system employ 8 antennas connected to two RFID readers. The details of the antenna setup for both schemes are described in §6. The users stand 2–5 meters away from the reader antennas in our experiments.[5] RF-

---

[4]Our further analysis reveals that, given the phase difference measurements, the grating lobes intersecting at the blue initial position are indeed the ones closest to the ground truth initial position as well. In other words, the initial positioning error is due to the error in the measurements of the signal phases, which are likely caused by random wireless noise and the multipath effect [22, 35].
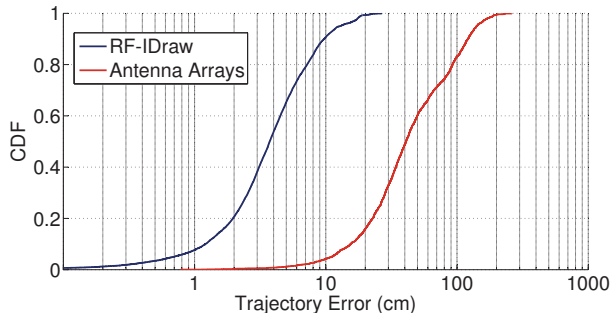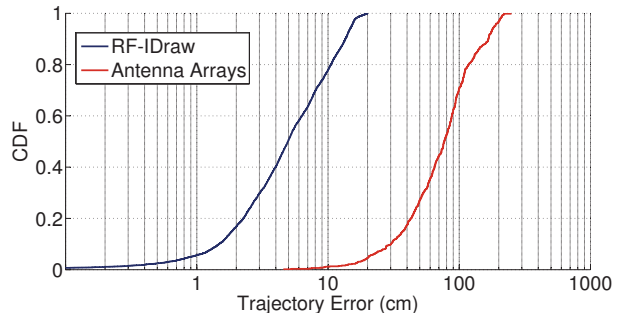
[5]The commercial RFID reader's range limits the range of our current prototype. Beyond 5 meters, we start to see significant message loss because the RFID cannot harvest enough energy to wake up.

(a) CDF of Trajectory Error Distance in LOS

(b) CDF of Trajectory Error Distance in NLOS

Figure 11—CDF of Trajectory Error Distance in LOS and NLOS: The median point-by-point error distance for RF-IDraw's reconstructed trajectory after removing the initial offset is 3.7 cm and 4.9 cm in LOS and NLOS respectively, $11\times$ and $16\times$ lower than the median errors (after removing the offset) for the antenna array based scheme in LOS and NLOS, significantly improving the accuracy of trajectory tracing. Furthermore, the errors for RF-IDraw are due to the coherent stretching and squeezing of the trajectory shape, rather than being random. Hence, despite this error, RF-IDraw's reconstructed trajectories of letters a few centimeters wide can still be successfully recognized in 97.5% of the cases by a handwriting recognition app, as we show in §9.1.

IDraw uses the multi-resolution positioning and trajectory tracing algorithms described in §5 to reconstruct the trajectory and estimate the initial position of the RFID. In the antenna array based system, each 4-antenna array measures an angle of arrival of the RFID, then the beams of the arrays are intersected to estimate the RFID position for each point on the trajectory, reproducing the trajectory.

In this section, we focus on analyzing the trajectory tracing and positioning accuracy of RF-IDraw and the antenna array based scheme. In §9, we further assess the capabilities of the two systems in the virtual touch screen application.

8.1 Trajectory Accuracy

Let us start by evaluating the accuracy of the two systems in reconstructing the RFID trajectory.

Metric: To focus on how much the reconstructed trajectory shape deviates from the actual trajectory shape, we remove a fixed offset between the reconstructed and the ground truth trajectories. The trajectory error for RF-IDraw is computed by removing the initial position offset from its reconstructed trajectory, and calculating the point-by-point position difference between its shifted reconstructed trajectory and the ground truth. For example, the trajectory error for the reconstructed trajectory in Fig. 10(b) would be the point-by-point differences between the two curves in Fig. 10(e). For the antenna array based scheme, since the errors along the trajectory are random and independent from each other, removing the initial position offset results in even larger error distances for other points on the trajectory. Hence, instead we compute the average position difference (a vector in space) between the reconstructed trajectory and the ground truth. We remove this average position difference from the reconstructed trajectory, then calculate the point-by-point differences between the shifted reconstructed trajectory and the ground truth. Note that, removing the average position difference essentially removes the DC offset and thus is favorable to the compared scheme.

Trajectory Accuracy Results in LOS: Fig. 11(a) shows the CDF of the trajectory errors for the two systems in LOS. The median error for RF-IDraw is 3.7 cm (point-by-point difference) and the 90th percentile is 9.7 cm. For the antenna array based system using the same number of antennas, the median error is 40.8 cm and the 90th percentile is 121.1 cm. Hence, RF-IDraw's trajectory tracing is $11\times$ more accurate than the antenna array approach in LOS.

As RFID manufacturers compete to increase the range, it is constantly growing and we expect RF-IDraw to benefit from that trend.

- Note that, the 3.7 cm accuracy does not fully capture the capability of RF-IDraw's trajectory tracing. Specifically, despite this error, RF-IDraw's reconstructed trajectory of letters a few centimeters wide can still be successfully recognized by a simple handwriting recognition app in 97.5% of the cases, as we evaluate in the virtual touch screen application in §9. This is because the point-by-point errors along RF-IDraw's reconstructed trajectory are not independent random errors. Instead, they have a structure and reflect coherent stretching, squeezing, and enlarging in the trajectory shape. Thus, although the trajectory error as defined here is a few centimeters, in most cases, the distortion does not affect the recognition of the letter or word (e.g., Fig. 10(c)).

- RF-IDraw's $11\times$ trajectory accuracy improvement over the antenna array based scheme is enabled by mainly two levers. First, the high-resolution and robustness to noise features of grating lobes allow RF-IDraw to trace the details of the RFID trajectory. Second, RF-IDraw's trajectory tracing algorithm tracks the continuous rotation of the grating lobes, resulting in shape resilience.

- Using the same number of antennas as RF-IDraw, the standard antenna arrays will have much wider beams which fail to provide the resolution or robustness needed to track the shape of a word written by a user in the air. Also, as opposed to the coherent deviations on RF-IDraw's trajectory, the errors the antenna array based scheme has are random and scattered all over the place, resulting in completely unrecognizable trajectories as §9 reveals.

Trajectory Accuracy Results in NLOS: We also evaluate the two systems' performance in non-line-of-sight experiments. The NLOS setup is in an office lounge area (8×12 m^2) divided into cubicles by 2.5 m tall, 20 cm thick separators made of two layers of wood. The reader setup is deployed behind these separators and we ensure that there is no LOS path from any of the reader antennas to the RFID on the user's hand. Similar to the literature on RFIDs [39, 37, 22, 12], we cannot do cross-room evaluation because the RFID reader cannot power an RFID from a different room.

Fig. 11(b) shows the trajectory error for both RF-IDraw and the antenna array based solution in NLOS. The median error of RF-IDraw is 4.9 cm and the 90th percentile is 13.6 cm, similar to LOS. In comparison, the antenna array based system's median trajectory error in NLOS is 76.9 cm and the 90th percentile is 166.7 cm.

- The degradation in trajectory accuracy caused by NLOS is more severe for the antenna array based scheme (40.8 cm in LOS v.s. 76.9 cm in NLOS) than for RF-IDraw (3.7 cm in LOS v.s. 4.9 cm in NLOS)). This is likely due to the different effects multipath [22, 35] has on the two methods. In NLOS, the direct path
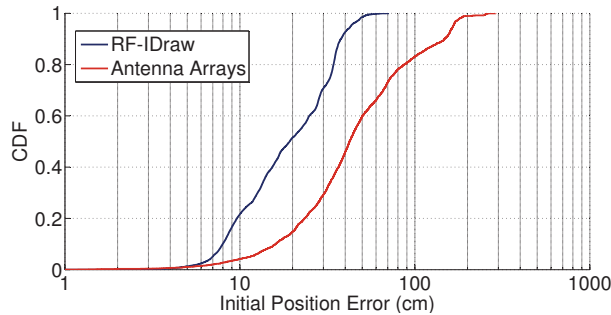
(a) CDF of Initial Position Error Distance in LOS
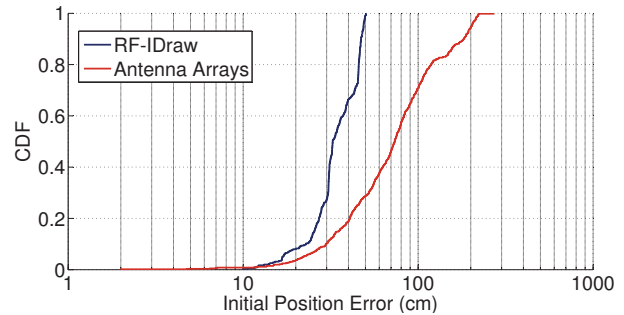
(b) CDF of Initial Position Error Distance in NLOS

Figure 12—CDF of Initial Position Error Distance in LOS and NLOS: The median initial position error of RF-IDraw is 19 cm and 32 cm in LOS and NLOS respectively, 2.2× and 2.3× lower than the median initial position error for the antenna array based scheme in LOS and NLOS. This localization accuracy improvement comes from RF-IDraw's use of trajectory tracing votes to refine its initial position estimate.

is attenuated and the dominant path may differ from the direct path. In this case, estimating the absolute position becomes challenging, as past work has shown [19, 39, 41]. As a result, the antenna arrays' performance significantly worsens. Yet the trajectory shape of RF-IDraw is more robust, because as long as there is a dominant path, the shape will still be discovered by following the dominant path.

In conclusion, RF-IDraw enables accurate trajectory tracing in both LOS and NLOS, outperforming the antenna array based technique using the same number of antennas by 11× and 16× in LOS and NLOS respectively. Note that, while several past RF localization techniques have also achieved centimeter scale accuracy [41, 39, 9], they either require a dense deployment of reference sources and/or a large number of antennas, or GHz of bandwidth. In contrast, RF-IDraw shows that, for the first time, one can trace the trajectory shape of an RF source with centimeter accuracy with a small number of antennas, without reference sources, and using commercial narrowband devices.

8.2 Initial Position Accuracy

Next, we evaluate the initial position estimation accuracy. Fig. 12 shows the CDF for initial position error for RF-IDraw and the antenna array based system. In LOS, the median initial position error of the antenna array based system is 42 cm and its 90th percentile is 148 cm; RF-IDraw's median initial position error is 19 cm, 2.2× better than AoA, and its 90th percentile is 38 cm, 3.9× better than AoA. In NLOS, the accuracy for both systems degrades, because of the attenuation of the direct path from the reader antennas to the RFID as well as the multipath caused by structures in the office lounge. The median initial position error for RF-IDraw in NLOS is 32 cm and the 90th percentile is 47 cm. For the antenna array based scheme, the median and 90th percentile accuracy is 74 cm and 183 cm respectively.

- It is worth noting that, the initial position accuracy for the antenna array based scheme is similar to its trajectory accuracy discussed in the previous section. This is because the antenna array based system estimates each position along the trajectory independently using the same method as it estimates the initial position. Therefore, the errors along the trajectory are random and independent from each other, and hence the trajectory error will be similar to the absolute position error.
- Using the same number of antennas, in both LOS and NLOS, RF-IDraw improves the initial position accuracy over the antenna array based technique by over 2.2×. This improvement stems from the fact that the overall trajectory vote is used to identify the best initial position, which refines RF-IDraw's accuracy, as shown in §7.2.
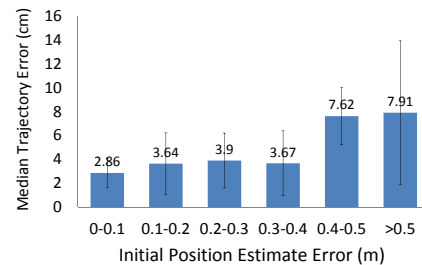

Figure 13—Initial Position Accuracy v.s. Trajectory Accuracy in RF-IDraw: When the initial position picked by RF-IDraw is within 40 cm of the actual initial position, this initial offset has minimal effect on the trajectory accuracy, due to the similarity between closeby grating lobes. If the initial position has a fairly large offset, the trajectory reconstructed will be more distorted, resulting in a higher trajectory error. Yet this higher error is mainly due to the enlarging of certain parts in the trajectory, and hence does not translate into high recognition error (e.g., red curve in Fig. 10(d)).

8.3 The Effect of Initial Position Accuracy on Trajectory Accuracy in RF-IDraw

As we have demonstrated through the microbenchmark experiment in §7, even when the initial position estimate has some offset, the shape of RF-IDraw's reconstructed trajectory is still preserved. Here let us further analyze the relationship between the trajectory accuracy and the initial position accuracy of RF-IDraw. Specifically, for each word written, we remove the initial position offset from RF-IDraw's reconstructed trajectory, then compute the median error for all the positions throughout the trajectory. Then we group the traces according to their initial position error, and study the trajectory error for each initial position error range.

Fig. 13 shows the trajectory error as a function of the initial position error.

- When the initial position error is below 40 cm, RF-IDraw's trajectory accuracy almost remains the same, i.e., around 3 cm. Hence, a reasonably small initial position offset does not affect the shape resilience of RF-IDraw.
- When the initial offset is fairly large, the trajectory error also increases, to a median of around 7–8 cm. To understand why this happens, recall the microbenchmark trace in §7, in particular Fig. 10(d). As we can see, if the initial position is off by a fairly large amount, while the shape of the reconstructed word is recognizable, the end of the trajectory is enlarged. Intuitively, this is due to the fact that we end up tracking a wrong grating lobe further away from the correct one, whose motion has the same trend as the correct one yet leads to a larger distortion in the shape as the RFID moves more. Indeed, in many cases where we observe
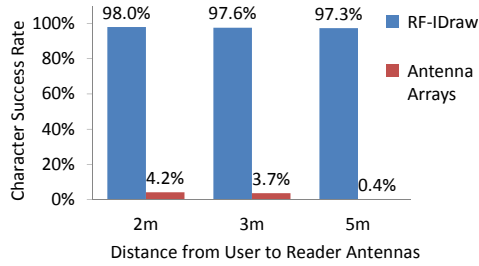
Figure 14—Character Recognition Success Rate: RF-IDraw's reconstructed trajectories for letters can be correctly recognized in 97.5% of the cases by the handwriting recognition Android app, whereas the character trajectories reproduced by the antenna array based system can only be recognized in less than 4% of the cases, equivalent to a random guess.

a high median trajectory error, it is due to the enlarging of the shape towards the end of the trace.

9. VIRTUAL TOUCH SCREEN APPLICATION

We use our prototype of RF-IDraw to demonstrate the feasibility of an RF-based virtual touch screen, which allows a user to input her commands to a desired computing device (e.g., an Android phone) by writing in the air. In particular, we feed each reconstructed trajectory as a set of instructions to an Android phone and the instructions emulate a sequence of touch screen events. Then we let the handwriting recognition function in the MyScript Stylus app [36] interpret it as text. We evaluate success rates for recognizing the character and word trajectories as reconstructed by RF-IDraw and the antenna array based system respectively.

9.1 Character Recognition Success Rates

The average width of a user's handwritten character is around 10 cm. Fig. 14 shows the success rate of correctly recognizing the character, as a function of the distance from the user to the reader antennas. As we can see, the character error rate almost remains the same at around 97%-98% when the user is 2 m, 3 m, and 5 m away from the reader antennas. The overall character recognition success rate for RF-IDraw's reconstructed trajectory is 97.5%, while the success rate for the antenna array based scheme is less than 4%.

The following points are worth noting:

- One may be wondering, with the 3–4 cm median trajectory accuracy as shown in Fig. 11(a) and Fig. 11(b), how can the trajectories reconstructed by RF-IDraw be successfully recognized for characters each of only a few centimeters wide. This is because the errors on the trajectory are not random errors from point to point. Instead, they are mostly due to the transform/distortion (e.g., stretching and squeezing) in the shape of the trajectory, as opposed to independent positioning errors. For example, the blue reconstructed trajectory (after removing initial position offset) in Fig. 10(e) deviates from the ground truth in the lower part of letter "c", resulting in trajectory error of a few centimeters. Yet because it is only a slight distortion instead of random errors, it does not affect the character recognition. In fact, different users write the same letter differently, and hence such type of distortion is naturally taken care of by the handwriting recognition app.
- The errors on the antenna array based scheme's reconstructed trajectory are random and incohesive. As a result, they lead to significant errors in recognizing the letters. In fact, even in the very rare cases where the trajectory is interpreted as the correct letter, the decision is more like a random guess by the software, because even a human could not recognize the letter.
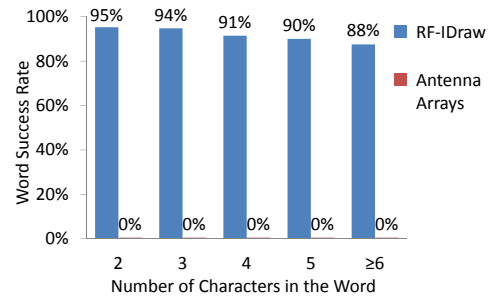

Figure 15—Word Recognition Success Rate: RF-IDraw's reconstructed trajectories for words can be correctly recognized in 92% of the cases by the Android app, enabling an effective virtual touch screen interface, far exceeding the capability of existing RF-based positioning system using the same number of antennas. None of the word trajectories reproduced by the antenna array based scheme is correctly recognized.
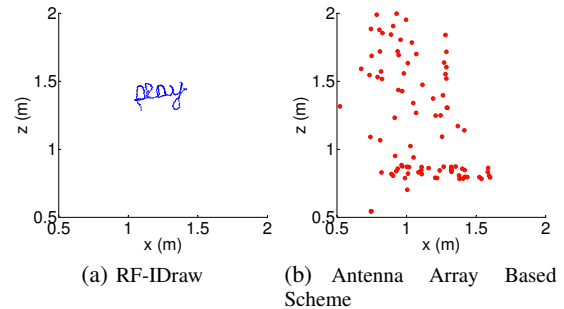

(a) RF-IDraw   (b) Antenna Array Based Scheme

Figure 16—Reconstructed Trajectories of "play" Written 5 m Away: RF-IDraw's reconstructed trajectory reproduces all the details in the user's writing. The antenna array based scheme's reconstructed trajectory is scattered all over the place, due to the antenna arrays' low-resolution beams and their high sensitivity to noise.

- Finally, the character recognition success rate of RF-IDraw's reconstructed trajectories holds similar for different distances. The reason for this is that the trajectory error at larger distances in many cases is due to the enlarging of certain part of the trajectory, which has fairly little effect on the recognition of a letter, and thus does not affect the recognition success rate.

9.2 Word Recognition Success Rates

Next, we study the success rate of recognizing the reconstructed trajectories for words. Overall, 92% of the word trajectories reconstructed by RF-IDraw were correctly recognized by the handwriting recognition app.

- Fig. 15 shows the word recognition success rate as a function of the number of characters in the word. It is expected that as the word gets longer, it is more difficult to recognize it correctly. Yet the word recognition success rate for RF-IDraw holds above 88% even for words consisting of 6 letters or more. Note that, the simple Android app that we use has only a basic handwriting recognition function suitable for a mobile phone. To enable a larger, full-fledged virtual touch screen, one could use advanced techniques in natural language processing to improve the success rate, especially for longer words, the dictionary of which is more confined and can be leveraged for better inference [40].
- None of the word trajectories reconstructed by the antenna array based scheme can be recognized correctly, i.e., 0% success rate. This is not surprising given its low character success rate above.
- Fig. 16 shows the trajectories for the same word "play" written by a user 5 m away from the reader antennas, as reconstructed by RF-IDraw and by the antenna array based system respectively.

While RF-IDraw is able to reconstruct the whole word, the antenna array based scheme's output fails to form any meaningful shape, because of its wide beams' low resolution and high sensitivity to noise.

In summary, using RF-IDraw's prototype, we enable a first-of-its-kind RF-based virtual touch screen in the air with 97.5% character recognition success rate and 92% word recognition success rate, far exceeding the capability of antenna array based technique with the same number of antennas.

9.3 Discussion

Finally, a few points are worth elaborating on, regarding the application of RF-IDraw:

- One could distinguish between two classes of in-the-air user interfaces. One class is based on a priori defining a few gestures like forward/backward motion [27], then using a machine learning approach to learn patterns and classify gestures into the learned categories. The alternative is an interface that is similar to having a pen for tablets. This interface does not require training or learning different user's motions. Yet it can trace a much richer set of gestures: one can create any command by drawing or writing, e.g., people can annotate slides in a meeting, draw icons/signs which would be interpreted by different computing devices, etc. While classification of a limited set of simple gestures may be sufficient for certain applications, we believe many emerging applications will benefit from an interface that can interpret a rich set of commands and does not rely on training, which is the approach we adopt in RF-IDraw.
- For applications that require selecting and manipulating items on a display, one can use RF-IDraw in a manner similar to operating a mouse to control a cursor on the screen. The user sees the cursor's position in real time and will naturally adjust her motion to reach the desired position based on the visual feedback.
- A limitation of our current implementation of RF-IDraw's virtual touch screen is that we manually segment the user's writing into words. We believe this can be addressed by using standard segmentation methods [26] in natural language processing, which would allow us to build a full-fledged virtual touch screen that can automatically process continuous streams of input.
- Finally, we note that the key idea of using grating lobes in RF-IDraw is transferable to other RF systems beyond RFID, such as WiFi and bluetooth. For example, one can potentially implement RF-IDraw on WiFi access points to trace the trajectories of nearby cellphones, which is one of our ongoing efforts. We acknowledge that the operating assumptions, constraints, and requirements of WiFi and bluetooth systems propose new challenges for applying RF-IDraw, which we plan to explore and address in future work.

10. CONCLUSION

This paper presents RF-IDraw, an accurate RFID-based trajectory tracing system that can transform any plane or surface into a virtual touch screen, allowing a user to interact with a desired computing device by writing her commands in the air. We believe RF-IDraw opens up a whole new class of applications in gaming and user interaction interface.

11. REFERENCES

[1] AN-900LH 900MHz antenna. rf-links.com/newsite/pdf/an900lh.pdf.
[2] Atacama Large Millimeter/submillimeter Array(ALMA). alma.mtk.nao.ac.jp/e/aboutalma/.
[3] Google Acquires Indoor/Outdoor Wireless Location Patents.
[4] MonkeyRunner API. developer.android.com/tools/help/MonkeyRunner.html.
[5] Omni-id exo 800. http://www.omni-id.com/products/.
[6] The corpus of contemporary american english. www.wordfrequency.info/.
[7] Ubi interactive. www.ubi-interactive.com.
[8] Vicon t-series. www.vicon.com/products/documents/Tseries.pdf.
[9] F. Adib, Z. Kabelac, D. Katabi, and R. C. Miller. 3d tracking via body radio reflections. NSDI, 2014.
[10] F. Adib and D. Katabi. See through walls with WiFi! SIGCOMM'13.
[11] Alien Technology Inc. ALN-9640 Squiggle Inlay. www.alientechnology.com.
[12] 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*.
[13] K. Chintalapudi, A. Padmanabha Iyer, and V. N. Padmanabhan. Indoor localization without the pain. MobiCom '10.
[14] Cisco. Cisco Announces Acquisition of ThinkSmart.
[15] EPCglobal Inc. EPCglobal Class 1 Generation 2.
[16] Forbes. Microsoft, Motorola, Nokia And RIM To Battle Google Over Indoor Location Market.
[17] Frost and Sullivan. Breakthrough Innovations in Indoor GPS, 2013.
[18] M. Hawes and W. Liu. Robust sparse antenna array design via compressive sensing.
[19] K. Joshi, S. Hong, and S. Katti. Pinpoint: Localizing interfering radios. NSDI, 2013.
[20] C. U. Keller. Interferometers. ATI 2010.
[21] R. Miesen et al. Holographic localization of passive uhf rfid transponders. In *IEEE RFID 2011*.
[22] P. Nikitin et al. Phase based spatial identification of uhf rfid tags. In *IEEE RFID 2010*.
[23] NOKIA. Accurate Mobile Indoor Positioning Industry Alliance, called In-Location, to promote deployment of location-based indoor services and solutions.
[24] S. J. Orfanidis. *Electromagnetic Waves and Antennas*. Macmillan Publishing Co., New York, 2010.
[25] P. Pal and P. Vaidyanathan. Coprime sampling and the music algorithm. 2011.
[26] R. Plamondon and S. N. Srihari. On-line and off-line handwriting recognition: A comprehensive survey. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2000.
[27] Q. Pu, S. Gupta, S. Gollakota, and S. Patel. Whole-home gesture recognition using wireless signals. MobiCom, 2013.
[28] A. Rai, K. K. Chintalapudi, V. N. Padmanabhan, and R. Sen. Zee: zero-effort crowdsourcing for indoor localization. Mobicom '12.
[29] M. T. Review. The Indoor Positioning System Era.
[30] SlashGear. Qualcomm Gimbal takes on Apple iBeacon for Micro-Location.
[31] Z. Tan, Y. Eldar, and A. Nehorai. Direction of arrival estimation using co-prime arrays: A super resolution viewpoint. 2013.
[32] TED. Free or cheap wii remote hacks.
[33] ThingMagic. Mercury6e rfid reader module. www.thingmagic.com/embedded-rfid-readers.
[34] A. R. Thompson, J. M. Moran, and G. W. Swenson. *Interferometry and Synthesis in Radio Astronomy*. Wiley-Interscience, 2001.
[35] D. Tse and P. Vishwanath. *Fundamentals of Wireless Communications*. Cambridge University Press, 2005.
[36] VisionObjects. Myscript stylus. www.visionobjects.com/en/myscript/.
[37] J. Wang, F. Adib, R. Knepper, D. Katabi, and D. Rus. Rf-compass: Robot object manipulation using rfids. MobiCom, 2013.
[38] J. Wang, H. Hassanieh, D. Katabi, and P. Indyk. Efficient and reliable low-power backscatter networks. SIGCOMM, 2012.
[39] J. Wang and D. Katabi. Dude, where's my card?: Rfid positioning that works with multipath and non-line of sight. SIGCOMM, 2013.
[40] W. Wang, A. Brakensiek, and G. Rigoll. Combination of multiple classifiers for handwritten word recognition. ICFHR'02.
[41] J. Xiong and K. Jamieson. Arraytrack: A fine-grained indoor location system. NSDI '13, 2013.