

Leveraging Electromagnetic Polarization in a Two-Antenna Whiteboard in the Air

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

ABSTRACT

Wireless sensing, tracking, and drawing technologies are enabling exciting new possibilities for human-machine interaction. They primarily rely on measurements of backscattered phase, amplitude, and Doppler signal distortions, and often require many measurements of these quantities—in time, or from multiple antennas. In this paper we present the design and implementation of PolarDraw, the first whiteboard in the air that sends differentially-polarized wireless signals to glean more precise tracking information from a tag. Leveraging information received from each polarization angle, our novel algorithms infer orientation and position of an RFID-tagged pen using just two antennas, when the user writes in the air or on a physical whiteboard. An experimental comparison in a cluttered indoor office environment compares two-antenna PolarDraw with recent state-of-the-art object tracking systems that use double the number of antennas, demonstrating comparable centimeter-level tracking accuracy and character recognition rates (88–94%), thus making a case for the use of polarization in many other tracking systems.

Keywords

RFID; Tracking; Electromagnetic Polarization

1. INTRODUCTION

We are rapidly moving toward a pervasively-sensed wireless world where most of our interactions with machines will be accomplished via gestures [3, 31] and writing in the air [38], and machines will be able to recognize our activities and pinpoint our location [1, 2, 16, 46]. On the industrial side, recent efforts include both short-range [12] and medium-range [11] indoor radar technologies, as well as mo-

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CoNEXT '16 December 12-15, 2016, Irvine, CA, USA

© 2016 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4292-6/16/12.

DOI: <http://dx.doi.org/10.1145/2999572.2999601>




Figure 1— Linearly-polarized electromagnetic waves propagate with a fixed electric field angle, exciting antennas or tags in proportion to alignment with this angle.




Figure 2— Recovered trajectory: *WoW, M, C, W, Z*.

re mature electronic whiteboard technologies based on ultrasound/infrared [25] and laser curtains [24]. While the use cases are compelling, the real-world utility of these sensing systems depends on several factors:

1. Infrastructure requirements— Systems that use angle-of-arrival information generally require a significant number of physically-separated antennas [3, 50], reducing deployability. Other systems require the simultaneous use of multiple widely-separated carrier frequencies [41], requiring simultaneous or added communication on these frequencies.

2. Pointing equipment— Some systems achieve millimeter accuracy, but require expensive (\$100) styluses in order to interact with ultrasound and infrared receivers [25].

3. Location and orientation estimation— In many cases, estimating the orientation of an object or person helps improve accuracy or activity recognition.

Electromagnetic wave *polarization* is a fundamental property of a wireless transmission, and refers to the orientation

Table 1— Infrastructure cost comparison. For a fair comparison of reader hardware, we include comparable products in the same product line as the ThinkMagic RFID reader hardware module that RF-IDraw used.

| Item | Unit cost | Quantity | Total cost |
|-------------------------|-----------|----------|------------|
| Reader (2-port) [39] | \$ 285 | 1 | \$ 285 |
| Antenna [20] | 79 | 2 | 158 |
| PolarDraw’s system cost | | | 443 |
| Reader (4-port) [40] | 398 | 1 | 398 |
| Antenna [53] | 135 | 4 | 540 |
| Tagoram’s system cost | | | 938 |
| Reader (4-port) [40] | 398 | 2 | 796 |
| Antenna [6] | 89 | 8 | 712 |
| RF-IDraw’s system cost | | | 1508 |

of the electrical field in the plane transverse to the wave’s propagation, as shown in Figure 1. In the context of mobile devices, polarization has received the most attention in recent years as increasing smartphone usage in an orientation facing the user (roughly level with the ground) has resulted in a loss of signal strength from a vertically-polarized transmissions, both indoors and outdoors [33]. Outdoors, the horizontally- and vertically-polarized paths from base station to mobile (of length *ca.* 2 mi.) are very different [21] due to differing reflectors. Indoors, for smaller Wi-Fi cell sizes, the same is true for longer non-line-of-sight paths, but shorter paths tend to share similar reflectors (and a more dominant line-of-sight path) when viewed from orthogonal polarizations [34].

It is well known among RFID practitioners that tag orientation affects RFID read rate, and prior academic work has leveraged these electromagnetic polarization differences to infer RFID tag orientation [17]. Our key observation in this work is that the similarity between horizontally- and vertically-polarized paths at modest ranges presents a unique opportunity for wireless motion tracking systems to leverage polarization for better location tracking accuracy. We use a simple RFID reader as a test case, but replace the reader’s standard circularly polarized antennas with linearly polarized antennas as shown in Figure 1, resulting in the linearly polarized transmission shown.

In this paper, we present the design and implementation of PolarDraw, the first motion tracking system that can accurately reproduce handwritten letters in the air or on a whiteboard with just *two* nearby antennas. Compared with the prior state-of-the-art system Tagoram, PolarDraw reduces the infrastructure cost by half (see Table 1), hence setting a new standard for minimal supporting infrastructure¹. PolarDraw estimates both the position and orientation of an RFID-tagged whiteboard pen, adding just pennies to the cost of each item to be used as a stylus. Thus PolarDraw meets each of

¹PolarDraw’s evaluation actually uses a much cheaper antenna (\$ 14 unit price) but we use a conservative cost estimation here for a generally-fair cost comparison.

the three preceding objectives for a practical and highly-deployable motion tracking system.

In brief, PolarDraw works as follows. Our system uses RFID antennas to measure the phase and amplitude of an RFID tag at a rate of *ca.* 100 Hz. Variations in the angle between tag and reader antenna result in fluctuations of power received from the tag, but as our feasibility study next (§2) shows, symmetry properties result in equal power changes when the tag rotates left or right. PolarDraw uses a second, linearly-polarized antenna to view the rotation from a 30° shift in perspective, thus overcoming this *rotational ambiguity* problem (§3.3). Rotation is just part of what determines the pen’s true trajectory, however. PolarDraw also estimates the displacement of the pen on the whiteboard using phase difference readings across both time and space (§3.4). Since they do not attempt to discriminate polarization, current trajectory tracing systems have not yet explored how best to fuse the above phase readings (which estimate pen displacement) with power readings (which estimate pen rotation), and doing so carefully is key to getting highly accurate results. We describe our novel algorithm, which incorporates a simple Viterbi-based probabilistic search (§3.5), to complete the design of PolarDraw.

Contributions. Although using polarization to infer tag orientation has been studied before, PolarDraw is the first system we are aware of that uses a power-orientation model to track a pen’s movement trajectory. The techniques introduced here introduce a new source of position information, and thus have the potential to be applied to many other indoor radar and indoor localization systems to enhance their accuracy. PolarDraw is also the first radio-based system we are aware of that simultaneously models changes in the pen’s displacement and orientation. As our experimental evaluation shows, PolarDraw benefits from the pen orientation estimation.

Experimental results. Our experiments measure:

1. Word recognition accuracy.
2. Letter recognition accuracy.
3. The general precision of tracking arbitrary writing.

Our experimental evaluation (§5) tests PolarDraw’s performance against the two leading motion tracking systems in the research literature mentioned above, RF-IDraw and Tagoram, with all three systems running in the same experimental environment. Our results in this cluttered, real-world office environment show two-antenna PolarDraw achieves median 10 cm accuracy, versus 8 cm for four-antenna RF-IDraw and Tagoram. Further microbenchmark experiments justify parameter choices and test the experimental limits of our system as tag to reader distance increases and writing size decreases.

Roadmap. The rest of this paper is organized as follows: §2 presents initial microbenchmark-style measurements in a cluttered office environment to establish the basic experimental possibility of measuring orientation through polarized transmissions. §3 presents the design of PolarDraw, after which we present our implementation (§4) and experimental results (§5). The remainder of the paper consists of related work (§6), discussion (§7) and conclusion (§8).




Figure 3— An initial feasibility study experimental setup and experimental results.

2. MEASURING POLARIZATION

To determine whether we can measure polarization in a real indoor environment where multipath is prevalent, we first describe empirical feasibility study, using the hardware setup shown in Figure 3(a). We situate an Impinj RFID reader connected to a linearly polarized antenna 2.5 m above an RFID tag placed on a turntable below, as shown in the picture.² In the first experiment, the tag rotates at a constant angular velocity on the turntable. In the second experiment, we manually translate the tag back and forth, keeping its orientation fixed. Figure 3(b) shows the received signal strength (RSS) and phase measured at the RFID reader during the tag’s rotation. As expected, RSS changes periodically during the tag’s rotation. It peaks at -24 dBm when the tag and the reader antennas are aligned, and drops gradually as the polarization mismatch angle between the two increases. Finally, the tag fails to acquire any power and there is no RSS reading when the orientation of the tag is perpendicular to the polarization angle of the reader antenna. On the other hand, since tag-to-antenna distance does not change during the tag’s rotation, the phase reading stays roughly constant. We note that the phase reading jumps when the polarization mismatch angle is around 90° and 270° : this is because the tag fails to acquire energy from the line-of-sight path due to the polarization mismatch. Nonetheless, it acquires energy along non-line-of-sight signal propagation paths, where the signal bounces off nearby objects, changing the measured phase angle. PolarDraw’s data pre-processing (§3.1) is designed to cope with and overcome these *spurious* readings from differently-polarized reflection paths.

Figure 3(c) shows RSS and phase during the tag’s back-and-forth movement over a distance of 8 cm. We observe a roughly constant RSS value, as expected since RSS is insensitive to small changes in distance over the tag-to-antenna link. In contrast, we can see the phase reading increases when the tag moves in one direction, remains stable when the tag stands still, and decreases when the tag moves back.

We conclude from the above empirical results:

1. At moderate distances in our line-of-sight indoor environment, RSS is sensitive to the polarization mismatch, yet it is insensitive to small changes in the length of the tag-to-antenna link.

²Our hardware setup is more completely described in §4.




Figure 4— PolarDraw’s high-level design: Two linearly-polarized antennas are mounted above a whiteboard or free space; users draw with RFID-tagged whiteboard pens.

2. With the exception of a “spurious” corner case when the tag antenna and reader antenna are orthogonal to each other, phase measurements are insensitive to polarization mismatch, yet are extremely sensitive to small changes in the length of the tag-to-antenna link.

Combining these two findings, we can conclude that in our experimental setup, separate measurements of RSS and phase can respectively estimate rotational and translational movement of the pen, with a reasonable degree of reliability. In the next section, we describe a design that extracts a high degree of reliability from these two signals.

3. DESIGN

PolarDraw comprises a UHF RFID reader connected to two linearly-polarized antennas mounted on a whiteboard as shown in Figure 4, or alternatively, simply tracking a “virtual whiteboard” in the air. Figure 5 shows PolarDraw’s data processing workflow: first, PolarDraw controls the RFID hardware to interrogate the tag collecting the phase and amplitude readings of each backscattered signal. These data are then delivered to the pre-processing module for data smoothing and spurious data rejection (§3.1). After that, the filtered data are fed into the tracking module, where PolarDraw uses a writing model (§3.2) to estimate pen moving direction (§3.3) and distance (§3.4). These pen moving direction and distance estimates finally feed into a Hidden Markov Model (HMM) for pen trajectory estimation (§3.5).

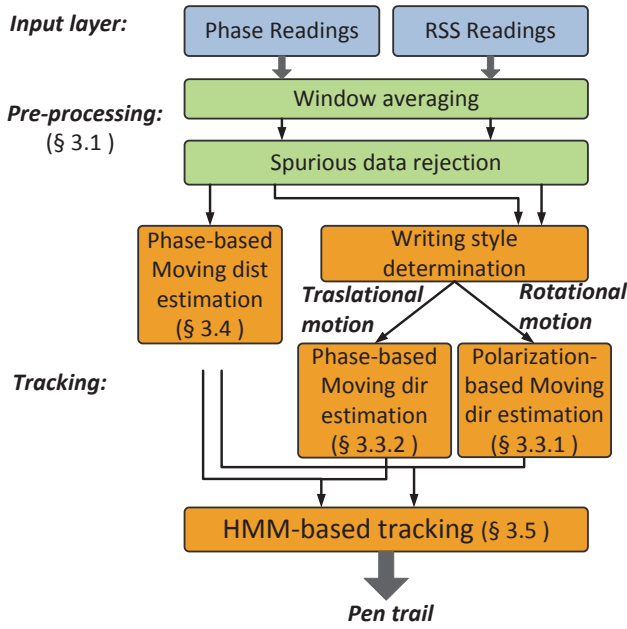


Figure 5— PolarDraw’s high-level design: Separate modules handle data input, pre-processing and pen tracking.

3.1 RFID data pre-processing

The pre-processing module contains two steps: data smoothing and spurious data rejection. In the first step, data smoothing aims to improve the stability of amplitude and phase samples by averaging away the phase and amplitude noise affecting the raw data. PolarDraw first divides time into windows of length 50 ms.³ Within each window, PolarDraw averages together the amplitude and phase readings belonging to that window. These averaged data form a new amplitude and phase series, indexed by the timestamp.

In the second step, PolarDraw rejects the spurious phase data described in §2 by comparing phase readings of adjacent windows. If the absolute value of their difference is larger than a threshold, PolarDraw rejects the readings at current window. Following the experimental setup in §2, we rotate the tag on the turntable at different speeds and test for the presence of spurious phase readings. Based on the result, we empirically set this angle threshold to 0.2 radians. The filtered data are then fed into the tracking module.

3.2 Modeling pen movement

A *pen trajectory* is composed of a set of discrete trajectory *fragments*, each of which can be characterized by a movement direction and distance on the (virtual or real) whiteboard. When a user is writing on a whiteboard, wrist movements tend to cause azimuthal rotations clockwise when the pen moves to the right, and counterclockwise when the pen moves to the left. The pen is of course free to rotate in three-space, and so we define the angular measurements of the pen’s elevation from the X–Z plane and azimuthal rotation

³We chose this averaging window size based on an empirical sensitivity analysis.

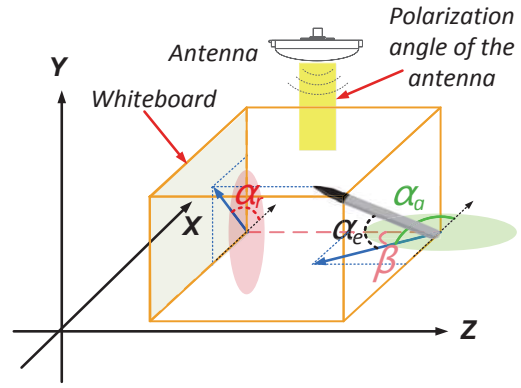


Figure 6— Modeling pen movement during writing.

measured from the X-axis and projected onto the X–Z plane summarized in Table 2 and illustrated in Figure 6.

| | |
|--------------|---|
| α_e : | Pen elevation angle, from the X–Z plane, |
| α_r : | Pen rotation angle, as projected on the X–Y plane, |
| α_a : | Pen azimuthal angle, as projected on the X–Z plane. |

Table 2— Pen angle as illustrated in Figure 6.

3.3 Estimating pen movement direction

PolarDraw estimates pen direction using both rotational and translational pen movements. PolarDraw’s rotational direction estimation estimates left-right directions only, and its translational direction estimation estimates movement in all four board directions. At each 50 ms timestep, PolarDraw first detects the dominant type of movement (translational or rotational) by observing the RSS trend: if the change in RSS reaches above some threshold δ ,⁴ then PolarDraw infers primarily rotational movement and tracks α_a (§3.3.1) to determine its direction. Otherwise, PolarDraw infers significant translational movement and instead inspects changes in phase to determine the pen’s direction of movement (§3.3.2).

3.3.1 Rotational movement direction estimation

The main idea behind rotational movement direction is as follows. We continuously estimate α_r , the pen rotation angle projected onto the whiteboard, and estimate a pen moving direction perpendicular to α_r , as illustrated in Figure 7.

The problem then reduces to the estimation of α_r . PolarDraw translates α_a to the pen rotation angle α_r using the following equation relating the angles shown in Figure 6:

$$\alpha_r = \pi - \arctan\left(-\frac{\sin \alpha_e}{\cos \alpha_e \cdot \cos \alpha_a}\right). \quad (1)$$

Since the antennas are directly above the pen, we can estimate α_a based on polarization mismatch (we describe this process next). There remains another variable free in Equation 1 however, α_e , the elevation angle. For simplicity of

⁴We test various thresholds, and empirically set $\delta = 2$ dBm, which optimizes the overall performance.

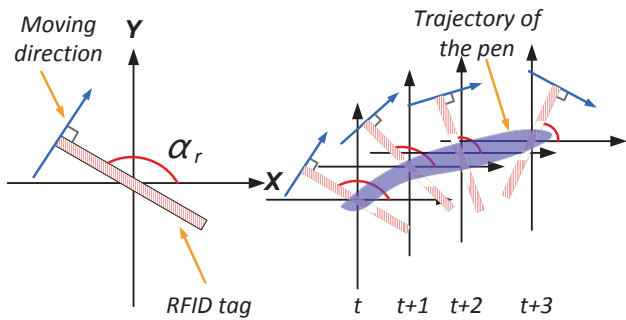


Figure 7— Tracking pen rotation and translation.




Figure 8— Polarization angle mismatch. The thick blue line indicates the polarization angle of the reader antenna. (a): rotation direction ambiguity. (b): azimuthal angle ambiguity. (c): breaking the rotation ambiguity by spatially separated antennas.

design we take a pragmatic approach: since pen elevation angle changes only slightly when rotating left to right while writing, we set α_e to be a constant determined by experiments described in §5.4.1 (there we also show that accuracy is insensitive to choice of α_e). Equation 1 can then be solved for α_r and hence pen movement direction.

The problem now reduces to the estimation of α_a . In designing PolarDraw's α_a estimation algorithm, we face the following two challenges. Firstly, as Figure 8(a) shows, both a clockwise and a counterclockwise rotation could result in an identical polarization mismatch angle β , hence the same end RSS. Consequently, we are unable to differentiate the rotation direction by simply observing the end RSS value alone. We term this challenge the *rotation direction ambiguity*. Secondly, as Figure 8(b) shows, within the range $[0, 180^\circ]$, there are always two azimuthal angles that result in the same RSS value. Therefore, we are also unable to estimate the azimuthal angle based on RSS alone. We term this challenge the *azimuthal angle ambiguity*.

Breaking ambiguities. An intuitive way is to refer to the tag's location. However, the tag's location is undetermined since the two-antenna based hyperbolic positioning scheme will yield multiple symmetric hyperbolas, and the tag's loca-




Figure 9— Writing motion RSS trends ($\gamma = 30^\circ$).


Table 3— RSS trend. (“-” means don't care.)

| Area | Direction | RSS trend | RSS changing rate |
|----------|---------------|---------------------------------------|-----------------------------------|
| Sector 1 | \Rightarrow | $s_i^1 \uparrow$ $s_i^2 \uparrow$ | $ \Delta s_i^1 < \Delta s_i^2 $ |
| | \Leftarrow | $s_i^1 \downarrow$ $s_i^2 \downarrow$ | $ \Delta s_i^1 < \Delta s_i^2 $ |
| Sector 2 | \Rightarrow | $s_i^1 \downarrow$ $s_i^2 \uparrow$ | - |
| | \Leftarrow | $s_i^1 \uparrow$ $s_i^2 \downarrow$ | - |
| Sector 3 | \Rightarrow | $s_i^1 \downarrow$ $s_i^2 \downarrow$ | $ \Delta s_i^1 > \Delta s_i^2 $ |
| | \Leftarrow | $s_i^1 \uparrow$ $s_i^2 \uparrow$ | $ \Delta s_i^1 > \Delta s_i^2 $ |

tion could be anywhere on any of these hyperbolas. In PolarDraw, we address the above challenges by jointly analyzing RSS trends from the two antennas. As Figure 8(c) shows, we orient the antennas so that their angles with the Z-axis are equal, denoting that measure γ . Let β_i^j denote the polarization mismatch angle between the tag and antenna j at time i . Then the antenna directions together with their perpendiculars naturally separate the rotation plane into three sectors 1, 2, and 3, as shown in Figure 8(c). During normal writing, pen azimuthal rotations are small relative to γ , and so we assume the azimuthal angle α_a is well within the union of these three sectors. Therefore, if we can determine in which sector the current α_a is, then we can break the azimuthal angle ambiguity.

Figure 9 shows RSS trends reported by the two antennas when a user writes on the whiteboard. Let s_i^j be the RSS reading reported by antenna j at time i , and $\Delta s_i^j = s_{i+1}^j - s_i^j$. As indicated in this figure, when the pen rotates in different direction within different sectors, the polarization mismatch angle β_i^1 and β_i^2 changes in a different way, which leads to a different direction of change as well as a different rate of change of RSS—Table 3 summarizes the result. From the table we find that by jointly analyzing RSS trends and changing rates, we can successfully estimate the moving direction of the pen as well as the range of the current azimuthal angle α_a that the pen points to, thereby solving both the rotation direction ambiguity and the azimuthal angle ambiguity.

Continuous azimuthal angle tracking. Let α_a^i be the azimuthal angle of the pen at time i . When the user begins writing, PolarDraw estimates in which sector the pen points and in which direction the pen points. Based on these, Polar-



(a) (b)— Pre-correction. (c)— Post-correction.

Figure 10— Ground-truth (a), and recovered pen trajectory before and after azimuthal angle correction.

Table 4— Phase changing trends during pen movement.

| Pen: | Up | Down | Left | Right |
|--------|--|--|--|--|
| Phase: | $\theta_i^1 \downarrow, \theta_i^2 \downarrow$ | $\theta_i^1 \uparrow, \theta_i^2 \uparrow$ | $\theta_i^1 \downarrow, \theta_i^2 \uparrow$ | $\theta_i^1 \uparrow, \theta_i^2 \downarrow$ |

Draw assigns an initial azimuthal angle as follows:

$$\alpha_a^0 = \begin{cases} \pi - \gamma, & \text{if clockwise and in Sector 1,} \\ \frac{\pi}{2} + \gamma, & \text{if clockwise and in Sector 2,} \\ \frac{\pi}{2} - \gamma, & \text{if clockwise and in Sector 3,} \\ \frac{\pi}{2} + \gamma, & \text{if counterclockwise and in Sector 1,} \\ \frac{\pi}{2} - \gamma, & \text{if counterclockwise and in Sector 2,} \\ \gamma, & \text{if counterclockwise and in Sector 3.} \end{cases} \quad (2)$$

PolarDraw then tracks α_a^i at each time i as follows:

$$\alpha_a^i = \begin{cases} \alpha_a^{i-1} - \Delta\beta, & \text{if clockwise rotation,} \\ \alpha_a^{i-1} + \Delta\beta, & \text{if counterclockwise rotation.} \end{cases} \quad (3)$$

where $\Delta\beta$ is a variable, indicating the average changing rate of the azimuthal angle over each 50 ms time window when human writes on the whiteboard. The assignment of $\Delta\beta$ is as follows:

$$\Delta\beta = \begin{cases} 6^\circ, & \text{if } |\Delta s_i^1| > \delta \text{ and } |\Delta s_i^2| > \delta, \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

We perform a sensitivity analysis of both δ and $\Delta\beta$, and set $\delta = 1.5$ dBm, $\Delta\beta = 6^\circ$, which jointly maximize overall system performance.

Initial azimuthal angle correction. In the tracking process, the initial azimuthal angle we assigned will likely deviate from its true value by some amount $\tilde{\alpha}_a$. Consequently, later azimuthal angle estimates suffer from this error $\tilde{\alpha}_a$ as well, resulting in an inaccurate tracking result.

PolarDraw corrects $\tilde{\alpha}_a$ based on detecting when the azimuthal angle of the pen crosses over the boundary of two sectors. Suppose at time i PolarDraw detects that the pen crosses over the boundary of two neighboring sectors based on the principle shown in Table 3. The azimuthal angle of the pen at time i should be approximately equal to the azimuthal boundary angle of these two sectors, denoted as $\hat{\alpha}_a^i$. Hence the difference between $\hat{\alpha}_a^i$ and α_a^i (the estimated azimuthal angle) indicates the initial azimuthal angle error $\tilde{\alpha}_a$. PolarDraw corrects the estimated azimuthal angles by subtracting $\tilde{\alpha}_a$ from all α_a^i , thus achieving a more accurate result. Figure 10 shows a snapshot of our system's output before and after azimuthal angle error correction.

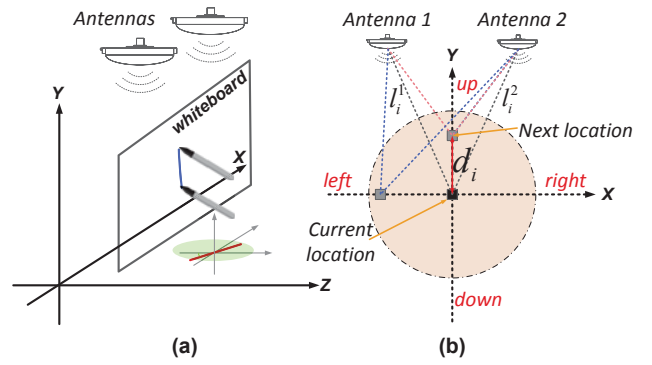


Figure 11— Writing without rotation. (a): the user writes on the whiteboard with a constant rotation angle; (b): geometry relationship of moving directions;

3.3.2 Translational movement direction estimation

Writing movement varies from moment to moment, and so sometimes the pen may move with negligible rotation, translation dominating its path, as shown in Figure 11(a). Consequently, the azimuthal angle α_a fails to reflect the true moving direction of the pen. When the pen moves in this primarily translational manner, PolarDraw quickly (within a 50 ms timestep) switches to exploit phase changes to estimate the pen's direction of movement. Specifically, we denote the distance between reader antenna j and the tag at time i as l_i^j . Let θ_i^j be the unwrapped phase reading reported by the antenna j at time i . As Figure 11(b) shows, when the pen moves up, both l_i^1 and l_i^2 decrease, resulting in $\theta_i^1 \downarrow$ and $\theta_i^2 \downarrow$. Conversely, when the pen moves down, both l_i^1 and l_i^2 increase, leading to $\theta_i^1 \uparrow$ and $\theta_i^2 \uparrow$. Table 4 shows measured phase trends when the pen moves in different directions, allowing us to decode four rough movement directions over the small timestep. In the next section, we describe how we determine the magnitude of this estimated movement.

3.4 Pen movement distance estimation

PolarDraw again exploits phase trends to estimate the moving distance of the pen. Our approach is based on an assumption that the moving distance during consecutive tag readings is within a half wavelength ($\lambda/2 \approx 16$ cm), where λ is the wavelength of the backscattered signal. This assumption holds in practice since our processing timestep is 50 ms, yielding a maximum detectable speed of 3.2 m/s—significantly exceeding normal writing speed.

Moving distance bounds. Let d_i be the moving distance of the pen at timestep i , and l_i^j be the distance between the pen and reader antenna j at the same timestep. Since phase cycles completely over a distance of λ , the change in distance between the pen and a reader antenna $\Delta l_i^j = l_{i+1}^j - l_i^j$ can be computed as follows:

$$\Delta l_i^j = (\theta_{i+1}^j - \theta_i^j) \cdot \lambda / (4\pi). \quad (5)$$

By the triangle inequality, $d_i \geq \max\{|\Delta l_i^1|, |\Delta l_i^2|\}$. Next, we denote the maximum moving speed of the pen as v_{max} (we

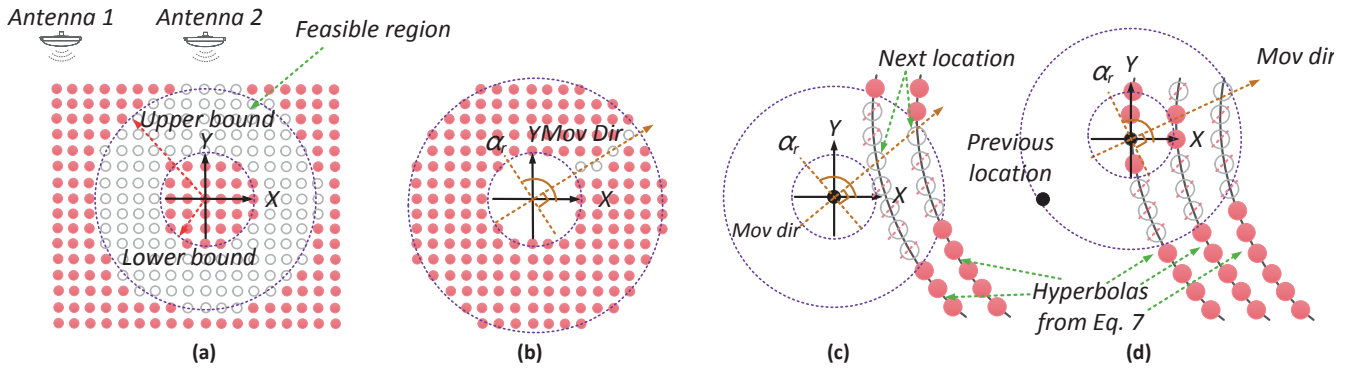


Figure 12— Moving distance inference; (a): Under the lower and upper distance bound. (b). Under the moving direction constraint; (c): Under the hyperbola constraint. (d): Tracking the second trajectory segment of the pen.

set v_{max} to 0.2 m/s in the experiment, which leads to empirically acceptable tracking accuracy.). Then an upper bound of the displacement is $v_{max}\Delta t$, where Δt is the 50 ms timestep.

Based on the foregoing analysis, we know the pen moving distance d_i should be within an annulus: $\max\{|\Delta l_i^1|, |\Delta l_i^2|\} \leq d_i \leq v_{max}\Delta t$ (as shown in Figure 12(a)), which we term the *feasible region*. Within the feasible region, the line segment between the current location and each block forms the displacement *candidate set*. Since the pen's next location should be aligned with the pen's current moving direction, we eliminate a large portion of infeasible locations in the feasible region, as shown in Figure 12(b). However, a large moving distance uncertainty still remains, as many blocks are aligned with the pen moving direction.

Antenna difference measurements. We also exploit the inter-antenna phase difference $\theta_i^2 - \theta_i^1$ to further reduce the uncertainty in our pen movement distance estimate. The relationship between phase readings and the tag-to-antenna distance can be formulated as follows:

$$\begin{cases} \theta_i^1 + 2k_1\pi = 4\pi l_i^1/\lambda \\ \theta_i^2 + 2k_2\pi = 4\pi l_i^2/\lambda \end{cases} \quad (6)$$

Subtracting the above equations and dropping subscripts:

$$\Delta l_i^{2,1} = \frac{\lambda}{4\pi} (\Delta\theta_i^{2,1} + 2k\pi), \quad (7)$$

where k is an unknown integer, $\Delta\theta_i^{2,1} = \theta_i^2 - \theta_i^1$ and $\Delta l_i^{2,1} = l_i^2 - l_i^1$. Let (x_1, y_1) and (x_2, y_2) be the location of two antennas, $\Delta l_i^{2,1}$ be the distance difference of these two antenna-to-tag links, we can construct a hyperbola with the two foci at the location (x_1, y_1) and (x_2, y_2) as our location estimation. Due to the phase ambiguity, there are multiple candidate hyperbolas. PolarDraw makes its movement distance estimate by intersecting the candidate hyperbolas with the location candidates already acquired, as shown in Figure 12(c) for the case of rotational moving direction estimation (we take an analogous approach for the case of translational moving direction estimation).

3.5 Tracking pen trajectory

Our design so far focuses on how to track the moving direction and the moving distance of the pen, separately. Now

we put them together and show how to leverage moving direction and the moving distance to infer the pen's trajectory fragment at each timestamp.

We approach the pen trajectory tracking problem as a discrete-time state estimation problem, where the state of the system at time t , X_t , is the location of the pen at that time. The measurements at time t , Y_t , are the phase and RSS readings reported by the two antennas. PolarDraw employs an HMM to find the most likely sequence of pen trajectory segments. The HMM consists of a set of interconnected states, each of which emits an observable output. The HMM is characterized by two probabilities: the *transition probability* between states and the *emission probability* over the position estimate.

Modeling the state transition probability. We divide the whiteboard into equal-sized blocks, denoted $B_{i,j}$. Then the state space of HMM consists of all blocks on the whiteboard. We consider the transitions between consecutive states to have equal probabilities.⁵ This state transition probability $Pr(X_{t+1} = (x_1, y_1) | X_t = (x_0, y_0))$ is defined as:

$$\begin{cases} \frac{1}{M}, & \text{if } \max\{|\Delta l_t^1|, |\Delta l_t^2|\} \leq \sqrt{(x_1 - x_0)^2 + (y_1 - y_0)^2} \\ & \leq v_{max}\Delta t, \\ 0, & \text{otherwise,} \end{cases} \quad (8)$$

where M is the the number of blocks within the feasible region (*i.e.*, gray circles in Figure 12(a)).

Modeling emission probability. At timestep t , the observation Y_t is represented as a tuple of phase and RSS readings: $[\theta_t^1, \theta_t^2, s_t^1, s_t^2]$. The emission probability for a given state represents the likelihood of seeing the pen located at the block conditioned on Y_t . We model emission probability by jointly considering both the moving direction constraint (Figure 12(b)) and the hyperbola constraint (Figure 12(c)). Let (x_0, y_0) and α_t' be the location and the rotation angle of the pen at time t , respectively. We construct the line passing through (x_0, y_0) with slope equal to the estimated pen moving direction to describe the approximate trajectory of the pen:

$$y = -\cot(\alpha_t') \cdot (x - x_0) + y_0. \quad (9)$$

⁵We leave more sophisticated motion modeling, such as the Kalman and Particle filters, for future work.

The pen’s next location is likely on this trajectory. On the other hand, the pen’s next location should also be on the hyperbolas we estimated based on the inter-phase difference of two antennas $\Delta\theta_t^{1,2}$. Equation 11 in the Appendix on p. , combines the moving direction and the hyperbola constraint to define the emission probability $Pr(X_{t+1} = (x_1, y_1) | Y_t = \theta_t^1, \theta_t^2, s_t^1, s_t^2)$.

Initial location estimation. PolarDraw finds the initial location of the tag using hyperbolic positioning. Since two antennas will determine multiple hyperbolas, PolarDraw randomly chooses a point on one of these hyperbola to bootstraps the relative location.

Trajectory Rotation. PolarDraw leverages Viterbi decoding to find the most likely pen trajectory. The computational complexity of Viterbi decoding is linear in the number of states and the number of observed variables, hence can be computed in real-time even with an embedded mini PC. As mentioned in §3.3, the initial azimuthal angle may have an error $\tilde{\alpha}_a$. Hence after the trajectory tracking, PolarDraw corrects the azimuthal angle error to acquire a more accurate pen trajectory. Let $\mathbb{P} = \{P_{x_1, y_1}, P_{x_2, y_2}, \dots, P_{x_T, y_T}\}$ be pen’s trajectory recovered by the Viterbi algorithm, where P_{x_t, y_t} is the location of the pen on time t . PolarDraw eliminates the impact of $\tilde{\alpha}_a$ on pen’s trajectory as follows:

$$\hat{\mathbb{P}} = \mathbb{P} \cdot \begin{bmatrix} \cos \tilde{\alpha}_r & -\sin \tilde{\alpha}_r \\ \sin \tilde{\alpha}_r & \cos \tilde{\alpha}_r \end{bmatrix} \quad (10)$$

where $\hat{\mathbb{P}}$ is the new trajectory of the pen; $\tilde{\alpha}_r$ is the rotational angle error induced by the azimuthal angle error.

4. IMPLEMENTATION

PolarDraw’s front-end hardware consists of an Impinj Speedway R420 RFID reader [15], two Laird tilt RF linearly polarized antennas [19], and an Avery Dennison AD-227m5 UHF passive RFID tag [7]. The passive tag is attached on a standard whiteboard pen. PolarDraw’s software components control the RFID reader to interrogate the RFID tag and process the tag readings for pen trajectory tracking. The tag interrogation module, implemented in Java, collects tag readings through the Low Level Reader Protocol (LLRP) [23] which the C#-based pen trajectory tracking module then processes.

Modulation scheme selection. A typical EPC GEN 2 reader supports different modulation schemes, with varying levels of robustness to background noise. PolarDraw round-robins all available modulation schemes, selecting the first with the standard deviation of phase variances at most 0.1 rad² for tag interrogation. We have conducted a sensitivity analysis to test other thresholds, and find this choice of phase variance threshold optimizes performance.

5. EVALUATION

In this section, we first present our experimental methodology (§5.1), after which we conduct end-to-end experiments in a laboratory environment to evaluate PolarDraw’s performance (§5.2). We then compare PolarDraw with two state-of-the-art motion tracking systems, Tagoram and RF-IDraw.

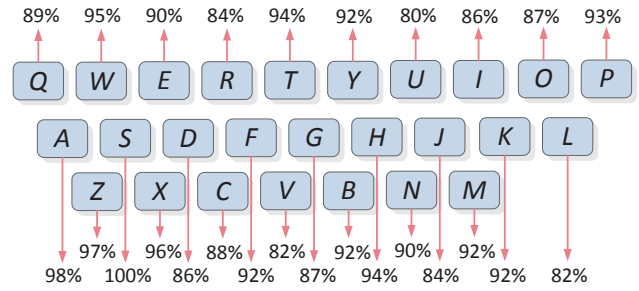


Figure 13— PolarDraw’s letter recognition accuracy.

Finally, we present micro-benchmark experiments that provide additional insight into which factors most impact PolarDraw’s performance (§5.4).

5.1 Methodology

To determine ground truth, we photograph the user’s writing on the whiteboard and leverage edge detection algorithm to extract writing trajectories. We use LipiTk [13] for hand writing recognition.

Performance metrics. We use the following metrics:

1. *Recognition Accuracy*: the fraction of successful character recognitions over the total number of characters.
2. *Similarity*: the *Procrustes distance* between the recovered trajectory and the ground truth trajectory. This metric determines a linear transformation (translation, rotation, and scaling) of the points in one trajectory that results in the closest match with to the points in another. The goodness-of-fit criterion is the sum of squared errors between the two trajectories.
3. *Confusion matrix*: each row shows the actual character and each column shows the estimated character.

Compared schemes. We compare PolarDraw with two other 2D tracking algorithms: RF-IDraw [45] and Tagoram [52]. RF-IDraw uses eight spatially separated antennas for RFID tag localization and tracking. Most COTS RFID reader support four antennas apiece, so we compare a four-antenna version of RF-IDraw with PolarDraw, for equal hardware resources. Tagoram adopts four antennas to locate and track an RFID tag based on the phase readings: our implementation uses two or four antennas, the former comparing Tagoram against PolarDraw with equal hardware resources.

5.2 End-to-end performance

We first present end-to-end results of studies in an office environment. The algorithm parameters are set according to the results in §5.4.

5.2.1 Character recognition accuracy

We first examine character recognition accuracy over the English alphabet. In these experiments, we invite a volunteer to write all 26 characters 100 times. The result is shown in Figure 13. PolarDraw achieves 93.6% recognition accuracy on average. Specifically, 15 out of 26 characters are correctly recognized with a probability higher than 90%. 21 out of 26 characters are correctly recognized with a probability higher than 85%. The remaining five characters have

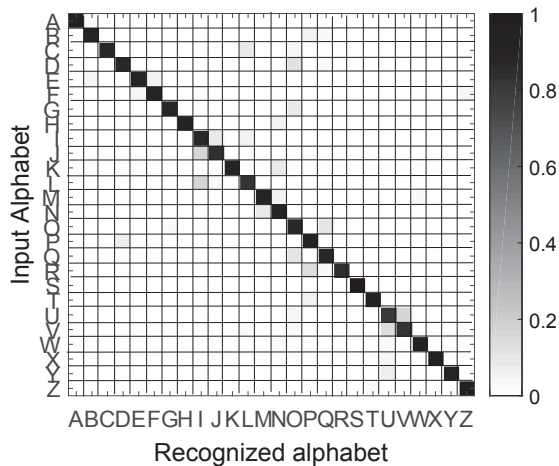


Figure 14— PolarDraw's letter confusion matrix.




Figure 15— Writing in air v. on the whiteboard.

a much lower recognition accuracy. However, their recognition accuracies are still above 80%. We believe by applying natural language processing techniques, we can further increase recognition accuracy.

5.2.2 Character recognition confusion matrix

To quantify whether character recognition performance is better for some characters over others, we present the character confusion matrix in Figure 14. The darker areas represent a higher rate of classification of the particular character. We observe from this figure that the rate of misclassifications varies for different ground truth characters. A large portion of errors is usually associated with those letters that share a very similar writing style. For example, the characters L and V are prone to be incorrectly recognized as I and U. We also note that recognition accuracy also varies with the complexity of the character. For example, those characters that can be written in a single stroke usually achieve a higher recognition accuracy.

5.2.3 Performance of in-air writing

We next remove the whiteboard and let the user write in the air. We conduct four groups of experiments here. In each group, we randomly choose 10 letters and let the user write each letter 10 times in air. For comparison, this user is also required to repeat these experiments but write on the whiteboard. Figure 15 shows the resulting recognition accuracy. As the result shows, PolarDraw performs consistently when the user is writing on the whiteboard, achieving a recognition accuracy around 91% on average. In contrast, PolarDraw's

Table 5— Recognition accuracy v. tag-reader distance.

| Distance (cm) | 20 | 40 | 60 | 80 | 100 | 120 | 140 |
|---------------|----|----|----|----|-----|-----|-----|
| Accuracy (%) | 77 | 83 | 87 | 90 | 91 | 90 | 88 |




Figure 16— Impact of both dynamic and static multipath on PolarDraw's performance.

performance declines about 8% when the user writes in the air. The reason is that without the whiteboard, writing is not confined to a two-dimensional plane, which leads to errors in the pen moving distance inference. However, PolarDraw still achieves over 80% recognition accuracy in this case.

5.2.4 Effect of tag-to-reader distance

Here we examine the impact of the tag-to-reader distance on PolarDraw's trajectory tracking performance. In these experiments, we vary this distance from 20 cm to 140 cm, with a step of 20 cm. Table 5 shows the recognition accuracy at each distance. Somewhat surprisingly, accuracy is relatively low (77%) when the tag is close to the reader antenna (e.g., 20 cm spacing)—this is because in this regime both polarization angle mismatch and the tag movement contribute to the RSS change. But as we increase the tag-reader distance, RSS becomes less sensitive to pen translational movements and PolarDraw accuracy increases. Recognition accuracy only slightly drops as we further expand the tag-to-reader distance from 100 cm to 120 cm and further to 140 cm. One possible explanation may be that the backscattered signal bounces off nearby objects and changes the polarization angle, which leads to an unusual RSS trend during human writing. We plan to investigate the use of more directional antennas to increase range in future work.

5.2.5 Effect of nearby interfering human motion

To examine the effect of multipath reflections on PolarDraw's pen tracking accuracy, we invite a volunteer to write on the whiteboard, while another person stands nearby or walks around the whiteboard to generate static or dynamic multipath interference on the backscattered signal. Figure 16 shows the recognition accuracy of PolarDraw under these different multipath conditions, as a function of the distance of the interfering user ("bystander") from the whiteboard. As the results show, PolarDraw's performance is insensitive to both dynamic and static multipath when the bystander is 90 cm away from the whiteboard. As the bystander stands or walks even closer nearby (60 cm and 30 cm), PolarDraw's performance only slightly decreases—PolarDraw still achieves

Table 6— Recognition accuracy of PolarDraw with and without polarization information.

| Algorithm | PolarDraw | w/o polarization |
|--------------|-----------|------------------|
| Accuracy (%) | 91 | 23 |




Figure 17— Tagoram/RF-IDraw experimental setup.

ves 83% and 87% recognition accuracy in these dynamic and static multipath regimes, respectively, demonstrating strong robustness.

5.2.6 Gain of using polarization

To drill down into the benefit of polarization on tracking accuracy, we implement a version of PolarDraw that tracks pen trajectory without polarization angle estimation. Table 6 presents the letter recognition accuracy of PolarDraw with and without polarization angle estimation. The data show that PolarDraw benefits substantially from polarization angle estimation—by a factor of around $4\times$ in terms of letter recognition accuracy.

5.3 Comparison v. RF-IDraw and Tagoram

We now compare PolarDraw directly against Tagoram and RF-IDraw. The hardware setup is shown in Figure 17.

5.3.1 Word recognition accuracy

We compare these three algorithms' ability on written words recognition. We divide the test words into four groups according to the word length. Within in each group, we randomly select 10 words from the Oxford English Dictionary (O.E.D) [28]. Figure 18 shows the word recognition accuracy of PolarDraw, Tagoram and RF-IDraw. When the word contains two characters, PolarDraw achieves a similar performance with both RF-IDraw and Tagoram, with a recognition accuracy over 91%. As the number of characters in the word increases, the performance of all three algorithms drops slightly. The performance of PolarDraw degrades slightly more than the other two algorithms, but with just two antennas, the minimum recognition accuracy of PolarDraw remains above 75%. With natural language processing techniques, we believe the recognition accuracy can be improved further. Furthermore, we find that the recognition accuracy of RF-IDraw here is lower than reported in [45]. This is due to the lesser number of antennas used in this experiments.




Figure 18— Recognition accuracy for various words.




Figure 19— The CDF of the Procrustes distance between the groundtruth and the trajectory recovered by algorithms.

5.3.2 Trajectory similarity

Here we directly examine the trajectory similarity of handwritten letters recovered by these three algorithms. We randomly choose five letters and invite one volunteer to write each letter 10 times with a 20 cm writing size. Figure 19 shows the cumulative distribution function (CDF) of the Procrustes distance between the recovered trajectories and the ground truth. The four-antenna versions of Tagoram and RF-IDraw achieve similar performance, with 90th percentile errors of 11.3 cm and 10.2 cm, respectively. PolarDraw's two-antenna performance is comparable at 13.8 cm.

We present examples of the pen trajectories recovered by the three algorithms in Figure 20. Compared with ground truth, we can see that all the recovered trajectories are stretched or rotated due to the localization and tracking errors. Comparing the trajectories recovered by these three algorithms, we find that these trajectories are distinct from each




Figure 20— Pen trajectory recovered by three algorithms.




Figure 21— Recognition accuracy across users: User 2 writes with a “stiff” style, rotating the pen only slightly.




Figure 22— PolarDraw's recognition accuracy over different tag-to-reader distance settings.

other, especially at the beginning and the ending part of the pen trajectory. Nevertheless, all of them preserve the basic profile of the handwritten letters.

5.3.3 Recognition accuracy across users

In this experiment we examine the recognition accuracy over different users. Figure 21 shows the resulting recognition accuracy. As the results show, PolarDraw, Tagoram and RF-IDraw all achieve roughly consistent and high performance for all four users, with two-antenna PolarDraw approaching the other two systems' accuracies for User 1. To test the limits of PolarDraw's performance, one of the users (User 2) was instructed to write in an unnaturally “stiff” style on the whiteboard. While slightly diminished, the resulting accuracy of PolarDraw remains at a comparably high level with just two antennas, demonstrating graceful degradation.

5.3.4 Impact of the tag-to-reader distance

We further examine the impact of the tag-to-reader distance on the trajectory tracking performance. In these experiments, we vary the tag-to-reader distance from 20 cm to 140 cm, with a step of 20 cm. Figure 22 shows the recognition accuracy at each distance. The recognition accuracy is relatively low (around 78%) when the tag is close to the reader antenna (e.g., 20 cm spacing). This is because both the polarization angle mismatch and the tag movement contribute to the RSS change. Hence PolarDraw less often correctly estimates the pen's direction of movement. As we increase the spacing between tag and the reader, the RSS tends to be less sensitive to the pen translation movement. Hence, PolarDraw could accurately estimate the pen moving direction, which results in an increased recognition accuracy, as shown in the result. The recognition accuracy drops as

Table 7— PolarDraw's recognition accuracy v. α_e .

| α_e ($^\circ$) | -45 | -30 | -15 | 15 | 30 | 45 |
|-------------------------|-----|-----|-----|----|----|----|
| Accuracy (%) | 91 | 91 | 92 | 91 | 93 | 90 |

Table 8— PolarDraw's recognition accuracy v. γ .

| γ ($^\circ$) | 15 | 30 | 45 | 60 | 75 |
|-----------------------|----|----|----|----|----|
| Accuracy (%) | 92 | 90 | 91 | 85 | 80 |

we further increase the tag-to-reader distance from 100 cm to 120 cm and further to 140 cm. A possible explanation may be that the backscatter signal bounces off nearby objects and changes the polarization angle, which leads to an unpredictable RSS trend during human writing. We plan to investigate that use of more directional antennas to even further increase range in future work.

5.4 Microbenchmarks

These microbenchmark experiments aim to give further insight into PolarDraw's performance. In particular, we seek to understand which factors will most affect tracking performance as we tune various operating parameters, to better understand how polarization angle mismatch helps to improve the pen trajectory tracking accuracy. The hardware setup of our microbenchmarks is shown in Figure 4. In the following experiments, we randomly choose 10 letters from the English alphabet, invite a volunteer to write each letter 10 times under different parameter settings, and measure the resulting letter recognition accuracy.

5.4.1 Choosing elevation angle

We first perform a sensitivity analysis examining how the pen elevation angle α_e (defined in §3.2) affects accuracy. Table 7 shows letter recognition accuracy for different pen elevation angles. As the result shows, PolarDraw achieves similar performance under different α_e settings. Based on these data we set $\alpha_e = 30^\circ$ in all other experiments.

5.4.2 Choosing inter-antenna angle

We now examine how the inter-antenna angle γ (defined §3.3) affects accuracy. In this experiment, we manually align the antenna orientation using a protractor. Table 8 shows letter recognition accuracy under different γ settings. PolarDraw achieves similar recognition accuracy when γ is relatively small (e.g., $\gamma = 15, 30$, and 45°). This is because the pen is prone to cross over the border of neighboring sectors (as shown in Figure 8) when the user writes on the whiteboard, which improves PolarDraw's accuracy. As expected, when γ increases, the probability that the pen crosses over the border of neighboring sectors decreases, resulting in a lower recognition accuracy. Informed by these experimental results, we set $\gamma = 15^\circ$ in all other experiments.

6. RELATED WORK

In this section we compare PolarDraw to prior work in multiple areas, chiefly whiteboard transcription systems, RF-

ID-based systems, and inertial sensor-based systems.

Whiteboard transcription systems.. These use cameras to record a whiteboard’s contents [10, 14, 51]. While they function well in a controlled lecture theater setting where infrastructure can be permanently installed, they depend on sufficient ambient light, and require a mostly unobstructed line of sight from a camera placed at a sufficient distance from the whiteboard, precluding deployment in some scenarios and complicating deployment in most. In contrast, PolarDraw can be deployed simply with two small antennas attached to a whiteboard. Electronic whiteboard technologies based on ultrasound-infrared [25] and laser curtains [24] also exist, but require expensive styluses in comparison with RFID tags, which cost cents apiece. Finally, the whiteboard transcription systems do not generalize to writing in the air, as we have experimentally demonstrated for PolarDraw.

RFID positioning and tracking systems.. Initial attempts in RFID positioning have used the RSS to estimate the distance between the reader and the tag. *PinIt* [44] localizes RFID tags in a non-line of sight environment by exploiting the multipath profile of each tag. *STPP* [37] uses RFID to spatially order tags without localizing them. However, the localization error of these two systems is relatively high (around 12 cm for PinIt, and 8 cm for STPP), hence they are unsuitable for motion tracing. Later work exploits fine-grained phase information of the backscattered signal for more accurate tag localization. *BackPos* [22] is a hyperbolic-based positioning system that requires at least three RFID antennas, but only achieves *ca.* 13 cm accuracy, which is unsuitable for tracking writing. *Tagoram* [52] introduces a holography-based tag tracking algorithm that achieves an accuracy of around 12 cm⁶ based on four reader antennas. In contrast, PolarDraw requires only two antennas and achieves comparable or better accuracy, depending on the number of antennas Tagoram operates with. *RF-IDraw* [45] is a pen-tracking scheme that achieves high accuracy via a non-linear spaced RFID antenna array, but requires a significant number (eight) of physically-separated antennas and rigid placement of antenna locations, reducing deployability. In contrast, PolarDraw requires minimal supporting infrastructure yet achieves high tracking accuracy.

Gesture recognition systems.. *E-Gesture* [30] employs the gyroscope and accelerometer to characterize eight kinds of predefined hand gestures. Similarly, *RisQ* [29] recognizes smoking gestures by using the accelerometer and gyroscope readings from a smart wristband. However, both systems only recognize simple and pre-defined hand gestures like “throw,” “draw,” “smoking” & *c.*, and so are not applicable to fine-grained finger motion or writing tracking.

Fine-grained motion tracking systems.. *PhonePoint Pen* [4] enables in-air human writing by tracking a smartphone’s motion with the gyroscope and accelerometer. *MoLe* [43]

⁶When the device’s location track is unknown *a priori*. When the device’s future path is known beforehand, Tagoram achieves one cm accuracy, but this does not apply to writing.

designs an inference algorithm to detect and identify the human typing through smart watch gyroscope and accelerometer readings. However, it requires a large amount of user data for model training. *WiDraw* [38] enables device-free in-air writing by exploiting changes in the AoA spectrum measured at many nearby Wi-Fi devices caused by human arm motion. However, to achieve five cm accuracy, WiDraw requires the support of 30 neighboring Wi-Fi devices and human writing 25 cm in width, considerably larger than PolarDraw. A Wi-Fi band software-defined radio based system by Chen *et al.* [9] passively tracks keystrokes adapting to interfering environmental motion, but requires five nearby antennas for best accuracy.

FingerIO [27] uses sound to track finger motion without the need for added RFID tags or any mobile device present, to a fine-grained precision of 8 mm. However when there is an interfering bystander in the environment at a range of 50 cm, the authors’ experimental results show an increase in FingerIO’s motion tracking error from 8 mm to 15 cm. On the other hand, our results in §5.2.5 demonstrate that PolarDraw suffers no significant degradation of tracking precision in the presence of interfering human motion at a lesser standoff distance (30 cm). *AAMouse* [55] uses acoustic ranging to track a smartphone’s trajectory with a median error of around 1.4 cm, but high frequency acoustic noises, especially sounds emitted by other smartphone-based styluses interfere with the system. The three-antenna *mTrack* system [48] uses one millimeter-wave transmit antenna and two receive antennas to track a pen with a 90th percentile error below 8 mm, if background motion interference is not present, but accuracy degrades by almost a factor of two to 1.2 cm in the presence of an interfering person walking at a two meter range or standing still at a one meter range ⁷. In contrast, our experiments show that PolarDraw can handle interfering walking-speed motion at a 30 cm range, hence demonstrate more robustness than the mTrack results. Besides, a back of the envelope calculation shows it will cost at least \$7k to implement and is thus not cost effective.

Compared to both acoustic and millimeter-wave tracking, our experiments demonstrate greater range (§5.2.4) and more robustness to nearby interfering motion (§5.2.5).

Wi-Fi localization and tracking.. Many of these proposals achieve meter-level localization accuracy [8, 26, 32, 54], hence are unable to recognize fine-grained human writing. Although later schemes achieve centimeter-level localization accuracy [2, 16, 18, 35, 36, 50, 56], they generally rely on software-defined radio and Wi-Fi chipsets on the mobile, and require more infrastructure than PolarDraw.

Human activity recognition.. *E-eyes* [47] detects and classifies different kinds of human activity at home through the use of commercial Wi-Fi APs. *WiHear* [42] recognizes human speech by analyzing the CSI pattern of Wi-Fi signals reflected by human lips. *WiKey* [5] decodes user input on keyboards by matching Wi-Fi signal patterns to fingerprints. These systems only work for predefined coarse-grained ac-

⁷mTrack’s experiments did not study closer writing to motion interference ranges.

tivities, and hence cannot be directly applied to fine-grained motion tracking and writing recognition since writing style varies from person to person.

7. DISCUSSION AND FUTURE WORK

We discuss limitations and opportunities for improvement.

Trade-off between infrastructure cost and performance. Prior systems such as Tagoram and RF-IDraw require four and eight antennas for tag motion tracking, respectively. In contrast, PolarDraw requires only two antennas to operate, hence reducing the system cost by half compared to Tagoram. While the recognition accuracy of PolarDraw decreases slightly compared to both Tagoram and RF-IDraw, we believe by applying natural language processing techniques, we can further increase recognition accuracy. Furthermore, Tagoram requires a relatively close antenna spacing, so that the tag is within the coverage area of all four antennas. In contrast, since PolarDraw relies on just two antennas, it can cover much larger geographical area and potentially reduce the deployment cost, *e.g.*, classroom whiteboard.

Extending to multi-user case. Although our experiments are conducted in the single-user case, we believe it can be easily extended to multiple users by separating phase readings of each tag. This would involve examining the tag ID and rejecting dynamic multipath phase readings caused by nearby user writings. We leave this for our future work.

Scaling to abrupt hand motions. We assume a smooth hand movement and a relative stable tilt angle, which may not true for all users. We believe with lightweight inertial sensors on the tag (*e.g.*, WISP [49]), it is possible to detect tile angle changes and fetch this angle into the HMM model for a more accurate trail estimation. The tag attached on the pen responds to the reader all the time, hence PolarDraw is unable to distinguish the words from a sentence. However, it is possible to detect whether the pen is touching the whiteboard or not by examining the accelerometer readings of this sensor-based tag. We leave this problem for our future work.

8. CONCLUSION

We have presented the design of PolarDraw, the first motion tracking system that can accurately reproduce and recognize handwritten letters in the air or on a whiteboard with just two nearby antennas. Our experimental results show that compared with other state-of-the-art multi-antenna tracking systems, PolarDraw reduces the infrastructure cost in half with a slightly degraded performance.

Acknowledgements

We sincerely thank our shepherd Prof. Neil Spring and the anonymous reviewers for their insightful comments. This material is based upon work supported by the National Science Foundation under Grant No. 1617161 and European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement No. 279976.

APPENDIX

Modeling the HMM emission probability. By jointly considering pen moving direction and the hyperbola constraint, we define the emission probability in Equation 11. In this equation, $\Delta\theta_{x_1, y_1}^{1,2}$ is the theoretical inter-phase difference of two antennas on location (x_1, y_1) . d_{max} is the maximum moving distance of the pen within the feasible region. This equation takes into account the likelihood that the location (x_1, y_1) is on the hyperbolas computed from $\Delta\theta_i^{1,2}$ and the likelihood that the location (x_1, y_1) is on the pen trajectory:

$$Pr(X_{t+1} = (x_1, y_1) | Y_t) = \left(1 - \frac{|\Delta\theta_t^{1,2} - \Delta\theta_{x_1, y_1}^{1,2}|}{4\pi} \right) \cdot \left(1 - \frac{|-\cot(\alpha_r^t) \cdot x_1 - y_1 + \cot(\alpha_r^t) \cdot x_0 + y_0|}{d_{max} \cdot \sqrt{\cot^2(\alpha_r^t) + 1}} \right) \quad (11)$$

References

- [1] F. Adib, Z. Kabelac, D. Katabi. Multi-person localization via rf body reflections. *NSDI*, 2015.
- [2] F. Adib, Z. Kabelac, D. Katabi, R. Miller. 3D Tracking via body radio reflections. *NSDI*, 2014.
- [3] F. Adib, D. Katabi. See through walls with Wi-Fi! *SIGCOMM*, 2013.
- [4] S. Agrawal, I. Constandache, S. Gaonkar, R. Roy Choudhury, K. Caves, F. DeRuyter. Using mobile phones to write in air. *MobiSys*, 2011.
- [5] K. Ali, A. X. Liu, W. Wang, M. Shahzad. Keystroke recognition using WiFi signals. *MobiCom*, 2015.
- [6] An-900lh antenna (<http://rf-links.com/newsite/pdf/an900lh.pdf>).
- [7] Avery-Dennison UHF RFID inlay *AD-227m5*.
- [8] P. Bahl, V. N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. *Infocom*, 2000.
- [9] B. Chen, V. Yenamandra, K. Srinivasan. Tracking keystrokes using wireless signals. *MobiSys*, 2015.
- [10] P. Dickson, W. Adrion, A. Hanson. Automatic capture and presentation creation from multimedia lectures. *Frontiers in Education Conf.*, 2008.
- [11] Emerald (<http://emeraldforhome.com>).
- [12] Google Project Soli.
- [13] LipiTk: online hand writing recognition toolkit.
- [14] L. He, Z. Zhang. Real-time whiteboard capture and processing using a video camera for remote collaboration. *IEEE Trans. on Multimedia*, 9(1), 198–206, 2007.
- [15] ImpinJ Speedway R420 reader.
- [16] K. Joshi, D. Bharadia, M. Kotaru, S. Katti. WiDeo: Fine-grained device-free motion tracing using RF backscatter. *NSDI*, 2015.
- [17] R. Krigslund, P. Popovski, G. Petersen. Orientation sensing using multiple passive RFID tags. *IEEE*

- Antennas and Wireless Prop. Letters*, **11**, 176–179, 2012.
- [18] S. Kumar, S. Gil, D. Katabi, D. Rus. Accurate indoor localization with zero start-up cost. *MobiCom*, 2014.
- [19] Laird antenna 915MHZ 1/2WV-7 RA RPSMA.
- [20] Laird pa9-12 antenna (<https://www.atlasrfidstore.com/laird-pa9-12-lp-outdoor-rfid-antenna-fcc/>).
- [21] W. Lee, Y. Yeh. Polarization diversity system for mobile radio. *IEEE Trans. on Commms.*, **20**(5), 912–923, 1972.
- [22] T. Liu, L. Yang, Q. Lin, Y. Guo, Y. Liu. Anchor-free backscatter positioning for RFID tags with high accuracy. *Infocom*, 2014.
- [23] LLRP Toolkit.
- [24] MimioProjector Touch Projector.
- [25] MimioTeach Interactive Whiteboard.
- [26] A. Musa, J. Eriksson. Tracking unmodified smartphones using Wi-Fi monitors. *SenSys*, 2012.
- [27] R. Nandakumar, V. Iyer, D. Tan, S. Gollakota. FingerIO: Using Active Sonar for Fine-Grained Finger Tracking. *CHI*, 2016.
- [28] Oxford English Dictionary.
- [29] A. Parate, M.-C. Chiu, C. Chadowitz, D. Ganesan, E. Kalogerakis. Risq: Recognizing smoking gestures with inertial sensors on a wristband. *MobiSys*, 2014.
- [30] T. Park, J. Lee, I. Hwang, C. Yoo, L. Nachman, J. Song. E-gesture: a collaborative architecture for energy-efficient gesture recognition with hand-worn sensor and mobile devices. *SenSys*, 2011.
- [31] Q. Pu, S. Gupta, S. Gollakota, S. Patel. Whole-home gesture recognition using wireless signals. *MobiCom*, 2013.
- [32] A. Rai, K. K. Chintalapudi, V. N. Padmanabhan, R. Sen. Zee: Zero-effort crowdsourcing for indoor localization. *MobiCom*, 2012.
- [33] T. Rappaport. *Wireless Communications: Principles and Practice*. Prentice Hall, second edn., 2002.
- [34] T. Rappaport, D. Hawbaker. Wide-band microwave propagation parameters using circular and linear polarized antennas for indoor wireless channels. *IEEE Trans. on Commms.*, **40**(2), 240–245, 1992.
- [35] S. Sen, R. R. Choudhury, B. Radunovic, T. Minka. Precise indoor localization using PHY layer information. *MobiSys*, 2011.
- [36] S. Sen, B. Radunovic, R. R. Choudhury, T. Minka. You are facing the Mona Lisa: Spot localization using PHY layer information. *MobiSys*, 2012.
- [37] L. Shangguan, Z. Yang, A. X. Liu, Z. Zhou, Y. Liu. Relative localization of RFID tags using spatial-temporal phase profiling. *NSDI*, 2015.
- [38] L. Sun, S. Sen, D. Koutsonikolas, K.-H. Kim. WiDraw: Enabling hands-free drawing in the air on commodity WiFi devices. *MobiCom*, 2015.
- [39] Thinkmagic micro embeded rfid reader module (<https://www.atlasrfidstore.com/thingmagic-micro-embedded-rfid-reader-module/>).
- [40] Thinkmagic m6e micro embeded rfid reader module (<https://www.atlasrfidstore.com/thingmagic-m6e-embedded-rfid-reader-module/>).
- [41] D. Vasisht, S. Kumar, D. Katabi. Decimeter-level localization with a single WiFi access point. *NSDI*, 2016.
- [42] G. Wang, Y. Zou, Z. Zhou, K. Wu, L. M. Ni. We can hear you with Wi-Fi! *MobiCom*, 2014.
- [43] H. Wang, T. Lai, R. Choudhury. MoLe: Motion leaks through smartwatch sensors. *MobiCom*, 2015.
- [44] J. Wang, D. Katabi. Dude, where’s my card? RFID positioning that works with multipath and non-line of sight. *SIGCOMM*, 2013.
- [45] J. Wang, D. Vasisht, D. Katabi. RF-IDraw: Virtual touch screen in the air using RF signals. *SIGCOMM*, 2014.
- [46] W. Wang, A. Liu, M. Shahzad, K. Ling, S. Lu. Understanding and modeling of WiFi signal based human activity recognition. *MobiCom*, 2015.
- [47] Y. Wang, J. Liu, Y. Chen, M. Gruteser, J. Yang, H. Liu. E-eyes: device-free location-oriented activity identification using fine-grained WiFi signatures. *MobiCom*, 2014.
- [48] T. Wei, X. Zhang. mTrack: High-precision passive tracking using millimeter wave radios. *MobiCom*, 2015.
- [49] Wisp 5 (<https://wisp5.wikispaces.com/wisp+home>).
- [50] J. Xiong, K. Jamieson. ArrayTrack: A fine-grained indoor location system. *NSDI*, 2013.
- [51] D. Xu, R. Yi. A computer vision based whiteboard capture system. *IEEE WACV*, 2008.
- [52] L. Yang, Y. Chen, X.-Y. Li, C. Xiao, M. Li, Y. Liu. Tagoram: Real-time tracking of mobile RFID tags to high precision using COTS devices. *MobiCom*, 2014.
- [53] Yap-100cp antenna (<http://sbrfid.com/yap100cp.pdf>).
- [54] M. Youssef, A. Agrawala. The Horus WLAN location determination system. *MobiSys*, 2005.
- [55] S. Yun, Y.-C. Chen, L. Qiu. Turning a mobile device into a mouse in the air. *MobiSys*, 2015.
- [56] Y. Zhu, Y. Zhu, B. Zhao, H. Zheng. Reusing 60GHz radios for mobile radar imaging. *MobiCom*, 2015.