



Pushing the Range Limits of Commercial Passive RFIDs

Jingxian Wang, *Carnegie Mellon University*; Junbo Zhang, *Tsinghua University*;
Rajarshi Saha, *IIT Kharagpur*; Haojian Jin and Swarun Kumar, *Carnegie Mellon University*

<https://www.usenix.org/conference/nsdi19/presentation/wangjingxian>

This paper is included in the Proceedings of the
16th USENIX Symposium on Networked Systems
Design and Implementation (NSDI '19).

February 26–28, 2019 • Boston, MA, USA

ISBN 978-1-931971-49-2

Open access to the Proceedings of the
16th USENIX Symposium on Networked Systems
Design and Implementation (NSDI '19)
is sponsored by



Pushing the Range Limits of Commercial Passive RFIDs

Jingxian Wang
Carnegie Mellon University

Junbo Zhang
Tsinghua University

Rajarshi Saha
IIT Kharagpur

Haojian Jin
Carnegie Mellon University

Swarun Kumar
Carnegie Mellon University

Abstract

This paper asks: “Can we push the prevailing range limits of commercial passive RFIDs?”. Today’s commercial passive RFIDs report ranges of 5-15 meters at best. This constrains RFIDs to be detected only at specific checkpoints in warehouses, stores and factories today, leaving them outside of communication range beyond these spaces. State-of-the-art approaches to improve the range of RFIDs develop new tag hardware that necessarily sacrifices some of the most attractive features of passive RFIDs such as their low cost, small form-factor or the absence of a battery.

We present PushID, a system that exploits collaboration between readers to enhance the range of commercial passive RFID tags, without altering the tags whatsoever. PushID uses distributed MIMO to coherently combine signals across geographically separated RFID readers at the tags. In doing so, it resolves the chicken-or-egg problem of inferring the optimal beamforming parameters to beam energy to a tag without any feedback from the tag itself, which needs this energy to respond in the first place. A prototype evaluation of PushID with 8 distributed RFID readers reveals a range of 64-meters to the closest reader and a $7.4\times$, $1.2\times$ and $1.6\times$ improvement in range compared to state-of-the-art commercial readers and other two schemes [10, 33].

1 Introduction

Conventionally, passive commercial RFID tags have a maximum range of about 5-15 meters. Passive RFIDs are limited in range owing to their limited cost, form-factor, and the FCC-mandated power limits of the RFID readers they harvest energy from. Indeed, in much of today’s factories [30] and warehouses [7, 11], RFID-tagged products can only be detected around specific checkpoints in the vicinity of an RFID reader, and are virtually undetectable at other points in between [21, 31]. Further, recent innovation on RFID-based localization and sensing [24, 25, 35, 56] remain constrained to a few meters around each reader in these large spaces.




Figure 1: (a) In traditional multi-antenna beamforming, a beam forms towards a particular direction; (b) PushID’s distributed MIMO beams create multiple energy peaks and troughs over space resulting in complex energy distributions.

In this paper, we ask – “Can we push the range limits of today’s commercial passive RFID tags without increasing the prevailing density of deployment of RFID readers?”. In particular, we aim to do so without modifying the RFID tags in any way by adding to cost and complexity [27, 29, 49], relays [34] or requiring batteries [37]. We further avoid using sophisticated multi-antenna or directional RFID readers that can indeed expand range [10, 33], yet are vulnerable to obstacles and limited by FCC transmit power limits [1]. More importantly, such systems require commercial RFIDs that are linearly polarized to be carefully oriented towards their location [23] to harvest sufficient energy.

We present PushID, the first system that uses distributed MIMO to increase the communication range of commercial passive RFID tags. PushID synchronizes both transmissions and receptions from multiple, distributed RFID readers to beamform power to RFID tags that are several tens of meters away from any individual reader antenna. In addition, PushID exploits the diversity in location and polarization of reader antennas to further improve range. Our detailed experimental evaluation on an eight-distributed multi-antenna reader testbed reveals that PushID achieves a $7.4\times$ improvement in range compared to the state-of-the-art commercial readers and 20% over the mean of distance improvement even when

compared to 8-antenna MIMO reader arrays [10], all while remaining compliant with FCC power limits for the readers.

PushID's main goal is to find optimal beamforming weights to beam power to an RFID tag at an unknown location and orientation. Since RFID systems are back-scatter [4], optimal beamforming weights amplify *both* the transmitted and received signals to/from RFID tags. At first blush, one may consider using channel reciprocity [19], where one infers the optimal beamforming weights based on the wireless channels of signals from the RFID tag to the reader antennas. However, in the context of passive RFIDs, this leads to a chicken-or-egg problem. To emit a response, a passive RFID tag needs to harvest sufficient energy from the beamformed signal of the RFID reader antennas. Yet to perform accurate beamforming, the readers need a response from the RFID tag in the first place. Indeed, naively iterating over all possible beamforming weights would take days to simply beamform power to one tag from a long distance.

PushID resolves this dilemma by developing a novel distributed blind beamforming approach to efficiently search through the space of beamforming weights without a response from the tag. At a high level, PushID models the received signal power at each point in space for different beamforming weights applied across reader antennas. It then identifies subsets of 3-D space where RFIDs placed would receive sufficient energy to respond. Unlike traditional beams of a directional antenna, in the context of distributed MIMO, these regions of space that receive sufficient energy are quite complex and they span the entire 3-D space as shown in Fig. 1. PushID shows that finding the optimal beamforming weights while minimizing overlap is analogous to the weighted set-cover problem, a well-known NP-complete problem [5]. PushID then develops heuristic approximation algorithms to efficiently search the entire space for RFID tags under a limited time budget. A key challenge in ensuring minimal overlap between patterns is the unknown multipath characteristics of the environment which can completely change the energy patterns that beamforming weights produce (See Fig. 3). PushID's approach to resolve this exploits responses from the RFID tags that are progressively detected as it applies various beamforming vectors. PushID uses these responses to better learn the nature and extent of multipath-richness in the environment. We show how this iteratively improves PushID's ability to efficiently look for and power other tags in the environment.

A second challenge PushID must resolve is achieving time and frequency synchronization across multiple distributed RFID readers to beamform coherently. PushID borrows from classic distributed MIMO architectures in the Wi-Fi context [19, 40] that one can treat a transmitter as the master and apply phase shifts to the remaining slave transmitters to emulate signals from the master. Yet, a key challenge in the RFID context is that transmissions from the readers are significantly longer than Wi-Fi, causing phase drifts to accumulate significantly even within one packet from the reader. PushID

resolves this by leveraging the full-duplex nature of RFID readers. Specifically, each PushID slave transmitter subtracts its own signal and tracks the drift in phase of a carefully chosen subset of remaining transmitters. It then uses these phase drifts to account for phase errors that accumulate within a packet dynamically. We show how our system eventually converges to tightly synchronized transmissions and receptions, even if some RFID readers are not in the range of the master reader.

Limitations: We note that PushID has two important limitations common to RFID-systems: (1) First, despite significant range improvements, a small fraction of RFID tags ($< 5\%$) are missed due to extreme shadowing or poor orientation; (2) Second, while PushID can handle modest mobility of tags (walking speeds), it struggles at higher speeds due to high dynamism in the multipath characteristics of the tags. We discuss and evaluate these limitations in Sec. 7.5-7.6.

We implement PushID on eight USRP N210 software radio readers, each connected to separate Jackson Lab Fury clocks and commercial Alien passive RFID tags. We perform our experiments on a 140 x 140 meter outdoor space and a 20 x 40 meter indoor space in both line-of-sight and non-line-of-sight scenarios. Our experimental results reveal that:

- PushID achieves a maximum communication range of 64 m, an improvement of $7.4\times$ that of commercial RFIDs and 20% over expensive 8-antenna MIMO.
- Even at short range, PushID achieves a mean throughput of 300 kbps at 8.5 meters ($2.6\times$ vs. commercial RFID).
- Our system detects over 95% of the tags in a 140×140 m area, while commercial readers can detect tags no further than 8.5 m at best.

Contributions: To our knowledge, this paper presents the first distributed MIMO system to power commercial passive RFID tags. We present a novel blind distributed beamforming algorithm to efficiently search through the space of beamforming weights as well as novel phase synchronization for RFIDs. A detailed prototype evaluation on an eight-antenna distributed reader testbed reveals a $7.4\times$ improvement in range compared to state-of-the-art commercial readers .

2 Related Work

RFID Communication and Sensing Systems: There has been much past research on RFID tags including ways to mitigate collisions [6], improve communication speed [41] and design a variety of localization and sensing solutions [24, 53]. However, all these solutions are limited to a range of at most 5-15 meters [32, 52] around the readers and thus have limited ability to locate, sense or communicate with RFID tags. Closely related to our system are recent solutions that use multi-antenna arrays connected to RFID readers to improve

range [10, 33]. While these systems improve range to commercial RFID tags (distances of up to 38 meters), our system varies in two ways: (1) First, systems with directional antennas are vulnerable to obstacles between the reader antennas and tags which can significantly attenuate the signal; (2) Second, they fail when the RFID tags are oriented poorly relative to the reader antennas. We see our system as an alternative approach using diversity of spatially distributed antennas within RFID ISM frequencies to extend the range of RFID and as complementary to [33]. We do think that both PushID and multi-antenna MIMO [42] as attractive solutions with different deployment cost/requirements. We show how by synchronizing signals across multiple RFID readers, PushID exploits the diversity in location and orientations of the reader antennas to significantly improve range compared to the state-of-the-art.

Wireless Power Transfer: Recent advances in wireless power transfer take two approaches: non-radiative near-field coupling and far-field RF radiation. Near-field coupling [22, 43] uses multiple coils to expand charging and communication range, yet is restricted to several tens of centimeters.

Far-field charging promises longer distance for wireless power transfer on the order of meters through innovative hardware design. Energy-harvesting RFID tags such as the WISP [46] promise a range of 18 meters. Past solutions [15, 16] have built wireless power transfer systems to deliver an enhanced energy in a targeted area. Recent work [48, 51] also demonstrates distances of kilometers by building custom backscatter tags with small batteries. Low-power WAN (LP-WAN) technologies [13, 14] including LoRa [47], SIGFOX [59] have explored battery-powered tags that can communicate over miles. In contrast, PushID strives to improve the range of current battery-free commercial RFID tags that cost a few cents by innovating at the RFID readers.

Blind Beamforming: In the RFID context, PushID needs to perform beamforming to tags whose wireless channels are unknown a priori. PushID builds upon blind beamforming [9], which is a set of theoretical beamforming solutions developed by the signal processing community in the presence of poor quality or even no channel state information available from the clients. Past work on blind beamforming with weak channel responses leverages its statistical features, for example, the cyclostationary property [45], spectral self-coherence [2]. In contrast, PushID does not have any channel information to leverage, given that tags are not a priori powered up. Other work on blind beamforming with zero channel feedback builds solutions without the need for carrier frequency synchronization [8, 33, 44]. PushID builds upon this past work but overcomes unique system-level challenges pertaining to the RFID context: (1) It accounts for feedback from neighboring RFID tags that are charged during the exploration of beamforming weights to refine the search in multipath-rich settings; (2) It deals with various synchronization challenges in distributed




Figure 2: Architecture of PushID

beamforming with long RFID packets.

3 Overview

PushID aims to power and communicate with commercial passive RFID tags via RF-backscatter from a team of distributed commercial RFID readers, where tags are beyond the communication range of any single reader. PushID achieves this by coherently combining signals across distributed RFID readers to maximize received signal power at one or more RFID tags, whose location and orientation are a priori unknown. We note that since RFID systems are back-scatter, PushID applies beamforming weights *both* on the transmitted and received signals to maximize energy to/from RFID tags. We assume all PushID RFID readers are connected to a wired backhaul, which allows them to coordinate transmissions and data that needs to be transmitted on the downlink. We further assume that while the locations of the RFID readers are known, the number of RFID tags, their locations and their environment are unknown.

At a high level, PushID’s system design is as follows (see Fig. 2): All RFID readers time and phase synchronize their transmissions on the air and iteratively apply various beamforming vectors in the hope of receiving responses from RFID tags in the environment. The readers then collect responses from various RFID tags in the environment and use the wireless channels they perceive in improving the search for other tags. The readers continue this process until they believe (with sufficient confidence) that they have covered the entire desired coverage area. To achieve this, PushID optimizes the following (related) properties: (1) Maximize the total number of RFIDs found within the coverage area under an overall time budget (which limits the number of beamforming vectors that can be iteratively attempted); (2) Maximize the throughput of signals from each tag.

The rest of the paper addresses the key challenges in designing the two main aspects of PushID’s architecture:

(1) Searching through the Beamforming Space: First, our system needs to identify optimal beamforming weights to

iteratively search over in order to power RFID tags in the entire space. While this problem can be trivially addressed using channel reciprocity for RFID tags that are in range, the key challenge is that tags are outside the coverage area of any single reader. Therefore, PushID must identify the smallest set of beamforming weights that can provide sufficient energy to all tags over the entire area of interest. The key to this is to effectively model multipath in the environment, which would change the set of beamforming weights to search over. Sec. 4 describes our approach.

(2) Synchronizing Distributed RFID Readers: Second, PushID should efficiently synchronize RFID readers that are spatially distributed indoors, without a shared clock between them. In contrast to past work in the Wi-Fi domain [40], our key challenge comes from the longer duration of RFID transmissions, over which packets can quickly lose synchronization in phase. Further, RFID transmissions, unlike Wi-Fi, are narrowband, which makes time synchronization challenging as well. Sec. 5 addresses these challenges.

4 Blind Distributed Beamforming

This section describes how PushID enables a team of RFID readers with an arbitrary geometry to find the optimal beamforming weights and beam power to all RFID tags in their coverage area, including those beyond the range of any single reader. We aim to achieve this without any response from these RFID tags to begin with or prior knowledge of their locations and orientations. For ease of exposition, this section assumes that all RFID readers experience no time, carrier frequency and phase offsets. We will explicitly deal with synchronizing distributed RFID readers in Sec. 5.

We make a few key observations about the scope and goals of PushID's approach:

- **Beamforming on both Downlink and Uplink:** We note that PushID seeks to amplify the received signal power from RFID readers to the tags and vice-versa. Specifically, since RFIDs operate on RF-backscatter, and owing to channel reciprocity [28], beamforming weights used on the transmit chain to power tags on the downlink can also be used to amplify their received signals on the uplink. For simplicity, the rest of this section discusses PushID in the context of maximizing energy on the downlink.
- **How much energy to beamform?:** We emphasize that the goal of PushID in this section is to simply beamform enough energy to detect an RFID tag. Once the RFID tag's response is received, we can then use the reciprocal channel to obtain the optimal beamforming vector to maximize data rate to that tag in future transmissions. As a result, the rest of this section will favor PushID formulations that maximize the energy RFIDs require over the entire space, as opposed to focusing on individual tags.

4.1 Exploring the Beamforming Space

A naïve approach to perform blind beamforming is an exhaustive search of the space of beamforming vectors in hope of eliciting a response. Prior work in the context of cellular multi-antenna arrays [36] constructs codebooks that progressively steer the beam along various discrete directions in hope of covering an entire cell efficiently. For example, consider a phased array of antennas (see Fig. 1(a)) where a transmitter could simply beam power iteratively along discrete angles to cover the space of interest.

However, the distributed RFID context makes such an elegant design challenging. First, given that RFID readers are widely separated and they form an arbitrary geometry relative to each other, beamforming weights distribute energy over the space in very complex patterns. Further, it is challenging to find weights that both cover the entire space of interest with sufficient energy and minimize overlap.

What do beamforming energy patterns look like?: To better understand how different beamforming weights from a distributed array of antennas impact the distribution of energy over the space, we perform a simple simulation. We consider four transmitters in the corners of a square with a one meter diagonal length. For simplicity, we consider that the transmitters are in 2-D free space and use standard wireless channel models [50]. We first apply a beamforming weight that allows signals from the transmitters to add up coherently at the center of the square. We plot the distribution of energy over the entire 2-D space around the square encompassing the transmitters in Fig. 3 (a). We notice that while the center of the square indeed receives maximum energy (denoted by bright yellow), there are multiple spots around the center that are also energized with a similar received signal strength. This shows that applying beamforming weights in a distributed array also focuses energy on unintended points in the space. This means that simply iterating beamforming weights to focus on individual points in the space would lead to much unwanted overlap and be grossly inefficient. We therefore conclude that:

OBSERVATION 1: PushID must seek to minimize overlap between energy patterns of beamforming weights it applies.

Problem Formulation and Optimization: Based on the above observation, we will now formulate PushID's core optimization problem that seeks to find the smallest group of beamforming vectors which energizes the entire space of interest with minimum overlap between them. At least to begin with, our system cannot rely on any feedback from RFID tags in the environment, given that none of them may have sufficient energy to respond. As a consequence, we have no prior information on the nature and extent of multipath in the environment. The rest of this section therefore assumes that line-of-sight paths to RFID readers dominate all other paths, and we explicitly account for multipath in Sec. 4.2.

At a high level, our approach shows that choosing the opti-




Figure 3: (a) Energy pattern when four transmitters focus beamforming at the center. (b) Energy pattern when we apply a phase shift ($\frac{2}{3}\pi$) to one of the transmitters. (c) Energy pattern when a strong reflector is placed along y-axis.

mal set of beamforming vectors is analogous to a well known combinatorial problem: the weighted set coverage problem. To see how, let us imagine that the 3-D space is divided into a grid of discrete blocks. Each beamforming vector effectively supplies sufficient energy to RFIDs in some subset of these blocks. Our goal is to find the smallest set of such beamforming vectors whose union is the universal set of all blocks in the grid. This is analogous to the weighted set cover problem, which seeks to find the smallest number of sets, each containing a few integers in the range $1, \dots, N$ whose union is precisely the universal set $\{1, 2, \dots, N\}$. Given that weighted set cover is NP-complete [55], we propose an efficient approximation algorithm, while presenting various optimization to reduce algorithmic complexity.

Mathematically, let us assume that the space of beamforming vectors has n discrete elements $\mathbf{B} = \{B_1, \dots, B_n\}$, and we aim to cover m discrete points in the space spanning the desired coverage area denoted by $\mathbf{G} = \{G_1, \dots, G_m\}$. Let the variable u_{ij} , $i = 1, \dots, n$ and $j = 1, \dots, m$ be one if the beamforming vector B_i provides energy to the point in space G_j . Given that we assume the L base station locations are known and no information on multipath is available (we discuss multipath in Sec. 4.2), we can determine u_{ij} as follows:

$$u_{ij} = \begin{cases} 1, & |B_i \cdot h_j|^2 > \tau \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

$$\text{where } h_j = \left[\frac{1}{d_{lj}} e^{-2\pi\sqrt{-1}\frac{d_{lj}}{\lambda}}, l = 1, \dots, L \right] \quad (2)$$

Where h_j is the vector of wireless channels from base stations to grid point j , λ is the wavelength, d_{lj} is the distance between the l^{th} base station and j^{th} grid point and τ is the minimum received energy required to energize an RFID tag.

Our objective is to find the smallest set of beamforming vectors that spans all m points in \mathbf{G} . We can state this mathematically as the following integer linear program based on the variable x_i which is 1 if and only if the i^{th} beamforming

vector is included in our optimal set:

$$\min \sum_{i=1}^n x_i$$

$$\text{s.t. } \sum_{i=1}^n x_i u_{ij} \geq 1 \quad \forall j \in \{1, \dots, m\} \quad (3)$$

$$x_i \in \{0, 1\} \quad \forall i \in \{1, \dots, n\} \quad (4)$$

The above formulation directly resembles the well-known weighted set cover problem [55], which processes a group of sets to find the smallest sub-collection whose union is also the union of the original group of sets. While this problem is known to be NP-complete [26], a reasonable polynomial-time approximation algorithm is to relax the above integer-program formulation into a linear program (LP). Specifically, we replace Eqn. 4 above by the following:

$$0 \leq x_i \leq 1 \quad \forall i \in \{1, \dots, n\} \quad (5)$$

We solve the optimization problem using standard linear programming [3] to obtain the optimal set $\{x_1^*, \dots, x_n^*\}$. We then output the chosen set of beamforming weights by applying randomized rounding [38] on the beamforming weights. This technique interprets the fractional part of the solution to the linear program as a probability distribution and then selects a solution by sampling this distribution. Randomized rounding is known to return a set of beamforming weights that is a valid set cover with probability at least $1/2$ [38]. Mathematically, to bound the probability, let ρ be a constant that satisfies: $e^{-\rho \log n} \leq \frac{1}{4n}$. Then randomized rounding proceeds for exactly $\rho \log n$ iterations, and in each iteration, it picks i^{th} beamforming weight with probability dictated by its respective solution x_i to the linear program.

Prior work has shown that the above approximation algorithm results in a set of beamforming vectors whose size is within a factor $O(\log n)$ of the optimum [17]. Our implementation uses the Ellipsoid LP-solver [57] with a worst-case complexity of $O(n^4)$, where n is the number of discrete beamforming vectors PushID's algorithm optimizes over.

Reducing Complexity and Search Space: To reduce the complexity of PushID, one must actively seek to reduce n , the number of beamforming weights that PushID considers in its optimization. Our key insight to this end is that while a large number of beamforming vectors are available, not all are created equally. To see why, let us revisit our example of the energy pattern from a beamforming vector that focuses energy at the center of a square in Fig. 3(a). We now slightly perturb the beamforming vector of one of the transmitters by choosing one of the transmitters and adding $\frac{2}{3}\pi$ to its phase, and we plot the updated energy pattern in Fig. 3(b). We make two observations about the resulting energy pattern. First, each local maximum of energy moves in different, complex ways. This is precisely why we need the optimization algorithm above to minimize overlap. Second, the size of each energized

region changes with maximum diffusing energy over wider spots. In practical terms, this means that the same amount of energy is spread out over a wider area than the previous case. Spreading energy over a wider space is good in that RFID tags over a wider region can be covered by a single beamforming weight. Yet, spreading energy over too thin over a wide area is likely to make the energy per unit area insufficient to activate tags. Indeed, the most ideal beamforming weights are those who diffuse energy in a Goldilocks zone between these two extremes. We observe the following:

OBSERVATION 2: PushID must favor beamforming weights with maximal total area where RFID tags remain powered.

PushID therefore aims to search over beamforming weights that meet the above criterion of maximizing area-of-coverage for RFID tags. Our approach begins with n beamforming weights chosen randomly, where n is dictated by available computing power. For each beamforming weight, we make incremental phase shifts and measure the gradient of net increase in coverage area. We then apply a gradient-based optimization that favors phase shifts which maximize coverage area. We implement Adadelta [58] to speed up the learning rate. The below algorithm summarizes our approach.

Algorithm 1 Gradient-Based Beamforming Vector Pruning

- 1: \mathbf{B} : random beamforming vectors. $t = 1, \dots, N$. D_i : energy of the set of points which could be covered by the i -th beamforming vector B_i .

Loop:

- 2: $Q \leftarrow |\bigcup_i^n D_i|$, where Q represents the number of points covered in \mathbf{G} by the beamforming vector set \mathbf{B}
- 3: $g^{(t)} \leftarrow \nabla_{\mathbf{B}}(Q)$
- 4: $E[g^{(t)}] \leftarrow \gamma E[g^{(t-1)}] + (1 - \gamma)g^{(t)^2}$
- 5: $\Delta \mathbf{B}^{(t)} \leftarrow -\frac{RMS[\Delta \mathbf{B}^{(t-1)}]}{RMS[g^{(t)}]} g^{(t)}$
- 6: $\mathbf{B} \leftarrow \mathbf{B} + \Delta \mathbf{B}$

while $Q < threshold$

return $arg\ max_{\mathbf{B}} Q$

We then feed the above n set of beamforming weights into our optimization algorithm to find the optimal set.

Design Decisions: We emphasize a few key design choices:

(1) **HOW FINELY TO DIVIDE SPACE?:** First, PushID must choose discrete points \mathbf{G} in the space to capture the area covered by a beamforming vector. Choosing too few would lead to coverage holes, while choosing too many would waste computation. PushID therefore samples the space at an interval empirically measured to be below the minimum distance between two adjacent energized regions across beamforming vectors in \mathbf{B} . We empirically find that this corresponds to a sampling distance of $\lambda/3$ in our experiments.

(2) **PICKING THE ENERGY THRESHOLD:** PushID chooses the threshold τ empirically by measuring the smallest amount of energy needed for an RFID tag to respond at its smallest




Figure 4: We simulate 10 transmitters deployed on a circle with a radius of 75 meters. Red points are the location of omni-directional transmitters. White points indicate the area that can activate the RFID tags (signal strength of energized area $\leq -12.8\text{dBm}$), black points are the opposite. From left to right, the plots represent energized pattern with 1, 330 and 450 beamforming vectors. The corresponding energized area is 34%, 80%, 97% of the enclosed area of transmitters.

data rate. Once the RFID tag is detected, future transmissions can use the reciprocal channel measurements from this tag to speed up data rates [50]. We note that τ must be calibrated conservatively to support all RFID tag models in the space.

(3) **IMPACT OF ORIENTATION:** PushID explicitly accounts for the reader’s antenna gain and polarization across spatial directions by applying a weight to each term of Eqn. 2: α_{lj} , which captures the attenuation in the l^{th} base station antenna when it faces the j^{th} grid point. We also account for the orientation of the tag by setting τ conservatively to the smallest amount of energy for a tag to respond should it be oriented most unfavorably relative to the readers.

(4) **RUN TIME:** We find that PushID’s run-time is primarily bottlenecked by the slow beamforming switch time of transmitters (4.5 ms for our hardware) as opposed to computation. PushID’s run time depends directly on the final number of beamforming vectors PushID must iterate over. This depends on the size of the space, placement of base stations and multipath. To get a sense for expected run time, we simulate ten RFID readers in a circle of radius 75 m (Fig. 4). We consider a threshold of $RSSI > -12.8\text{ dBm}$ for the RFID tags to respond. Without PushID, we find that only 33% of the total area of interest is covered. However, after only 450 iterations we find that nearly 97% of the area is covered by PushID. This maps to a total time of 2.0 s for 97% coverage. Sec. 7.6 discusses these tradeoffs in experiments (we observe 4 s run time due to multipath and change in layout).

(5) **ADAPTING TO NEW INFORMATION:** Our approach so far arrives at a static set of beamforming weights. However PushID can benefit from the new information in the channel response of RFID tags, as they are detected. In particular, we are interested in learning about the extent and nature of multipath, which can impact the optimal set of beamforming vectors. Sec. 4.2 below deals with this explicitly.




Figure 5: A reflected path can be modeled by a virtual source (r_1, θ_1, ψ_1) – the mirror image of the tag about the reflector.

4.2 Accounting for Multipath

While so far our discussion considers free space, the presence of multipath can considerably change the set of beamforming weights to efficiently search over a given area. To see why, we revisit our example in Fig. 3(a), and this time we add a strong reflector and re-evaluate the energized regions of space as shown in Fig. 3(c). We notice that the resulting energy heatmap varies considerably from the free-space heatmap, both in the number, size and placement of the hotspots. Note that the same set of reflectors can influence the energy perceived at different tags differently. As a result, we conclude:

OBSERVATION 3: PushID must account for multipath given its impact on the optimal set of beamforming vectors.

PushID’s high-level approach to do so uses the responses from RFID tags in the environment that are progressively detected. Indeed, in the absence of any response, PushID has no information about multipath to work with, and therefore assumes a free-space channel. As responses from RFID tags are collected, PushID progressively computes information of the location and orientation of dominant reflecting surfaces in the environment. It then uses this information to update its optimization algorithm, specifically, it modifies the energy patterns corresponding to our beamforming weights.

Finding Dominant Reflectors: To compute the location of dominant reflectors, PushID adapts the MUSIC algorithm while accounting for the arbitrary geometry of the RFID readers. Specifically, the algorithm takes as input wireless channels along the various frequencies of operation of an RFID tag (note that RFID tags naturally hop through a range of frequencies in the 900 MHz ISM band). It then measures the polar coordinates: (r, θ, ϕ) representing the mirror-image of the RFID tag along dominant reflectors by measuring the power of the received signal $P(r, \theta, \phi)$ of signals received from these coordinates. Mathematically, we write:

$$P(r, \theta, \phi) = \frac{1}{|a(r, \theta, \phi)^\dagger E_n E_n^\dagger a(r, \theta, \phi)|} \quad (6)$$

$$\text{where: } a(r, \theta, \phi) = [e^{4\pi j |r-r_i| \cos(\theta-\alpha_i) \cos(\phi-\beta_i)/\lambda}]_{i=1, \dots, N}$$

Where (r_i, α_i, β_i) are the polar coordinates of the transmitters, λ is the signal wavelength, j is the square root of -1 ,

E_n are the noise eigenvectors of $h_{obs} h_{obs}^\dagger$, h_{obs} represents the vector of observed wireless channels of the tags and $(\cdot)^\dagger$ is the conjugate-transpose operator. Our algorithm computes the top- s ($s = 5$ in our implementation) local maxima in $P(r, \theta, \phi)$ to define the set of polar reflector coordinates: $\{(r_k, \theta_k, \phi_k), k = 1, \dots, s\}$.

Folding Multipath into the Optimization: At this point, we model how the energy patterns of beamforming weights change due to our knowledge of reflecting surfaces. We use a ray-tracing model [50] to account for how multipath changes received signal power. Mathematically, we rewrite Eqn. 2 in the definition of u_{ij} as:

$$h_j = \left[\sum_{k=1}^s 1/d_{ljk} e^{-2\pi\sqrt{-1}d_{ljk}/\lambda}, l = 1, \dots, L \right] \quad (7)$$

Where λ denotes the wavelength and d_{ljk} represents the distance traversed by the ray emanating from base station l to grid point j when reflecting off reflector at coordinates (r_k, θ_k, ϕ_k) . This formulation effectively removes the free-space assumption in our optimization to explicitly account for ambient reflectors.

Accounting for past vectors: We note that as new information about multipath emerges, one must account for how this impacts the coverage area of beamforming vectors used previously and therefore invoke the optimization to fill gaps in coverage. Mathematically, let us denote x_i^* as an indicator function on which beamforming weights were used previously. Then we can rewrite Eqn. 3 as:

$$\sum_{i=1}^n x_i u_{ij} \geq 1 - \sum_{i=1}^n x_i^* u_{ij} \quad \forall j \in \{1, \dots, m\} \quad (8)$$

Modeling fleeting and small reflectors: While the above formulation assumes reflectors impact all RFIDs in the coverage area equally, in practice, this may not be the case. Specifically, reflectors have a higher probability of impacting nearby RFIDs compared to RFID tags that are further away. Similarly, reflectors that were computed in the past may no longer exist at the same location and orientation in the future. To account for these effects, PushID employs the exponential weighting method [12] to progressively reduce the contribution of reflectors to the optimization with increasing distance from the reflector or time elapsed since detection. Specifically, we re-write Eqn. 7 as:

$$h_j = \left[\sum_{k=1}^s w_k / d_{ljk} e^{-2\pi\sqrt{-1}d_{ljk}/\lambda}, l = 1, \dots, L \right] \quad (9)$$

Where $w_k = f_1^{d_{ljk}} f_2^{t_k}$ and $f_1, f_2 < 1$ are constants (empirically set to 0.9 in our implementation) and t_k is the time elapsed since the measurement of reflector k was made.

5 Distributed Synchronization for RFIDs

In this section, we consider a classic challenge for distributed MIMO systems: accurate time and frequency synchronization with multiple distributed RFID readers for coherent beam-forming. In particular, we actively estimate and correct for carrier frequency and timing offsets, which would otherwise cause transmissions across RFID readers to often combine incoherently. We build upon the classic distributed MIMO architecture used in the context of Wi-Fi [19, 40] while accounting for new challenges in the RFID context.

Quick Primer on Distributed MIMO: At a high level, past distributed MIMO systems for Wi-Fi [19, 40] use a master-slave architecture where multiple slave transmitters attempt to transmit in-phase with a master transmitter. Prior to transmitting each data packet, the master sends a short beacon containing a known preamble. Slave transmitters estimate their phase relative to this beacon to account for frequency offsets between the master’s clock and their own clocks. In addition, slaves exploit the relatively wide bandwidth of Wi-Fi to estimate phase shifts due to timing offsets. Slaves then apply phase shifts compensating for these offsets when they transmit data packets in-tandem with the master. Of course, during the data packet transmission itself, small additional phase drifts can accumulate owing to residual time and frequency offsets. As the duration of packets is short, the slope of such phase drifts can be readily corrected for.

Challenges in the RFID context: RFIDs bring two important challenges for distributed MIMO: (1) First, RFID packets last for a much longer time ($\sim 100\times$) than Wi-Fi packets [40], ensuring greater phase drift due to frequency offset. This is because tags need to harvest enough energy in order to respond to the readers’ queries, and this takes more time as the distance between the RFID tags and readers increases. (2) Second, RFID transmissions are narrowband (20 kHz), meaning that resolving timing offsets is extremely challenging.¹ The rest of this section tackles each of these challenges.

5.1 Frequency Offset Compensation

PushID’s key idea to compensate for frequency offsets leverages the full-duplex nature of RFID readers. Specifically, each PushID slave cancels out its own signal to recover the signal from the master reader. By measuring how the phase of this node drifts over time, PushID can correct for phase drifts that accumulate since the initial synchronization.

Correcting for Drift: For simplicity, let us begin with the case of two RFID readers – one master and one slave reader. To achieve that, initial synchronization in phase, we borrow from MegaMIMO [40]’s phase synchronization protocol

¹Past proposals for Wi-Fi like SourceSync [39] cannot be directly used in the narrow-band RFID context, as they assume availability of phase measurements over wide bandwidths (20 MHz) to estimate timing offsets.




Figure 6: PushID constructs a spanning tree of the RFID readers in a distributed manner to guide synchronization.

where the RFID slave reader applies an initial phase shift to synchronize with the master (see Alg. 2). However, as frequency offset accumulates over time, the readers will notice that the phase of the master drifts.

PushID forms a closed loop to compensate for residual frequency offsets. Specifically, the slave RFID reader observes the change in phase from the master over a time interval $(t, t + \Delta t)$, where t is the most recent time of synchronization and Δt is the time elapsed since then. Should we observe a change in phase of the master transmitter since then, i.e. $phase(t + T) - phase(t)$, the slave applies the negative value of this phase offset to its own transmission. Note that this would, in effect, remove residual phase errors, allowing the two readers to combine their signals coherently at RFID tags. Further, note that due to channel reciprocity, the same phase shifts would ensure coherent combining on both the downlink and uplink.

Algorithm 2 Initial Frequency Synchronization

- 1: f_M : master’s oscillator frequency, f_{S_i} : slave i oscillator frequency.
 - 2: $\mathbf{H}(\mathbf{t}) = (h_1 e^{j2\pi(f_M - f_{S_1})t}, h_2 e^{j2\pi(f_M - f_{S_2})t})^T$
 - 3: Decompose $\mathbf{H}(\mathbf{t})$ we have $\mathbf{H}(\mathbf{t}) = \mathbf{R}(\mathbf{t})\mathbf{H}\mathbf{T}(\mathbf{t}) = \begin{pmatrix} e^{-j2\pi f_{S_1}t} & 0 \\ 0 & e^{-j2\pi f_{S_2}t} \end{pmatrix} \begin{pmatrix} h_1 \\ h_2 \end{pmatrix} e^{j2\pi f_M t}$
 - 4: $\mathbf{H}(\mathbf{t}) = e^{j2\pi f_M t} \mathbf{R}(\mathbf{t})\mathbf{H}\mathbf{T}(\mathbf{t}) e^{-j2\pi f_M t} = \begin{pmatrix} e^{j2\pi(f_M - f_{S_1})t} & 0 \\ 0 & e^{j2\pi(f_M - f_{S_2})t} \end{pmatrix} \mathbf{H} = \mathbf{A} \mathbf{H}$
 - 5: Apply frequency compensation $\mathbf{H}^{-1}\mathbf{A}^{-1}$ at slaves.
-

Scaling the system: A key challenge, however, is scaling the above system beyond two readers. This is because upon canceling one’s own transmission, each reader would perceive a linear combination of all other readers in its vicinity and not that of the master alone. This means that should a phase drift occur simultaneously for multiple readers (which is likely), these readers would be misled by out-of-sync transmissions from the others among them. To make matters worse, some readers may be beyond the communication range of the master and therefore may not be in a position to synchronize directly with the master as the last resort.

To mitigate this problem, we seek to assign to each RFID

reader a unique reference reader to which it may synchronize, should the master not be within its vicinity. Specifically, we assign indexes $1, \dots, n$ to each RFID reader, where RFID reader 1 denotes the master and all other readers denote slaves. We assume that these indexes are known a priori by the readers and can be constantly broadcasted so that any reader that fails can be removed from consideration in the optimization. Each RFID reader aims to synchronize its phase relative to the neighbor with the smallest index. In effect, our synchronization scheme creates a spanning tree of RFID readers (see Fig. 6), provided the graph of all readers is a connected one, which we assume. This spanning tree is designed to ensure that all readers eventually synchronizes to the root – the master RFID reader.

At this point each slave RFID reader subtracts its own signal and tracks the phase of the remaining linear combination. Should this change beyond a threshold, the RFID reader requests all its children and descendents in the spanning tree to cease transmission and then attempts to re-synchronize its transmission with the remaining active readers. We note that the master RFID reader never stops its transmission. It is easy to see that this scheme ensures that all RFID readers eventually transmit in-sync with the master.

5.2 Time Synchronization

PushID performs a two-step time synchronization process, a coarse synchronization to align symbols and a fine-synchronization that leverages the phase of signals from the master across frequency.

Coarse Synchronization: PushID synchronizes slaves with the master using the known preamble of the *Query* command transmitted by the reader to initialize an inventory. To maximize time resolution, slaves receive this signal from the master at a high sampling rate. Slaves then apply correlation with the known preamble to obtain the index of the master’s signal. We then only consider correlation coefficients above a threshold to reject outliers (we reject the bottom 6%). We repeat this process over five preambles and choose the result with the maximum correlation coefficient.

Fine Synchronization: To compensate for drift in timing offsets, PushID exploits the phase of signals from the master RFID reader a function of frequency. Specifically, recall that RFID transmissions hop between a wide range of frequencies in the 900 MHz ISM band spanning a total of 26 MHz. Each slave RFID reader estimates the phase of signals, having subtracted its own signal, across frequencies. PushID then monitors for any change in the slope of the phase of this signal across frequencies between measurements. Specifically, recall that any time offset of Δt between the two readers results in a frequency-dependent phase shift of $\Delta\phi = 2\pi f\Delta t$. As a result, PushID can estimate timing drifts by applying a least-squares linear regression [54] of ϕ as a function of f and




Figure 7: Depicts one round of PushID’s search for the unknown tags. Without a tag response, PushID starts another round by applying the next beamforming vector B_{i+1} .

obtain the resulting slope m . Any drift in timing offset can simply be computed as $m/2\pi$ and corrected for. We note that our system can scale akin to frequency offset compensation above, for more than two readers in the network. Specifically, when any RFID reader goes out-of-sync in time, it informs all its descendents to stop transmitting before attempting to re-synchronize.

6 Implementation and Evaluation

We implement PushID on a testbed of USRP N210 software radios with SBX/WBX daughterboards operating as RFID readers. We feed an omnidirectional and planar antenna to the antenna ports of each node for full-duplex use. All our readers are SISO, unless specified otherwise. Each USRP connects with an independent Jacksonlab Fury clock which could lead frequency and timing offset among the nodes. At the backend, each USRP is connected via Ethernet cables to a 64-bit Dell computer running Ubuntu 16.04. We also assume dedicated socket-based TCP connections between the reader nodes. Our RFID tags are commercial passive Alien Squiggle RFID tags. We measure a maximum range of up to 8.7 meters with our setup using one reader antenna.

PushID software: PushID is fully implemented in UHD/C++ including beamforming and distributed synchronization. In addition, we use an in-house UHD/Gnuradio based C++ RFID emulator to decode signals from the tags. We implement the set-cover based optimization in the cloud on a cluster of 64-bit core i7 Ubuntu machines and report the optimal beamforming weights to the reader nodes.

We ensure that all PushID RFID readers implement ASK modulation with PIE encoding to align with the specifications of the Gen2 RFID protocol. Apart from delivering energy, PushID readers also actively transmit messages which specify the tag’s modulation format, encoding scheme and backscatter frequency. The RFID tags in our experiments employ ASK modulation and FM0 encoding. The protocol flow of PushID is shown in Fig. 7.

Testbed: We evaluate PushID in two testbeds indoors and outdoors. (a) *Outdoor Testbed:* We deploy PushID around a football field (140×140 meters) with 8 transmitters. (b) *Indoor Testbed:* We deploy four-transmitter based PushID

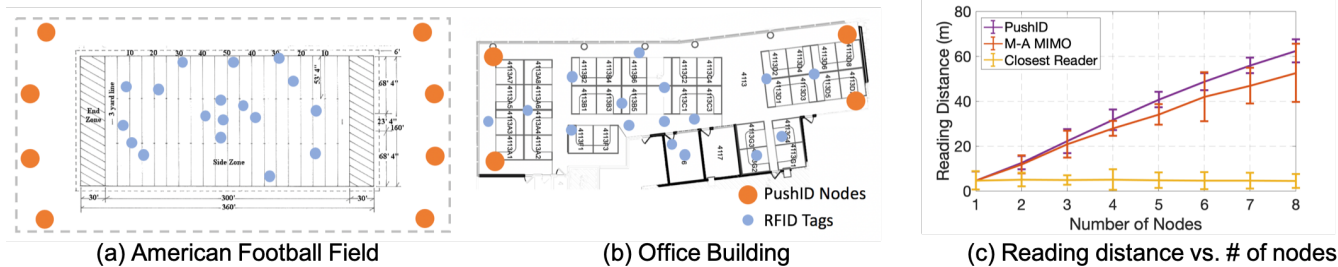


Figure 8: (a) **Outdoors**: We deploy PushID in a football field with one case of the transmitters placed as shown above. (b) **Indoors**: We deploy a four-transmitter based PushID on a floor (20 × 40 meters) of an office building covering multiple rooms and cubicles (c) Plots maximum reading range vs. # of readers. M-A MIMO – Multi-antenna MIMO.




Figure 9: Time and Frequency Synchronization Accuracy

across a floor (20 × 40 meters) of an office building covering multiple rooms and cubicles. We note both testbeds are multipath rich due to stands/partitions in the former and cubicles/furniture in the latter blocking the direct path of some readers. We mount readers in various positions including different elevations. We put RFID tags in various positions and orientations that face towards different angles. Fig. 8(a) and (b) shows the candidate locations of RFID tags (represented by blue dots) and readers (represented by orange dots). We note that unless specified otherwise, all results incorporate an equal amount of data points (over 1000 RFID tag locations considered among them) from both testbeds with our core results evaluating how system accuracy changes in line-of-sight vs. non-line-of-sight relative to all readers.

Baseline: We compare PushID against two competing schemes: (1) *Closest Reader*: We assume that each reader independently decodes signals and each tag receives energy from the closest reader; (2) *Multi-antenna MIMO*: We assume that all reader antennas are co-located and synchronized by an external clock [10]. We note that unless specified otherwise, error bars in graphs denote standard deviation.

7 Results

7.1 Synchronization Accuracy

We evaluate the accuracy of PushID in achieving accurate frequency and time synchronization between base stations.

Method: We consider a testbed of up to eight USRP N210-

based RFID readers, one designated as the master and use PushID to synchronize the slaves to the master at high accuracy. The RFID readers are placed in various arbitrarily chosen geometries and different relative distances between the slaves and master reader. We measure two quantities of interest: (1) The error in time synchronization; (2) The error in frequency synchronization of signals at a USRP N210 receiver that compares the phase of the signals received from the master and slave(s) post PushID’s synchronization. We perform this experiment in both our indoor and outdoor testbed in which about half of the slaves on average are in non-line-of-sight relative to the master and some slaves (>50% of nodes >50 m away from master) synchronize via multiple hops using PushID’s spanning tree approach.

Results: Fig. 9(a) shows the mean and standard deviation (error bars) in accuracy of time synchronization for different ranges of distance between a slave RFID reader and the master. We note that, as expected, the mean error with increasing distance also increases: PushID achieves a mean error of 0.12, 0.53, 2.69 μ s with the range of 10 to 30, 30 to 50 and 50 to 70 meters. However, we note that even the worst-case error is much smaller than one Nyquist time sample, given the narrow bandwidth of RFIDs (20 kHz). This means that PushID achieves the required level of time-synchronization accuracy to perform efficient distributed MIMO.

Next, Fig. 9(b) depicts the accuracy in frequency synchronization with increase in the number of RFID readers in the network. We find as expected the mean and variance of phase synchronization errors increase modestly as more readers join the network. PushID achieves a mean error of 0.0001 Hz in frequency offset overall across experiments. We note that this error corresponds to a phase shift of 0.0005° over the duration of a typical RFID packet and therefore minimally impacts the throughput of PushID’s distributed MIMO architecture, as observed in Sec. 7.3.

7.2 Range vs. Number of Nodes

In this experiment, we evaluate the maximum distance for a number of slaves that can detect the response from the tag.

Method: We deploy up to eight RFID readers in various geometries (starting from co-located and with progressively




Figure 10: Throughput vs. (a) Distance of closest reader; (b) number of readers; (c) CDF with changing tag orientation.

increasing spacing) with tags placed in both involving line-of-sight and non-line-of-sight relative to the reader. Note that neither the location or number of RFID tags placed in the environment are known a priori to the readers. We consider, in aggregate over 1000 RFID tag locations across experiments. Across experiments, we note the distance between the RFID tag and its closest reader. Our experiments consider distances of up to 80 meters. Our goal is to estimate the maximum distance at which an RFID tag can be detected at the readers.

Results: Fig. 8(c) shows the maximum range of RFID tags with respect to the number of reader nodes (up to eight). As expected, PushID amplifies the received signal power from RFID readers to the tags and vice-versa, the PushID’s range increases quasi-linearly as the number of RFID readers increases. We note that the rate of increase does dip (gradually) with increasing number of readers due to the increasing impact of time and frequency synchronization errors as reported in Sec. 7.1 above. We further notice a surprising decreasing trend in the standard deviation with maximum distance. We find that this stems from the robustness of PushID to orientation in the presence of multiple distributed readers all oriented in diverse spatial directions. Our results show that PushID achieves a maximum range to an RFID tag of 64 meters, a gain of about $7.4\times$ vs. commercial RFID and 20% over multi-antenna MIMO. We also note that the standard deviation of multi-antenna MIMO is large in various experimental settings. In contrast, PushID has better resilience and stability across experiments which gives more spatial diversity that benefits the poor polarization sensitivity of the RFID dipole antennas. We also notice that sometimes the multi-antenna MIMO has better performance than PushID when tag’s orientation favors the direction of collocated readers in the LOS setting.

7.3 Throughput vs. Distance and Scale

In this experiment, we evaluate the impact of PushID on the throughput as we vary the number of RFID readers and the distance between the tag and its closest reader.

Method: We measure the throughput by first measuring the SNR of each RFID tag measured from the eight RFID readers after coherently combining signals to and from the tag. We

then adapt the ESNR metric [18] to the RFID context to estimate the maximum data rate achievable for the received SNR. Note that once PushID’s algorithm is applied to detect a tag, we use channel reciprocity to maximize throughput to each detected tag in this experiment. We deploy PushID in both the outdoor (Fig.8(a)) and indoor scenario (Fig.8(b)) and consider tags in both line-of-sight and non-line-of-sight relative to the readers. Our RFID tags use FM0 modulation which allows for data rates over 45 kbps.

Throughput vs. Distance: Fig.10a shows the increase in throughput vs. distance in line-of-sight and non-line-of-sight settings and compares it against the baseline system that connects to the closest RFID reader. We observe that for the baseline, as expected, a reader has a maximum range of about 8.7 m across both line-of-sight and non-line-of-sight settings and we see that performance drops to zero throughput beyond this distance. PushID, with a 67.5 meter maximum range on average outperforms the baseline significantly in line-of-sight and 58.9 meters in non-line-of-sight. Further, as expected the throughput of PushID drops down as the distance increases due to lower signal-to-noise ratio. We note that quite significantly, PushID’s performance with eight transmitters increases the throughput of RFID tags $2.6\times$ when compared to the baseline closest-RFID reader scheme at its maximum range of about 8.7 m.

Throughput vs. Number of Reader Nodes: As expected, with the increasing number of reader nodes we observe a gradual (logarithmic) increase in network throughput of covered RFID tags on average for PushID. There is a similar, although much more modest increase with reader nodes for the baseline owing to an increase in coverage area. However, our system observes a net mean throughput gain of $2.6\times$ over a network of 8-nodes over the baseline closest-reader system.

7.4 Impact of Orientation

Method: In this experiment, we model the distribution of the throughput of an RFID tag progressively oriented along various directions in 100 locations with 8 readers and compare three schemes: (1) PushID; (2) A 8-antenna MIMO scheme; (3) The closest reader baseline.

Results: Fig. 10c plots the CDF of throughput across

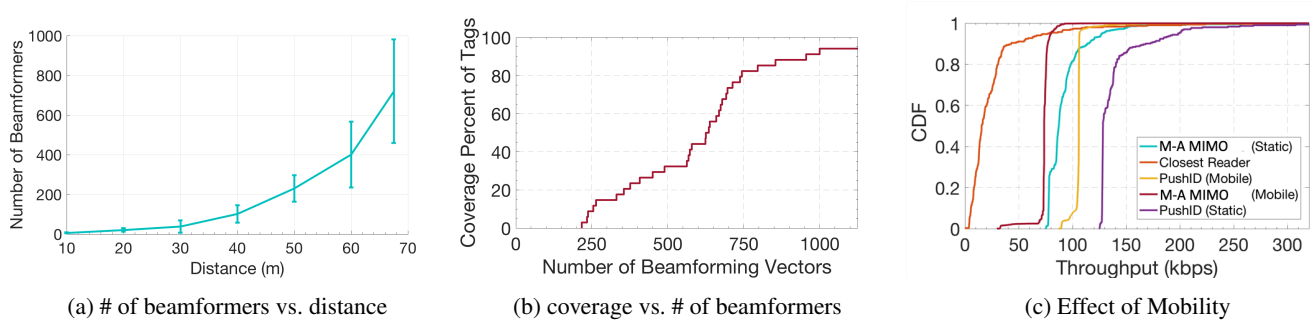


Figure 11: (a) Distance of closest reader; (b) % of tags covered; (c) CDF of throughput with mobility

schemes with changing tag orientation. We observe that PushID outperforms both multi-antenna MIMO (by $1.54\times$ median) and the closest reader baseline (by $7.4\times$ median). This is because, as explained in Sec. 4.1, PushID readers are oriented variously and therefore much more robust to change in orientations of the tags. In contrast, multi-antenna MIMO with reader antennas co-located loses performance when tags are oriented away from the MIMO reader, unlike PushID.

7.5 Impact of Mobility

Method: We deploy two tags in the environment in our indoor testbed, each placed initially at the same 100 randomly chosen initial locations at varying distances (up to 60 m) from the closest reader: (1) static RFID tags; and (2) an RFID tag moved around at walking speeds by volunteers. We measure the performance of PushID 8-antenna MIMO scheme and our closest reader baseline from eight RFID readers and compare performance.

Results: Fig. 11c plots the CDF of throughput across schemes for static and mobile tags. As expected, throughput dips in the presence of mobility across schemes. However, we note that PushID achieves gains over the baseline despite mobility ($1.6\times$ for static and $1.4\times$ for mobile over multi-antenna MIMO). PushID’s robustness to mobility stems from two reasons: (1) The mobility of the RFID tag ensures that the tag is highly likely to move out of coverage holes. As a result, during its trajectory, PushID has a higher likelihood of detecting the tag, compared to a static tag. This counteracts to help recover some of the loss in performance owing to changing multipath in PushID’s algorithms. (2) Once the tag is first detected, PushID can use channel reciprocity to rapidly continue beamforming to the tag and thereby respond to its mobility. We however highlight (as stated in limitations in Sec. 1) that mobility at very high speeds would significantly deteriorate PushID’s performance and gains, just as it would deteriorate commercial RFID systems.

7.6 Convergence and Coverage

Method: In this experiment, we measure the convergence time of PushID’s algorithm and how it is impacted by the distance of RFID tags and its trade-off with total area covered.

We note that the initial set of beamforming vectors PushID uses can be found offline and future updates take minimal time overhead, PushID’s main computational bottleneck is the rate at which beamforming weights can be applied by the USRP hardware which is 4.5 milliseconds in our testbed. We therefore measure convergence time in terms of the number of beamforming vectors that needs to be applied. Once again, we consider RFID tags at a wide range of distances to the closest reader with eight RFID readers and run PushID.

Results: Fig. 11b shows that the percentage of tags discovered increases with increasing number of beamformers for distances from the closest reader over 60 m. We emphasize here that it is rare for our system to detect all tags, since some tags remain virtually undetectable due to their location, orientation or shadowing – a natural limitation of our system (highlighted in Sec. 1) and indeed most wireless systems (for e.g. even the best cellular networks have deadspots [20]). Beyond 95% coverage, we see diminishing returns upon applying more beamforming vectors. Fig. 11a measures the trade-off between the distance of the tag and the number of beamformers needed to find $> 95\%$ of tags in the area. We find that in the worst case at maximum distance, we need 980 beamformers (4.4 seconds for a USRP N210).

8 Conclusion and Future Work

This paper presents, to our knowledge, the first distributed MIMO system to power commercial passive RFID tags. PushID develops a blind distributed beamforming algorithm to efficiently search through the space of beamforming vectors. It further develops a novel phase synchronization algorithm to synchronize distributed RFIDs. A detailed prototype evaluation on an eight-antenna distributed reader testbed reveals a $7.4\times$ improvement in range compared to state-of-the-art commercial readers. While this paper focuses on using existing commercial tags without modifications, we believe future work needs to explore algorithms that innovate on radio design and light-weight computation on the tags themselves to improve PushID’s performance.

Acknowledgements We thank anonymous NSDI reviewers and our shepherd, Fadel Adib, for their feedback and insights. We would like to thank NSF (grants 1718435, 1657318, 1823235 and 1837607) for their support.

References

- [1] Regulatory status for using rfid in the epc gen2 (860 to 960 mhz) band of the uhf spectrum. https://www.gs1.org/docs/epc/uhf_regulations.pdf, 2018. (Accessed on 08/27/2018).
- [2] AGEE, B. G., SCHELL, S. V., AND GARDNER, W. A. Spectral self-coherence restoral: a new approach to blind adaptive signal extraction using antenna arrays. *Proceedings of the IEEE* 78, 4 (April 1990), 753–767.
- [3] ALEVRAS, D., PADBERG, M., AND PADBERG, M. W. *Linear optimization and extensions: problems and solutions*. Springer Science & Business Media, 2001.
- [4] ATLASRFIDSTORE. Impinj rhcp far field rfid antenna (fcc/etsi). <https://www.atlasrfidstore.com/impinj-rhcp-far-field-rfid-antenna-fcc-etsi/>. (Accessed on 04/30/2018).
- [5] BAR-YEHUDA, R., AND EVEN, S. A linear-time approximation algorithm for the weighted vertex cover problem. *Journal of Algorithms* 2, 2 (1981), 198–203.
- [6] BIRARI, S. M., AND IYER, S. Mitigating the reader collision problem in rfid networks with mobile readers. In *Networks, 2005. Jointly held with the 2005 IEEE 7th Malaysia International Conference on Communication., 2005 13th IEEE International Conference on* (2005), vol. 1, IEEE, pp. 6–pp.
- [7] BISWAL, A. K., JENAMANI, M., AND KUMAR, S. K. Warehouse efficiency improvement using rfid in a humanitarian supply chain: Implications for indian food security system. *Transportation Research Part E: Logistics and Transportation Review* 109 (2018), 205–224.
- [8] BLETSAS, A., LIPPMAN, A., AND SAHALOS, J. N. Simple, zero-feedback, distributed beamforming with unsynchronized carriers. *IEEE journal on selected areas in communications* 28, 7 (2010).
- [9] CARDOSO, J.-F., AND SOULOUMIAC, A. Blind beamforming for non-gaussian signals. In *IEE proceedings F (radar and signal processing)* (1993), vol. 140, IET, pp. 362–370.
- [10] CHEN, S., ZHONG, S., YANG, S., AND WANG, X. A multiantenna rfid reader with blind adaptive beamforming. *IEEE Internet of Things Journal* 3, 6 (2016), 986–996.
- [11] CHEN, Y.-C., CHU, C., CHEN, R.-S., SUN, H. M., AND JU, P. Rfid-based bonded warehouse for science park. *International Journal of Radio Frequency Identification Technology and Applications* 5, 1 (2018), 1–23.
- [12] CROWDER, S. V. A simple method for studying run-length distributions of exponentially weighted moving average charts. *Technometrics* 29, 4 (1987), 401–407.
- [13] DONGARE, A., NARAYANAN, R., GADRE, A., LUONG, A., BALANUTA, A., KUMAR, S., IANNUCCI, B., AND ROWE, A. Charm: exploiting geographical diversity through coherent combining in low-power wide-area networks. In *Proceedings of the 17th ACM/IEEE International Conference on Information Processing in Sensor Networks* (2018), IEEE Press, pp. 60–71.
- [14] ELETREBY, R., ZHANG, D., KUMAR, S., AND YAĞAN, O. Empowering low-power wide area networks in urban settings. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication* (2017), ACM, pp. 309–321.
- [15] FAN, X., DING, H., LI, S., SANZARI, M., ZHANG, Y., TRAPPE, W., HAN, Z., AND HOWARD, R. E. Energyball: Wireless power transfer for batteryless internet of things through distributed beamforming. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 2 (July 2018), 65:1–65:22.
- [16] FAN, X., ZHANG, Z., TRAPPE, W., ZHANG, Y., HOWARD, R., AND HAN, Z. Secret-focus: A practical physical layer secret communication system by perturbing focused phases in distributed beamforming. In *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications* (April 2018), pp. 1781–1789.
- [17] FEIGE, U. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM (JACM)* 45, 4 (1998), 634–652.
- [18] HALPERIN, D., HU, W., SHETH, A., AND WETHERALL, D. Predictable 802.11 packet delivery from wireless channel measurements. *ACM SIGCOMM Computer Communication Review* 41, 4 (2011), 159–170.
- [19] HAMED, E., RAHUL, H., ABDELGHANY, M. A., AND KATABI, D. Real-time distributed mimo systems. In *Proceedings of the 2016 ACM SIGCOMM Conference* (2016), ACM, pp. 412–425.
- [20] HARROLD, T., AND NIX, A. Intelligent relaying for future personal communication systems.
- [21] HSIA, K.-H., WU, M.-G., LIN, J.-N., ZHONG, H.-J., AND ZHUANG, Z.-Y. Development of auto-stacking warehouse truck. *JRNAL* 4, 4 (2018), 334–337.
- [22] JADIDIAN, J., AND KATABI, D. Magnetic mimo: How to charge your phone in your pocket. In *Proceedings of the 20th annual international conference on Mobile computing and networking* (2014), ACM, pp. 495–506.

- [23] JIANG, C., HE, Y., ZHENG, X., AND LIU, Y. Orientation-aware rfid tracking with centimeter-level accuracy. In *Proceedings of the 17th ACM/IEEE International Conference on Information Processing in Sensor Networks* (2018), IEEE Press, pp. 290–301.
- [24] JIN, H., WANG, J., YANG, Z., KUMAR, S., AND HONG, J. Rf-wear: Towards wearable everyday skeleton tracking using passive rfids. In *Proceedings of the 2018 ACM International Joint Conference and 2018 International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers* (New York, NY, USA, 2018), UbiComp '18, ACM, pp. 369–372.
- [25] JIN, H., WANG, J., YANG, Z., KUMAR, S., AND HONG, J. Wish: Towards a wireless shape-aware world using passive rfids. In *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services* (New York, NY, USA, 2018), MobiSys '18, ACM, pp. 428–441.
- [26] KARP, R. M. Reducibility among combinatorial problems. In *Complexity of computer computations*. Springer, 1972, pp. 85–103.
- [27] KIM, D., AND YEO, J. Dual-band long-range passive rfid tag antenna using an amc ground plane. *IEEE Transactions on Antennas and Propagation* 60, 6 (2012), 2620–2626.
- [28] KUESTER, D., AND POPOVIC, Z. How good is your tag?: Rfid backscatter metrics and measurements. *IEEE Microwave Magazine* 14, 5 (2013), 47–55.
- [29] LEE, J., KWON, H., AND LEE, B. Design consideration of uhf rfid tag for increased reading range. In *2006 IEEE MTT-S International Microwave Symposium Digest* (June 2006), pp. 1588–1591.
- [30] LIUKKONEN, M., AND TSAI, T.-N. Toward decentralized intelligence in manufacturing: recent trends in automatic identification of things. *The International Journal of Advanced Manufacturing Technology* 87, 9–12 (2016), 2509–2531.
- [31] LOO, C. H., ELSHERBENI, A. Z., YANG, F., AND KAJFEZ, D. Experimental and simulation investigation of rfid blind spots. *Journal of Electromagnetic Waves and Applications* 23, 5–6 (2009), 747–760.
- [32] MA, Y., HUI, X., AND KAN, E. C. 3d real-time indoor localization via broadband nonlinear backscatter in passive devices with centimeter precision. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking* (2016), ACM, pp. 216–229.
- [33] MA, Y., LUO, Z., STEIGER, C., TRAVERSO, G., AND ADIB, F. Enabling deep-tissue networking for miniature medical devices. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication* (2018), ACM, pp. 417–431.
- [34] MA, Y., SELBY, N., AND ADIB, F. Drone relays for battery-free networks. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication* (2017), ACM, pp. 335–347.
- [35] MA, Y., SELBY, N., AND ADIB, F. Minding the billions: Ultra-wideband localization for deployed rfid tags. In *ACM MobiCom* (2017).
- [36] NADEEM, Q.-U.-A., KAMMOUN, A., AND ALOUINI, M.-S. Elevation beamforming with full dimension mimo architectures in 5g systems: A tutorial. *arXiv preprint arXiv:1805.00225* (2018).
- [37] PILLAI, V., HEINRICH, H., DIESKA, D., NIKITIN, P. V., MARTINEZ, R., AND RAO, K. S. An ultra-low-power long range battery/passive rfid tag for uhf and microwave bands with a current consumption of 700 na at 1.5 v. *IEEE Transactions on Circuits and Systems I: Regular Papers* 54, 7 (2007), 1500–1512.
- [38] RAGHAVAN, P. *Randomized Algorithms*.
- [39] RAHUL, H., HASSANIEH, H., AND KATABI, D. Sourcesync: a distributed wireless architecture for exploiting sender diversity. *ACM SIGCOMM Computer Communication Review* 41, 4 (2011), 171–182.
- [40] RAHUL, H. S., KUMAR, S., AND KATABI, D. Jmb: Scaling wireless capacity with user demands. In *Proceedings of the ACM SIGCOMM 2012 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication* (New York, NY, USA, 2012), SIGCOMM '12, ACM, pp. 235–246.
- [41] SARANGAN, V., DEVARAPALLI, M. R., AND RADHAKRISHNAN, S. A framework for fast rfid tag reading in static and mobile environments. *Computer Networks* 52, 5 (2008), 1058–1073.
- [42] SHEPARD, C., YU, H., ANAND, N., LI, E., MARZETTA, T., YANG, R., AND ZHONG, L. Argos: Practical many-antenna base stations. In *Proceedings of the 18th annual international conference on Mobile computing and networking* (2012), ACM, pp. 53–64.
- [43] SHI, L., KABELAC, Z., KATABI, D., AND PERREAULT, D. Wireless power hotspot that charges all of your devices. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking* (2015), ACM, pp. 2–13.

- [44] SKLIVANITIS, G., ALEXANDRIS, K., AND BLETSAS, A. Testbed for non-coherent zero-feedback distributed beamforming. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on* (2013), IEEE, pp. 2563–2567.
- [45] SLOCK, D. T. M. Blind fractionally-spaced equalization, perfect-reconstruction filter banks and multichannel linear prediction. In *Proceedings of ICASSP '94. IEEE International Conference on Acoustics, Speech and Signal Processing* (April 1994), vol. iv, pp. IV/585–IV/588 vol.4.
- [46] SMITH, J. R., SAMPLE, A. P., POWLEDGE, P. S., ROY, S., AND MAMISHEV, A. A wirelessly-powered platform for sensing and computation. In *International Conference on Ubiquitous Computing* (2006), Springer, pp. 495–506.
- [47] SORNIN, N., LUIS, M., EIRICH, T., KRAMP, T., AND HERSENT, O. Lorawan specification. *LoRa alliance* (2015).
- [48] TALLA, V., HESSAR, M., KELLOGG, B., NAJAFI, A., SMITH, J. R., AND GOLLAKOTA, S. Lora backscatter: Enabling the vision of ubiquitous connectivity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 3 (2017), 105.
- [49] TRAN, N., LEE, B., AND LEE, J.-W. Development of long-range uhf-band rfid tag chip using schottky diodes in standard cmos technology. In *Radio Frequency Integrated Circuits (RFIC) Symposium, 2007 IEEE* (2007), IEEE, pp. 281–284.
- [50] TSE, D. Fundamentals of wireless communication.
- [51] VARSHNEY, A., HARMS, O., PÉREZ-PENICHER, C., ROHNER, C., HERMANS, F., AND VOIGT, T. Lorea: A backscatter architecture that achieves a long communication range. In *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems* (2017), ACM, p. 18.
- [52] WANG, J., AND KATABI, D. Dude, where’s my card?: Rfid positioning that works with multipath and non-line of sight. In *ACM SIGCOMM Computer Communication Review* (2013), vol. 43, ACM, pp. 51–62.
- [53] WANG, J., VASISHT, D., AND KATABI, D. Rf-idraw: virtual touch screen in the air using rf signals. In *ACM SIGCOMM Computer Communication Review* (2014), vol. 44, ACM, pp. 235–246.
- [54] WOLD, S., RUHE, A., WOLD, H., AND DUNN, III, W. The collinearity problem in linear regression. the partial least squares (pls) approach to generalized inverses. *SIAM Journal on Scientific and Statistical Computing* 5, 3 (1984), 735–743.
- [55] YANG, J., AND LEUNG, J. Y.-T. A generalization of the weighted set covering problem. *Naval Research Logistics (NRL)* 52, 2 (2005), 142–149.
- [56] YANG, L., CHEN, Y., LI, X.-Y., XIAO, C., LI, M., AND LIU, Y. Tagoram: Real-time tracking of mobile rfid tags to high precision using cots devices. In *Proceedings of the 20th annual international conference on Mobile computing and networking* (2014), ACM, pp. 237–248.
- [57] ZAK, S. H., UPATISING, V., AND HUI, S. Solving linear programming problems with neural networks: a comparative study. *IEEE Transactions on Neural Networks* 6, 1 (1995), 94–104.
- [58] ZEILER, M. D. Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701* (2012).
- [59] ZUNIGA, J. C., AND PONSARD, B. Sigfox system description. *LPWAN@ IETF97, Nov. 14th* (2016).