

Design and Implementation of an RFID-Based Customer Shopping Behavior Mining System

Zimu Zhou, *Student Member, IEEE*, Longfei Shangguan, *Student Member, IEEE, ACM*, Xiaolong Zheng, *Member, IEEE, ACM*, Lei Yang, *Member, IEEE*, and Yunhao Liu, *Fellow, IEEE*

Abstract—Shopping behavior data is of great importance in understanding the effectiveness of marketing and merchandising campaigns. Online clothing stores are capable of capturing customer shopping behavior by analyzing the click streams and customer shopping carts. Retailers with physical clothing stores, however, still lack effective methods to comprehensively identify shopping behaviors. In this paper, we show that backscatter signals of passive RFID tags can be exploited to detect and record how customers browse stores, which garments they pay attention to, and which garments they usually pair up. The intuition is that the phase readings of tags attached to items will demonstrate distinct yet stable patterns in a time-series when customers look at, pick out, or turn over desired items. We design ShopMiner, a framework that harnesses these unique spatial-temporal correlations of time-series phase readings to detect comprehensive shopping behaviors. We have implemented a prototype of ShopMiner with a COTS RFID reader and four antennas, and tested its effectiveness in two typical indoor environments. Empirical studies from two-week shopping-like data show that ShopMiner is able to identify customer shopping behaviors with high accuracy and low overhead, and is robust to interference.

Index Terms—Shopping behavior, RFID, backscatter communication.

I. INTRODUCTION

SHOPPING behavior analysis is of great importance in understanding the effectiveness of marketing and merchandising campaigns [2]. Deep shopping behavior data can help retailers capture customers' preferences, test new arrivals, and adjust marketing strategies. Mining customer shopping behavior in online stores is achievable by analyzing click streams and shopping carts [3], [4]. However, physical store retailers lack effective methods to identify customer behaviors. The only available information is the sales history, which fails to reflect customer behaviors before they check out, *e.g.* how customers browse the store, which products they show an

interest in, and what products they match up. Therefore, it is essential to explore new ways of capturing customer behaviors in physical stores.

Previous efforts have exploited cameras to monitor customer shopping behaviors in clothing stores [5], [6]. However, such methods involve sophisticated computer vision techniques to recognize and analyze arm motions. Alternative methods track customer routes in stores to mine hot zones and popular products [7], [8]. For example, the more customers traverse a route, the higher attention the items along this route gain. However, these approaches still fail to provide high-fidelity shopping behaviour information such as product browsing, pick-up actions and trying on clothes.

RFIDs are emerging as an essential component of Cyber Physical Systems. Many well-known garment manufacturers (*e.g.*, Abercrombie & Fitch, Calvin Klein, Decathlon) adopt passive RFIDs for sales tracking and anti-counterfeiting [9]. We envision the adoption of RFIDs will sweep the clothing market in the near future, and in this paper, we explore the feasibility of mining customer behavior in physical clothing stores with RFID devices. Through analyzing the shopping processes in clothing stores, we abstract three behavior mining functionalities essential to retailers: discovering popular categories, identifying hot items and excavating correlated items.

- **Popular category** represents the clothes that customers frequently *stop by*. Popular category indicates the items that attracts customers at first glance, and provides information *e.g.* trends in consumer preferences, to attract more customers into the shops.
- **Hot items** are the clothes frequently *picked out* or *turned around* by customers. They indicate customers' deeper interests after the first glance. This information can help retailers develop strategies to transfer the hot items into the final purchase.
- **Correlated items** are the clothes that are frequently paired with or tried on together, which facilitates retailers to infer customer shopping habits and adopt bundle-selling strategies to boost profits.

These shopping data reflect *which items the customers browse through, they show an interest in, and they match up*. Through jointly analyzing these three kinds of shopping data with a sales history, retailers can acquire a much deeper business value. For example, items which are seldom tried-on may indicate the designs of these items are not run of the mill, while a hot item with unsatisfactory sales volume may suggest an unacceptable price, which indicates the need for a sales promotion or discount.

Manuscript received August 8, 2016; revised December 22, 2016; accepted March 20, 2017; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor S. Jin. Date of publication April 12, 2017; date of current version August 16, 2017. This work was supported by the National Natural Science Foundation of China through the General Program under Grant 61672240. This paper was presented at ACM SenSys 2015 [1]. (*Corresponding author: Xiaolong Zheng.*) (Zimu Zhou and Longfei Shangguan are co-first authors.)

Z. Zhou is with the Computer Engineering and Networks Laboratory, ETH Zurich, 8092 Zürich, Switzerland (e-mail: zimu.zhou@tik.ee.ethz.ch).

L. Shangguan is with the Department of Computer Science, Princeton University, Princeton, NJ 08544 USA (e-mail: longfeis@cs.princeton.edu).

X. Zheng and Y. Liu are with the School of Software and TNLIS, Tsinghua University, Beijing 100084, China (e-mail: xiaolong@greenorbs.com; yunhao@greenorbs.com).

L. Yang is with the Department of Computing, The Hong Kong Polytechnic University, Hong Kong (e-mail: young@tagsys.org).

Digital Object Identifier 10.1109/TNET.2017.2689063

In this paper, we present ShopMiner, a customer shopping behavior sensing system using commercial off-the-shelf (COTS) RFID devices. The principle is that the phase readings of RFID tags attached to clothes items demonstrate distinct yet stable patterns when customers stop beside, pick out, turn around, or pair up items. Specifically, customers are likely to stand still for a while in front of attractive items, and hence block the wireless links between reader antennas and the items. Thus the phase reading of the popular items show a distinct pattern from the unpopular ones (*i.e.*, not viewed by customers). Similarly, the phase reading of hot tags will change dramatically when customers pick them out or turn them around to inspect the front design. The correlated items are brought together by one customer, thus experiencing the same moving route and showing similar phase changes.

ShopMiner's design harnesses these spatial-temporal phase reading correlations. Our key techniques include a foreground/background segmentation scheme for popular category detection, a statistical model for hot item identification, and a clustering algorithm for correlated items excavation. We implement ShopMiner on COTS RFID devices including an ImpinJ R420 reader, four Yeon antennas model YAP-100CP and multiple Alien UHF passive tags. Experimental results show that ShopMiner can detect popular items with a True Positive Rate (TPR) of 92%, identify hot items with an accuracy of 94% and 87% for pick-out and turn-around, and achieve over 85% accuracy for correlate item excavation in multi-user case.

Contributions

(1) ShopMiner is a unified sensing framework. Despite recent works on RFID-based shopping behavior sensing [10], [11], none has incorporated the three key factors that are essential to retailers, *i.e.*, which items the customers browse through, they show a interest in, and they pair up. (2) As a long-term running system, ShopMiner optimizes computation and storage overhead. We design a hierarchical architecture and a set of algorithms for multistage behavior detection. (3) We implement ShopMiner on COTS RFID devices, and conduct comprehensive experiments in two shopping-like scenarios. Empirical studies show that ShopMiner achieves over 90% TPR for customer behavior detection.

Roadmap

In the rest of paper, we present the scope, design, implementation and evaluation of ShopMiner in Section II, Section III, Section IV and Section V. Section VI reviews related work and Section VII concludes this paper.

II. SCOPE

We envision ShopMiner can be deployed in clothing stores to monitor customer behaviors without body instruments. Fig. 1 shows the typical shopping process in clothing stores before a customer checks out. It contains the following steps: browsing or standing still in front of attractive items; examining interesting items by picking them out or turning




Fig. 1. (a): browsing and then stopping in front of the item of interest; (b), (c): inspecting the item more closely; (d): matching up and trying items on in a fitting room.

them around; taking the desired items and trying them on in a fitting room. Through identifying and counting items that are most viewed, picked out and turned around by customers, as well as matched up items, retailers can discover popular categories, hot items, and correlated pairs for better trading strategies and tie-in promotions.

III. DESIGN

This section presents the design of ShopMiner. We assume each garment is attached to an RFID tag. Note that the design of ShopMiner is catered for clothing stores where garments are hung on racks. While some techniques are dedicated to clothing stores, for instance, to differentiate pick-out and turn-around actions, others can be applied in other shops such as bookstores (Section V-F).

A. Discovering Popular Category

Popular items are the garments that customers frequently stop beside and look at. Such information indicates customers' first impression of products.

1) *Exploiting Body Blocking Effect*: We exploit the blocking effect of a human on multiple tag-to-antenna links to detect customer stopping beside and to infer popular categories. As Fig. 2(a) shows, when a customer (User 1) stands still in front of one garment, his/her body tend to block the Line-Of-Sight (LOS) link between the reader antenna and the item. On average, the link will experience high shadowing losses.

As an illustration, we deploy an ImpinJ R420I RFID reader with three directional antennas and 16 RFID tags (T_1 to T_{16}) as Fig. 2(b). The antennas are placed 3m away from the tags (tag i is attached to garment i) and 0.8m above the floor. One volunteer was asked to walk along route 01, stand still for 8s in front of garment 04, and walk away. Fig. 3 plots the phase measurements of tag 04 to tag 07. During the first 4s, the phase readings of all the four tags stay at a different yet stable value. From 4s to 12s, the phase of tag 04 changes to another stable level, manifesting an obstacle between tag 04 and the antenna. During the last 6s, the phases of tag 05, 06 and 07 change sequentially, indicating that the LOS paths of these tags are blocked sequentially. This experiment shows the possibility of using the body shadowing effect to detect customer presence.

2) *Modeling Multipath Effect*: The multipath effect is common indoors and can affect phase readings. As shown in Fig. 2(a), a customer near the rack can create a new path.

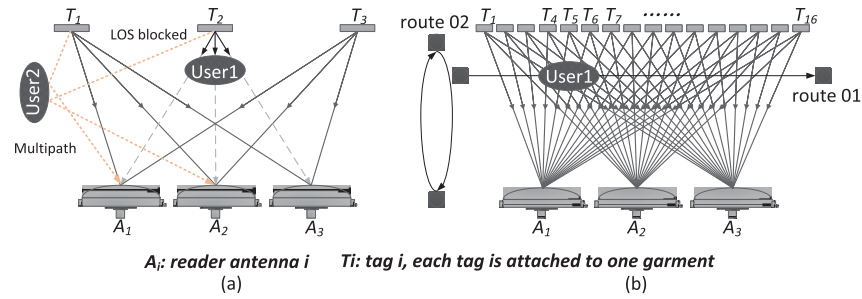


Fig. 2. Illustrative settings of popular category discovery. (a): LOS path blocked/multipath. (b): moving routes of a user.




Fig. 3. Phase trends of four tags.




Fig. 4. Phase readings affected by multipath effect.

Such changes of multipath components can affect the measured phase of RFID tags.

To model the impact of multipath on the phase readings, three volunteers walk back and forth along route 01 and 02 for 10min and we collect over 20,000 phase readings from 48 links. We randomly pick five tags and plot their phases in Fig. 4. As shown, the multipath-induced phases roughly follow a Gaussian distribution.

In summary, both body blocking and the multipath effect can introduce phase dynamics. The phase dynamics exerted by multipath is small and follows a Gaussian distribution, whereas the phase changes dramatically when the LOS path is blocked.

3) *Popular Category Detection Scheme*: To detect the items in a popular category, we apply a Gaussian model based change point detection scheme from background detection in image processing [12]. Specifically, we first detect dramatic changes in the phase streams of each link, and combines the results from multiple links of the same tag to infer whether the item that the tag attached to is in a popular category.

Gaussian model based change point detection: We first consider the phase stream of the i^{th} link of a tag (t). The phase stream is split into successive phase segments of length d using a window of size $|w|$. We denote $r_{i,j}$ as the phase reading of the i^{th} link of the tag collected within the j^{th} window (w_j). Since each tag can be interrogated multiple times in one

window, we use the average phase reading within one window as the phase of this tag in this window. We set $|w| = 0.02s$ and $d = 50$, respectively, which empirically balances computational efficiency and detection granularity. As shown in Fig. 4, the phase of each link shows a distinct Gaussian distribution, we thus create one Gaussian model for each link. ShopMiner then examines each phase reading $r_{i,j}$ in a phase segment by comparing it against the corresponding distribution. Given a phase reading $r_{i,j}$, and the Gaussian model $N_i(\mu_i, \sigma_i^2)$ of link i , we formulate the following hypothesis test with H_0 (change) and H_1 (stable):

$$\begin{cases} H_0 : r_{i,j} \notin (\mu_i \pm \frac{\sigma_i}{\sqrt{k_i}} \cdot z_{\alpha/2}) \\ H_1 : r_{i,j} \in (\mu_i \pm \frac{\sigma_i}{\sqrt{k_i}} \cdot z_{\alpha/2}) \end{cases} \quad (1)$$

where $0 < \alpha < 1$, and k_i is the sample size. $(\mu_i \pm \frac{\sigma_i}{\sqrt{k_i}} \cdot z_{\alpha/2})$ stands for the confidence region with the confidence level $(1 - \alpha)$. For example, H_1 under $\alpha = 0.05$ indicates that the phase reading $r(i, j)$ exhibits notable change (and thus the corresponding link is blocked) with a probability of 95%.

Since there may be multiple links between a tag t and the antennas as in Fig. 2, we further conduct a majority voting to decide whether a tag t is blocked by a customer. In case of a tie, we count it as ‘blocked’ to avoid missing potential items in popular categories.

Detecting popular category: If a tag t is determined as ‘blocked’ via the above change point detection in a time window w_j , it does not necessarily mean the tag is in a popular category. Only tags that are blocked after consecutive time windows are decided as in a popular category, because they attract the customers to stand in front of them for a reasonably long time. Thus we formulate the following hypothesis test with H_0 (in popular category) and H_1 (not in popular category):

$$\begin{cases} H_0 : s_t \geq \theta \\ H_1 : s_t < \theta \end{cases} \quad (2)$$

where s_t is the number of consecutive phase segments that tag t is blocked and θ is a pre-defined threshold. Note that θ is mainly determined by body blockage of the LOS path. Thus it is affected by the length of links and is robust to factors such as shop size, layout and reader placements. Within the range of RFID readers (about 6m), a fixed threshold can scale to other shops without per-shop calibration.

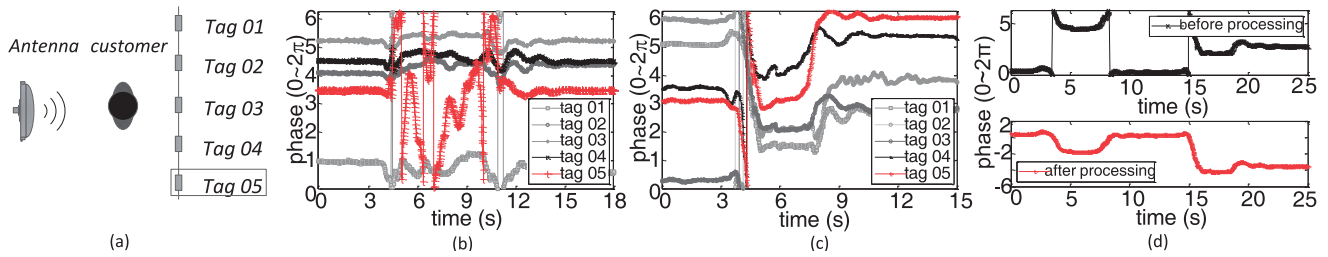


Fig. 5. Illustration of hot item identification: (a) experimental settings; (b) phase patterns of pick-out; (c) phase patterns of turn-around; (d) phase trend before/after de-periodicity.

Model training and updating: To precisely detect the change point, an accurate Gaussian model is required for each tag. Initially, the parameters of m Gaussian models are computed when there are no customers in the shop (*e.g.*, before the opening). Afterwards the phases of the newly detected change points are input into the model to update the model parameters to adapt to environmental changes.

B. Identifying Hot Items

ShopMiner identifies hot items by detecting and counting the following customer actions:

- **Turn-around:** A customer observes an item by turning it around from the side-view to the front-view.
- **Pick-out:** A customer takes a close look at an item by picking it out from the clothing rack.

The two actions indicate different levels of interest in items. Customers often turn around an attractive item for its initial appearance, and pick it out for a closer and more detailed inspection when they show greater interest. Hence ShopMiner detects and identifies these two actions separately.

1) Detecting Pick-Out and Turn-Around Actions: Both the turn-around and the pick-out actions will induce motion to tags, which will exhibit notable phase changes. As shown in Fig. 5(a), we place five garments on a rack and ask a volunteer to (1) pick up item 05, closely inspect it, and then put it back; (2) turn over item 05 to see its front, hold it for a while and then put it back. The phase readings of the five tags are shown in Fig. 5(b)-(c).

As Fig. 5(b) shows, initially the phase readings of these five tags all remain stable. When the volunteer *picks out* item 05 at 4s, its phase jumps abruptly until the volunteer puts it back at 12s. Similarly, as shown in Fig. 5(c), when the volunteer *turns around* item 05 at 4s, its phase also changes significantly. The phase then keeps stable throughout [4s, 8s], indicating that the item is held steadily by the volunteer. As the volunteer puts the item back, another fierce phase change occurs. The measurement indicates that it is possible to detect pick-out and turn-around actions by observing the phase changes of tags. That is, identifying whether a period of phase changing occurs. Note that the phase readings also experience notable changes when a person blocks the links, which we exploit for popular category discovery. However, as shown in Fig. 3, the phase changes due to blocking effect are usually within 2π , while those induced by pick-out or turn-around

often exceed 2π (Fig. 5). A partial explanation is that the changes of the LOS path (changes of distance and orientation due to pick-out/turn-around) may incur significantly more notable phase variations than the blocking of the LOS path (the LOS path remains the same but with different attenuation due to body blocking). Therefore, it is possible to distinguish the phase trends of popular category and those of hot items.

2) Differentiating Pick-Out and Turn-Around Actions: In clothing stores, it is common that clothes are hung compactly on the rack, with their side-views facing the customers (Fig. 1). Consequently, when the customer picks out one garment and closely inspects it, the nearby items will be struck unintentionally, causing them to vibrate. Hence the phase trends of these items will exhibit a minor yet *different* variation tendency to the desired item (tag 01 - 04 in Fig. 5(b)). In contrast, when the customer turns one item (say item 05) to see its front, the surrounding items will be forced to turn as well. As a result, the phase trends of these items will show a *similar* tendency to tag 05 (tag 01 - 04 in Fig. 5(c)). Thus it is viable to distinguish pick-out and turn-around actions by jointly considering the phase readings of nearby tags, *i.e.*, comparing the similarity of their phases.

3) Hot Item Identification Scheme: We first design a segmentation-based pick-out/turn-around detection scheme and then present a peer-assisted identification scheme to differ these two actions.

Segmentation-based detection: ShopMiner performs segmentation on the phase readings to detect whether a pick-out/turn-around action occurs. Denote the phase readings as $S = (s_i) \in R^{1 \times N}$, where N is the number of discrete time points within a window. We first calculate the discrete probability distribution function (PDF) of phase values within each window. Given two consecutive windows w_i and w_j and their PDFs P and Q , we then compute the *KL-divergence* of the two PDFs:

$$D_{KL}(P||Q) = \sum_i P(i) \cdot \ln \frac{P(i)}{Q(i)} \quad (3)$$

The *KL-divergence* quantifies the similarity of phase trends within two consecutive windows. We denote the period when there is a pick-out/turn-around action as a *motion period* and the remaining as a *silent period*. Within a silent period, the phase value will stay at a stable level. Hence the *KL-divergence* of two consecutive windows within the silent period should be small. In contrast, if at least one window




Fig. 6. Pipeline of hot item identification: (a) raw phases; (b) segmentation; (c) de-periodicity; (d) normalization.

is within the motion period, the PDF of these two windows should be notably different, thus leading to a large KL-divergence. ShopMiner checks $D_{KL}(P||Q)$ to detect whether the current window is within the silent period. After finding all windows within the silent periods, we can extract the motion periods accordingly. We set the window size to 0.5s to cover the majority duration of one pick-out/turn-around action.

Improving detection accuracy: Since both pick-out and turn-around actions will introduce phase turbulence to nearby items, their phase readings will vary as well and cause false alarms in the segmentation-based scheme. In a measurement of 200 pick-out and turn-around actions, we observe a false alarm rate of 40%. To improve the detection accuracy, we further process the phases in the following two steps.

De-periodicity: The phase value reported by the reader API is a periodic function ranging from 0 to 2π . So when the phase decreases to 0, it will jump to 2π (Fig. 5(d)). We term this abrupt phase change as a phase hop. ShopMiner adopts the method in [11] to handle phase hops, which adds or subtracts 2π when a phase hop occurs. The phase trend before/after de-periodicity is shown in Fig. 5(d).

Variance comparison: Suppose m tags are detected in the segmentation step. For each tag i , we denote its phase readings within the motion period as $S_i = (s_j) \in R^{1 \times N_i}$. N_i is the length of the sample group, and may vary from tag to tag due to multipath [13] and random access principle of ALOHA protocol. We thus split each sample group into N frames. The frame length is set to 0.1s. Since multiple samples may locate within a frame, we compute their mean (denoted as s'_j), and use it as the phase value of this frame. Then we compute the variance of S_i as follows:

$$\text{Var}(S_i) = \frac{1}{N} \sum_{j=1}^N (s'_j - \mu)^2 \quad (4)$$

After computing the variance of each tag, the one with the highest variance is denoted as the desired tag. The rationale is as follows. The motion of nearby items is indirectly driven by human actions. The driven power will be absorbed by the clothes and diminishes rapidly over time. Hence the desired item will experience notably higher turbulence than the undesired ones, and show a larger variance.

Peer-assisted identification: To distinguish pick-out and turn-around actions, we jointly consider the desired tag with the tags nearby. The observation is that the phase readings of nearby tags demonstrate a similar variation for turn-around, yet behave differently for pick-out. Specifically, for each of the

m phase trends, we zoom out the local dissimilarity of phase samples by normalizing this phase trend (Fig. 6(d)). After the normalization, we splice these m phase trends consequently into a single phase trend, say $S = (s_j) \in R^{1 \times N}$. Then auto-correlation is performed on S :

$$\chi(m, \tau) = \frac{\sum_{k=0}^{k=\tau-1} [s_{m+k} - \mu(m, \tau)][s_{m+k+\tau} - \mu(m+\tau, \tau)]}{\tau \cdot \sigma(m, \tau) \cdot \sigma(m+\tau, \tau)} \quad (5)$$

where $\mu(k, \tau)$ and $\sigma(k, \tau)$ are the mean and standard deviation of the phase samples $\langle s_k, s_{k+1}, \dots, s_{k+\tau-1} \rangle$, respectively. In our case, τ equals the number of data samples within the motion period, and is known a prior. Generally, the phase trend S , if connected by k similar phase trends, should behave like a periodic signal, hence having a high auto-correlation. Hence we can ascertain whether the motion period is caused by turn-around, based on the auto-correlation coefficient:

- if $\chi(m, \tau) \geq \delta$, then *action* = turn-around;
- if $\chi(m, \tau) < \delta$, then *action* = pick-out;

We test various thresholds over 1000 measurements, and find a threshold $\delta = 0.65$ optimal for our case. This threshold is impacted by the layout of shops as well as link distance. While some efforts have explored modeling the impact of human presence on wireless links to ease parameter transferring to a different link [14], it remain open how to transfer parameters for diverse activities and in multipath-rich environments as in our scenario. Hence we recommend re-calibration of this threshold for different shops. Note that reordering the garments, which occurs frequently in clothing stores, does not affect the performance, since ShopMiner does not rely on the sequence or position of garments for pick-out/turn-around detection.

C. Excavating Correlated Items

Our correlation analysis aims to find the garments that are usually tried on together, *e.g.*, dress shirt and tie usually tie in together, while people buying suit pants often consider dress shoes. Previous efforts [10] proposed an RSS-based localization technique for correlated item discovery, based on the intuition that correlated items held by the same customer should be in close proximity. However, such a method is error-prone due to the following two reasons. (1) Items around the customer may also be in close proximity to the items in hand, and hence will be mistaken as correlated items. (2) Customers block/generate propagation paths dynamically in clothing stores, hence dampening the resolution of location-based schemes.

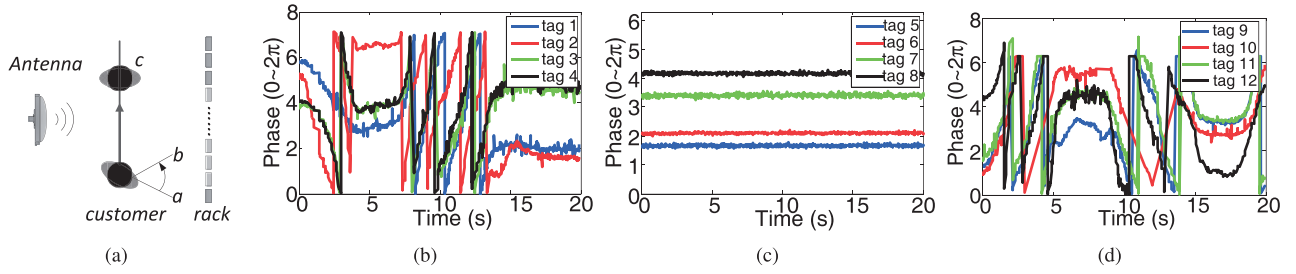


Fig. 7. Illustration of correlated item discovery: (a) routes of a customer; (b) phase trends of correlated items; (c) phase trends of stationary tags nearby; (d) phase trends of another group of correlated items.




Fig. 8. Phase clustering for correlated item discovery.

1) *Spatial-Temporal Correlation of Phase Features*: ShopMiner explores the spatial-temporal correlation of phase readings to discover correlated items. The observation is that correlated items, either in hand or in a shopping bag, follow a similar moving pattern with the customer, hence experiencing consistent temporal phase patterns. As an illustration, a volunteer is asked to carry four garments and walk to the fitting room along a route as Fig. 7(a). Fig. 7(b) plots the phase readings of the four tags, which exhibit similar temporal patterns. Specifically, when the volunteer walks from *a* to *b*, because the distance between the tag and antenna first increases and then decreases, the phase trend shows a symmetric profile. As the volunteer walks to *c*, the phase readings change continuously within $[0, 2\pi)$ then eventually remain stable when the customer reaches *c* and stays there for a while. On the other hand, comparing Fig. 7(b) with Fig. 7(c)-(d), the tags within different categories have diverse temporal phase profiles, which naturally set them apart. The measurement validates the feasibility of using phase trend similarity for correlated item discovery.

2) *Clustering Correlated Items*: Given a set of phase trends (x_1, x_2, \dots, x_n) , we aim to partition the n phase trends into m (m is unknown a priori and $m \leq n$) sets $S = \{S_1, S_2, \dots, S_m\}$, such that the within-cluster sum of squares is minimized:

$$\arg\min_S \sum_{k=1}^m \sum_{i,j \in S_k} T(x_i, x_j) \quad (6)$$

where $T(x_i, x_j)$ is the distance between x_i and x_j .

We design a heuristic algorithm that iteratively partitions the phase trends into different categories. Fig. 8 shows an example of the algorithm with ten tags to be classified. In the first iteration, the algorithm randomly picks one tag as the

pivot (tag 01 in the example), and computes the distance between its phase profile with those of the remaining tags. Tags whose phase profiles are sufficiently similar to that of the pivot are clustered. The algorithm then randomly picks another pivot and repeats this process on the remaining tags until the within-cluster sum of squares is minimized. Once the algorithm terminates, we get the corresponding tag set.

Distance metric $T(x_i, x_j)$: The desired distance metric needs to deal with two issues: (1) *phase trend inconsistency*, which means that phase trends may be of different lengths due to multipath effect. (2) *interrogation time inconsistency*, which indicates the phase values are sampled in different time slots due to the random access protocol. To address these issues, we use the Dynamic Time Warping (DTW) metric [15], which allows two time series that are similar but locally out of phase to align in a non-linear manner.

Handling phase scaling: Since phase is proportional to the tag-to-antenna distance, it may vary slightly from tag to tag within the same category due to different positions of tags. As shown in Fig. 7(b), *V-shapes* within these four tags are the scaled version of each other. This scaling problem may degrade the performance of DTW. Therefore we employ a variant of DTW, called Derivative Dynamic Time Warping (DDTW) [16], which exploits the same principle of DTW yet uses the derivative of phases as input. DDTW tolerates the phase difference in the Y-axis by inputting the derivatives of phases rather than the absolute values.

IV. IMPLEMENTATION

This section describes the practical issues in implementing ShopMiner.

A. Hardware

We prototype ShopMiner using COTS UHF RFID devices, including ImpinJ Speedway R420 RFID readers¹ and Yeon circularly polarized antennas. Each garment has an Alien passive RFID tag model AZ-9634 attached. The reader is connected to a local server via an Ethernet cable. To minimize the influence of network latency, we time-stamp each tag reading by the reader's local clock.

¹Our evaluations were conducted in China where, by default, RFID readers select one fixed channel to operate. To enable ShopMiner in other countries e.g. the US where readers are required to hop channels continuously according to FCC regulations, we recommend the phase calibration scheme in [17] to eliminate phase discontinuity.

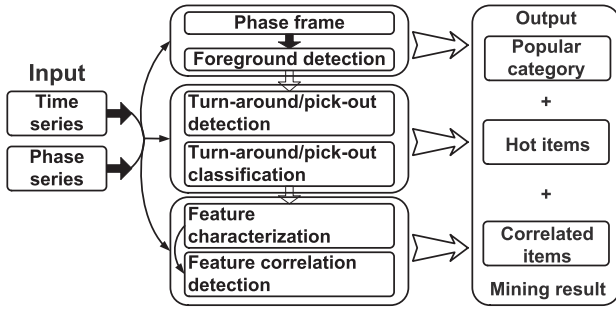


Fig. 9. Work flow of ShopMiner.

B. Software

We implement the software component of ShopMiner in Java. Fig. 9 shows the work flow. At the lowest level is the data collection module, which is integrated with the Octane SDK and continuously interrogates the nearby tags to read phase at a rate of 340 readings per second. The tag readings are grouped according to the tag ID (96bit or 128bit identifier) and stored in the local database. Initially, the data processing module fetches the phase readings, and feeds them into the popular category discovery module. After discovering popular categories, ShopMiner collects their tag IDs and performs pick-out/turn-around identification on the readings of these tags and any nearby tags to detect hot items. ShopMiner finally clusters hot items to excavate the correlated items. The software runs on a Lenovo PC with an Intel Core i7-4600U 2.1GHz CPU and 8GB RAM.

C. Modulation Scheme Selection

A typical EPC Gen 2 reader supports multiple pre-configured tag interrogation modes. Each mode differs in modulation, resulting in different reading rates and sensitivity to radio interference. For instance, the DenseReaderM4 mode has a low reading rate but high resistance to interference, while the MaxMiller mode yields a high reading rate yet only works in clear RF environments. In ShopMiner a high reading rate is required to track multiple items. However, the dynamic environments in clothing stores can induce severe interference to RFID interrogation, leading to noisy phase readings. Thus it is crucial to balance the reading rate and the resistance to interference. We empirically search for the balance as follows. ShopMiner round-robins all the selectable modulation schemes in order from the highest reading rate to the lowest. The first scheme with the standard phase variances $\text{Var}(\theta) \leq \delta$ is selected for tag interrogation. We conduct extensive experiments to test various thresholds, and set $\delta = 0.1$ for optimized performance.

D. Reader Deployment and Reader Collision

Mainstream commercial RFID readers can support four antennas, with each antenna effectively covering an area of $4m \times 4m$. Hence the retailer can monitor an area of nearly $60m^2$ with one reader, which costs less than 1000 USD. However, multiple readers are required for large stores or critical regions, *e.g.* racks for new products where more customers




Fig. 10. Illustration of collisions. Case 1: tag-to-tag collision; Case 2: reader-to-reader collision.

are expected. Since the coverage of multiple readers may overlap, collision is an important issue in a multi-tag-multi-reader system. As Fig. 10 shows, there are two types of collisions: The first is tag-to-tag collision, where multiple tags respond to the reader at the same time slot, making the reader unable to resolve any ID. The second is reader-to-reader collision. For a tag residing in the overlap between the interrogation zones of two adjacent readers, it may hear from both readers, and fail to resolve the command from either reader if the readers broadcast concurrently.

We adopt retransmission to resolve tag-to-tag collision, a common approach in slotted ALOHA protocol. To minimize the reader-to-reader collision, we propose a reader scheduling algorithm based on the Maximal Weighted Independent Set (MWIS). Given an undirected graph $G(V, E)$, an independent set of G is a subset $S \subseteq V$ such that no node pair (u, v) are neighbours in G , and every node $w \notin S$ has a neighbour in S . In the context of ShopMiner, we denote each RFID reader as a vertex in the graph. The weight on each edge indicates the number of tags that can be interrogated by both end point readers. ShopMiner employs the algorithm in [18] to find the MWIS, and schedules the reader to interrogate tags. The time complexity is asymptotically logarithmic in n (the number of nodes), which will not introduce significant computational overhead. To balance the workload of each reader, we sequentially wake up the readers in the independent set (algorithm's output) and the readers in the complement of this independent set. The initial weight of each edge is set based on the intersection of the tag set acquired by each reader. The weight of each edge is then updated every T minutes, where T changes with time. For instance, more customers browse stores when getting off work. Hence it is more likely to see frequent items change during these periods. Thus T should be small to guarantee the system correctness.

E. Filtering Interference from Employees

ShopMiner is designed to work in clothing stores in a self-service mode, where customers are expected to browse through the shops freely without the company of sales assistants. However, in many clothing stores, employees still frequently put clothes from the fitting rooms back to racks. Since ShopMiner cannot differentiate clothes held by employees from customers using RF signals, employees may impose interference to ShopMiner, especially for correlated items. We thus implement an interference filtering mechanism in ShopMiner to filter items held by employees as follows.




Fig. 11. Snapshot of data storage in ShopMiner.

We assume that each employee wears an RFID tag with a known ID to ShopMiner. ShopMiner then detects employees holding clothes items in the same principle as correlated item discovery in Section III-C. This is because the phase readings of both the employee tag and item tags still exhibit similar temporal patterns. Therefore, we adopt the same clustering method for correlated item detection, but further separate the clusters with tag IDs known to be attached to employees, and exclude these clusters for correlated item discovery.

While employees bringing items back to the racks can induce interference for correlated items discovery, employees bringing items from the warehouse may indicate customers' interest in the garments, as customers tend to ask employees for items of a different color or size that are not shown on the racks. To utilize this information, we need to check whether the employee has been to the warehouse. This can be achieved by deploying an additional reader within or at the door of the warehouse, which monitors the comings and goings of employee tags. Items co-located with the employee from the warehouse (detected by phase reading clustering) are identified as hot items.

F. Reducing Computation and Storage Overhead

Reduce Computation Overhead: Since hundreds of items may be picked out or turned around frequently in a clothing store, it is important to minimize the amount of items for correlated item discovery. As most of the items picked out or turned around may be hung directly back on the rack without being tried on, they do not belong to any category of correlated items and hence can be filtered out. In ShopMiner, only the tags satisfying the following two conditions will be considered for correlated item discovery: (1) The item has been picked out or turned around; (2) The item is sequentially identified by different RFID readers (indicating the item is in motion). Such a heuristic prunes a large portion of uncorrelated items and boosts the efficiency of correlation analysis.

Reduce Data Storage: In ShopMiner, RFID readers periodically interrogate tags to continuously monitor customer behaviours and the acquired data accumulate gradually. Fig. 11 shows a snapshot of the data acquired with time when an RFID reader monitors 5 items. The data volume increases linearly with time and rapidly accumulates to over 10Mb after 30s. That is, ShopMiner will generate 24Gb data monitoring only 100 items each hour, which is intolerable for practical deployment. To reduce the data storage overhead,




Fig. 12. Prototype deployment and testing environments.




Fig. 13. Illustration of testing scenarios with three item-to-customer distances ((a)–(c)) and three different numbers of customers ((d)–(f)).

ShopMiner runs the popular category discovery and hot item identification modules online. Only data of hot items will be stored in the database for off-line correlated item discovery. As Fig. 11 shows, after such data pruning, ShopMiner incurs significantly lower storage overhead, which accumulates to only 0.2Mb after 30s. That is, there will be only 480Mb data to 100 tags in one hour, which is negligible compared to the unfiltered 24Gb data. We admit that for large-scale shops, tag populations would be high and hence the data volume still accumulate rapidly. Thus we recommend performing off-line correlated item discovery each hour and discarding any accumulated data after each round of analysis.

V. EVALUATION

In this section, we introduce the experiment scenario and detail the system performance.

A. Scenario

We evaluate the performance of ShopMiner in two typical indoor environments: (a) an office of $26 \times 14 \text{ m}^2$ to mimic a large store and (b) a twin-bedroom apartment of $13 \times 9 \text{ m}^2$ to mimic a small clothing store (Fig. 12). In both test environments, we hang 20 garments on a clothing rack (2m long and 1.4m high). The space between each adjacent garment is about 5cm. The location of the clothing rack is denoted as the dashed squares in Fig. 12. Each antenna is of 26cm x 26cm and placed on a bracket 0.8m above the floor. The center of each antenna is around 1m apart. We also test ShopMiner with different item-to-customer distances and numbers of nearby customers as in Fig. 13.

B. Popular Category Discovery

We evaluate the performance of popular category discovery in terms of *granularity*, which represents the minimum number

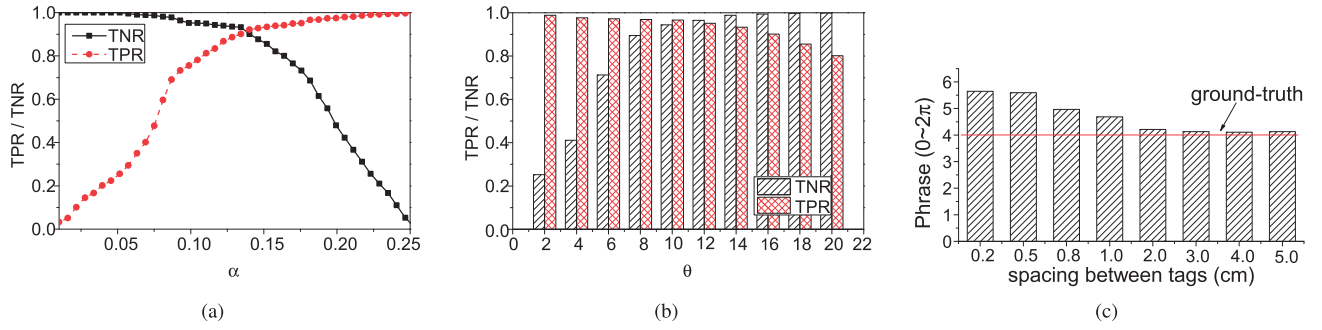


Fig. 14. Impact of (a) confidence level (b) threshold γ and (c) tag spacing on popular category discovery.




Fig. 15. Robustness of popular category discovery with different (a) numbers of antennas, (b) item-to-customer distances, (c) numbers of customers and (d) item-to-customer distances with multiple customers.

of items that ShopMiner can detect when a customer stands still in front of a garment.

Impact of Confidence Level ($1-\alpha$): We vary α from 0.01 to 0.2, and calculate the true positive rates (TPR) and the true negative rates (TNR). We define TPR (TNR) as the fraction of correctly identified body blocking (no body blocking) events among all body blocking (no body blocking) events. Fig. 14(a) plots the TPR and TNR under various confidence levels. ShopMiner achieves a balanced TPR and TNR of over 91% using a confidence level of 0.86, which is used afterwards.

Impact of Threshold θ : As shown in Equation 2, a large θ causes ShopMiner to mistake browsing events for a non-browsing event (*e.g.* one customer stands still for only a short time to examine an item of interest, then moves on). Conversely, a small θ causes ShopMiner to mistaken non-browsing events as a browsing event (*e.g.*, a customer walks slowly around the rack and unintentionally blocks the LOS path). We define TPR (TNR) as the fraction of correctly identified browsing (non-browsing) events among all browsing (non-browsing) events. Fig. 14(b) plots TPR and TNR under a range of thresholds θ . ShopMiner achieves a balanced TPR and TNR of over 92% using a threshold of 12s, which we use in the subsequent evaluations.

Impact of Tag Spacing: If two tags are closely placed, they will both backscatter signals from the reader, thus affecting the actual phase readings of either tag [17]. Such a coupling effect from neighboring tags restricts the minimal spacing of tags for ShopMiner to properly operate. Fig. 14(c) shows the impact of tag spacing on phase readings. We vary the spacing of two tags from 0.2cm to 5cm, and plot the phase readings of one tag as well as its ground truth phase reading calculated by propagation models. As shown, the smaller the tag spacing,

the larger the phase reading deviates from the ground truth. From our evaluation, we recommend a minimal tag spacing of 2cm to avoid strong interference from neighboring tags.

Granularity: In this experiment, a customer browses through a rack of clothes and stands still in front of the ones of interest. The distance between the customer's route and the rack is 0.3m. We repeat the experiment 50 times with different numbers of antennas. Fig. 15(a) shows the detection granularity with different numbers of antennas. The detection granularity is about 6 pieces of clothes with one antenna. With more antennas, the average granularity improves and peaks at 3.2 pieces of clothes with four antennas, yet with larger variance in granularity. This is because the achievable granularity is limited by both the number of antennas (*i.e.*, number of links) as well as the size of human body. As shown in Fig. 2(b), customers tend to stand close to the racks when browsing garments and block links of multiple items, making it difficult to achieve a granularity of 1 piece even with sufficient antennas. Even worse, the closer spacing between garments, the more difficult to achieve finer-grained granularity. Although such a granularity fails to precisely reveal the specific item that the customer is browsing, it can remarkably narrow down the scope of candidates. Furthermore, as garments of the same style are often hung close, such a granularity can guarantee that ShopMiner identifies those popular category. We use four antennas in the following experiments.

Robustness: We vary the distance between the item and the customer (termed as item-to-customer distance) d from 0.2m to 0.6m, and examine how it affects the detection granularity. The antenna is put 2m away from the rack. As Fig. 15(b) shows, the detection granularity decreases moderately with the increase of the item-to-customer distance. Specifically, when the customer is with close proximity (0.3m) to the items,

the detection granularity maintains at 2.9 pieces on average. As we increase d to 0.6m, the granularity drops to 5 pieces on average. This may be because the customer blocks more links between the undesired items and the antenna, hence leading to coarser detection granularity.

We also test the detection granularity with multiple customers. The distance between the rack and customers is about 0.3m. Fig. 15(c) shows the granularity with different numbers of customers browsing clothes simultaneously. When there is only one customer, ShopMiner performs best, with a detection granularity of 3.2 pieces on average. With more customers, *e.g.*, $n = 3$, the detection granularity drops to 5.3 pieces on average, as more customers will introduce more complex multipath reflections. We also investigated the impact of item-to-distance in multi-customer scenarios. Fig. 15(d) plots the granularity with the following settings. Group 1: three customers at a distance of 0.3m from the rack for comparison. Group 2: three customers at a distance of 0.6m from the rack. Group 3: three customers with each at a distance of 0.3m, 0.6m and 0.9m from the rack, respectively. As shown, multiple customers further away (0.6m) results in an average granularity of 6.4 pieces of clothes. The detection granularity is worse with multiple customers at diverse distances (Group 3), indicating that complicated signal propagation will lead to performance degradation in detection granularity.

Summary: ShopMiner can detect popular categories with a TPR of 92%. The detection granularity degrades slightly with the increase of item-to-customer distance and the number of customers. Deploying more antennas help improve the detection granularity but the granularity is also limited by the size of human body.

C. Hot Item Identification

We evaluate hot item identification in terms of *TPR* and *FPR*. TPR is defined as the proportion of successfully detected pick-out and turn-around actions among all pick-out and turn-around actions. FPR is defined as the proportion of mistaken pick-out or turn-around actions over all non-pick-out and non-turn-around actions.

Accuracy: In this experiment, a volunteer randomly picks out or turn around different garments for 20 times. Two other volunteers hang out around the rack as interference. The experiment was repeated 20 times in two scenarios by 10 different volunteers. The ground-truth is recorded by video. We first plot the Receiver Operating Characteristic (ROC) curves for *detecting* pick-out and turn-around actions in Fig. 16. ShopMiner achieves a balanced detection accuracy of 92% with a FPR of 13%. Table I further shows the confusion matrix to *distinguish* pick-out and turn-around actions. ShopMiner achieves TPRs of 94% and 96% for turn-around and pick-out actions, respectively. Turn-around actions are misclassified as pick-out with a 7% probability, while pick-out actions are misclassified as turn-around with a 5% probability. The result demonstrates that the auto-correlation based detection scheme can successfully distinguish pick-out and turn-around actions with high accuracy. The misclassification cases are mainly due to the diversity in pick-out and turn-around actions, as well




Fig. 16. ROC curve for pick-out/turn-around detection.




Fig. 17. Phases of dropping off garments and picking out two adjacent garments by two customers.

TABLE I
CONFUSION MATRIX OF PICK-OUT/TURN-AROUND IDENTIFICATION

Ground-truth	Predicted					
	Turn around			Pick out		
	1	2	3	1	2	3
Turn around	187	184	178	13	16	22
Pick out	9	10	13	191	190	187

as some corner cases *e.g.* accidentally dropping off garments. As shown in Fig. 17, the phase trend of dropping off garments will deviate from those of its neighboring tags, which will be misclassified into a pick-out action. Some pick-out actions may not exhibit notable changes (as the one in red), which leads to higher similarity of phase trends among neighboring tags, and will be misclassified into a turn-around action. Conversely, if a customer half-turns a garment so that only one neighboring garment is forced to vibrate, the similarity of phase trends will be low, which leads to a misclassification into pick-out.

Robustness: As shown in Fig. 16, the detection accuracy of pick-out/turn-around actions decreases slightly with more customers, yet ShopMiner still achieves an accuracy of 85% with a FPR of 22% with three customers. This is because multiple customers introduce complex propagation environment, which introduces phase disturbance to each item. Similarly, as shown in Table I, the misclassification rate rises moderately with the increase of customers. Specifically, when there are two customers in front of the same rack, ShopMiner achieves an average misclassification rate of 8% and 5% for turn-around and pick-out, respectively. This index increases slightly with three customers, and finally peaks 11% and 7% with four customers. This is because when multiple customers turn around garments on the same rack, the garments near the desired one will be pushed by multiple customers. Thus their phases will change irregularly, which degrades the identification performance.

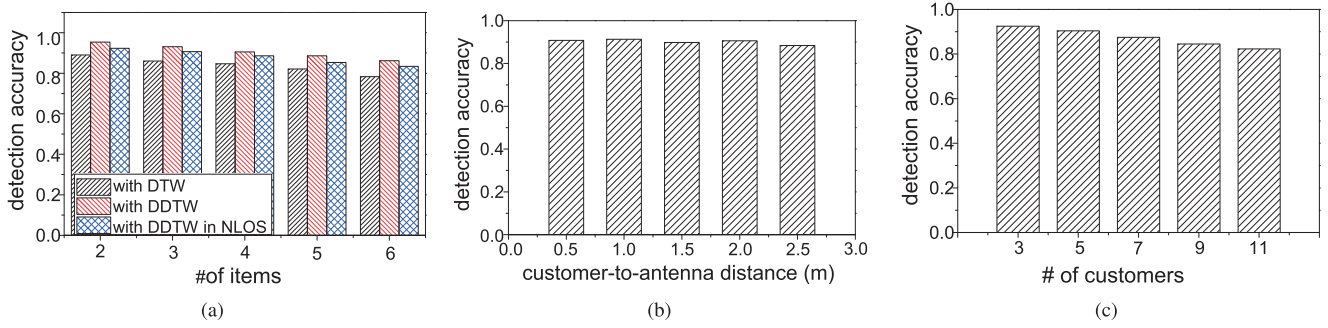


Fig. 18. Performance of correlated item excavation with different (a) numbers of items (b) customer-to-antenna distances and (c) numbers of interfering customers.

Summary: ShopMiner achieves an overall accuracy of 92% with a FPR of 13% for pick-out/turn-around detection. The TPR of action identification is 96% and 94% for pick-out and turn-around, respectively. ShopMiner is insensitive to up to three customers.

D. Correlated Item Excavation

We evaluate the correlation item excavation in terms of *detection accuracy*, which is defined as the proportion of correctly identified correlated items over all correlated items.

Accuracy: In this experiment, a volunteer carries different numbers of garments and walks around the reader's reading zone in two settings: (1) there is no blockage between the volunteer and the reader, *i.e.* LOS condition; (2) there is one person standing close to the reader to create NLOS propagation between the volunteer and the reader. The experiment was repeated 200 times by 10 different volunteers for each setting. Fig. 18(a) shows the detection accuracies using DTW and DDTW metrics. As shown, the detection accuracies decrease slightly with the number of correlated items in both LOS and NLOS settings. Comparing the performance using DTW and DDTW metrics, with one or two items, the detection accuracy retains around 91% and 83% for the DDTW and DTW metrics, respectively. With more garments, the performance gap under different distance metrics becomes larger, and finally peaks at 10%. It manifests that the DDTW metric significantly improves the detection accuracy. Comparing the performance under LOS and NLOS propagation, the accuracy for correlated items discovery in the NLOS setting is around 5% lower than that in the LOS setting, but still remains above 82% with six items. Such a drop in accuracy is because when the LOS is blocked by the person, there will be an abrupt change in phase readings, and the changes are different for each tag, which will decrease the similarity of phase trends among the tags.

Robustness: We evaluate the impact of the customer-to-antenna distance on the correlated items detection accuracy in Fig. 18(b). When the customer stands close to the antenna (0.5m or 1m), the detection accuracy retains above 90%. The accuracy drops to 88% when the customer is 2m from the antenna. Hence the increase of customer-to-antenna distance will not led to significant performance degradation.

We further examine the impact of customer population on the detection accuracy. In this experiment, we arrange




Fig. 19. Employee detection accuracy.

different number of customers to walk around, with each customer bringing different number of clothes. As shown in Fig. 18(c), ShopMiner achieves an overall accuracy of 92% when there are three customers. The detection accuracy then decreases slightly with more customers. However, ShopMiner still achieves an detection accuracy over 85% with 11 volunteers, showing that ShopMiner is robust to customer population changes.

Summary: ShopMiner achieves an overall detection accuracy of over 93% with one customer, and that of over 85% with six customers for correlated items discovery. Also ShopMiner is insensitive to the number of customers and the antenna-to-customer distance.

E. Employee Recognition

As discussed in Section IV-E, employees in the shops cause false correlated item discovery and ShopMiner avoids the impact of employees by filtering correlated items with tag IDs from employees. Hence it is crucial to correctly read the tags attached to employees. We evaluate the reading rate of an RFID tag attached to a volunteer with different on-body tag placements. Fig. 19 shows the detection accuracy of employee tag with several common tag placements of *on chest*, *on waist*, *on wrist with motion* and *on wrist without motion*. As shown, the chest/waist placement yield detection rates of above 85%, while wrist placement with motion achieves only 40% accuracy. This is because tags attached on wrist are likely to contact with the skin, which dramatically affects the backscatter signal. The results suggest employees wear an RFID tag on chest or waist to ensure high detection accuracy so as to eliminate their interference on correlated item discovery.




Fig. 20. Bookstore scenario.




Fig. 21. Performance in bookstore case.

F. Extending to Bookstores

In this experiment, we explore extending ShopMiner to bookstores, another common scenario where items are of similar material and shape, and are tidily placed. Fig. 20 illustrates the experimental settings, where a pile of books are placed on a shelf. We attach an RFID tag on the spine of each book. To avoid the strong coupling effect as in Fig. 14(c), we choose relatively thick books for evaluation, which results in a tag spacing of roughly 2cm to 4cm. Different from clothing stores, a customer usually picks out a book from the shelf directly without turning it around. Hence we only evaluate the performance of pick-out detection and correlated item excavation. In this experiment, two volunteers randomly picked out different numbers of books from the shelf. For each number of books, the volunteers performed the pick-out actions for 100 times. Fig. 21 shows the results of pick-out detection and correlated item excavation (one correlated item is actually pick-out detection). The pick-out detection accuracy is around 95%, and the correlated item detection accuracies are all above 90% for up to 5 books. Such a accuracy for books is comparable to that for garments, indicating the feasibility to apply ShopMiner in bookstores.

To further evaluate whether ShopMiner can be extended to other stores, we select 10 daily goods of different material and sizes. An RFID tag is attached to each item and we evaluate the reading rates of tags within one second. Table II summarizes the results. As shown, the reader fails to read metal items, which limits the applicability of ShopMiner to grocery stores and electronics stores. Items made of paper, plastic, wood, *etc.*, have similar or slightly higher reading rates compared with books and clothes, indicating that ShopMiner may be extended to monitor the customer interactions with these items. Items of smaller sizes *e.g.* pencil and marker pen have lower reading rates, and we may encounter more severe interference for closely placed small items. In summary, the reading rate of ShopMiner is affected by the material and size of the items.

TABLE II
READING COUNTS OF DIFFERENT ITEMS

Item	Reading Count	Item	Reading Count
Paper Box	130	Plastic Cup	128
Wooden Box	126	Glass Cup	117
Iron Box	0	Pottery Cup	120
Cotton Cloth	108	Pencil	70
Book	113	Marker Pen	103

Therefore, it is non-trivial to extend ShopMiner to grocery stores, where the products are of diverse material and sizes. The metal objects can also interfere with the tags attached to other items, thus leading to missed tag readings.

VI. RELATED WORKS

While wireless signals such as Wi-Fi are prevalent for activity recognition in smart homes and smart offices [19]–[23], RFID-based solutions are preferred in human-object interaction such as shopping behavior monitoring and recognition. ShopMiner is inspired by this trend of wireless sensing, and is particularly related to the following categories of research.

Offline Shopping Behavior Mining: Despite the academic and commercial success in online shopping data acquisition, there have been few for offline shops. You *et al.* [24] utilized smartphones to record shopping time in physical stores. Kanda *et al.* [25] designed a sensor network in retail stores to track and cluster consumer locations, and further infer hot items. We advance this area through RFID technologies, and acquire more comprehensive data of shopping behaviors.

Shopping Behavior Recognition With Wearable Sensors: Lee *et al.* [26] examined the impact of six kinds of activities on customer behaviors using smartphone sensing. Sen *et al.* [27] also utilized smartphones to detect pre-defined shopping activities in retail stores. Rallapalli *et al.* [28] proposed a customer tracking and browsing behavior sensing system using smart glasses. IRIS [29] combined a smartphone and a smartwatch to recognize item-level interactions and further infer episode-level attributes in grocery shops with high accuracy. While these works provide shopping behavior data, they require access to customers' phones or wearable devices, which may degrade the shopping experience.

RFID-Based Context Sensing: OTrack [30] designed an RSS-based RFID system to track the tag order for baggage sorting. STPP [31] designed a phase profiling technique for the relative localization of RFID tags. Tagoram [32] tracked mobile RFID tags at a centimeter accuracy using hologram techniques. Our work is inspired by these works in tag tracking, but our focus is to leverage the phase pattern to infer customer behaviors, rather than the location of tags.

RFID Systems in Physical Stores: Melia *et al.* [33] deployed an RFID system in an apparel retail store for both operational (*e.g.*, inventory management) and experiential (*e.g.*, interactive fitting room) enhancement. Our work also utilizes RFID technology, yet focuses on the retailer side, *i.e.* mining shopping behaviors. CBID [10], Tagbooth [11] and IDSense [34] are the most closely related works. CBID [10] exploited Doppler effects to detect customer behavior in shops. Tagbooth [11]

used RSS patterns to identify the pick-up actions in retail stores. IDSense [34] extracted phase features of a single tag to classify human-item interactions such as *still*, *translation*, *rotation* and *swing* via machine learning. Our work differs from these works in three aspects. (1) We incorporate three key factors that are essential to retailers, *i.e.*, which items the customers browse, they show a interest in, and they pair up. CBID and Tagbooth only detect pick-ups and infer correlated items, and IDSense only infers hot items by counting the number of human-item interactions. (2) We harness phase information to mining the customer shopping behavior, which is more accurate than RSS or Doppler measurements. (3) Although both IDSense and our system utilize phase information, we jointly leverage phase patterns of *neighboring* tags, which enables finer-grained customer gesture detection, while IDSense only uses the phase of a *single* tag.

VII. CONCLUSION

In this paper, we present the design, implementation and evaluation of ShopMiner, an RFID-based shopping behavior mining system for physical clothing stores. With an RFID tag attached to each garment, ShopMiner could detect which garments customers stop beside, pick out, turn around, or pair up. Such shopping behavior data could benefit retailers to discover popular categories, hot items, and correlated pairs for better trading strategies and tie-in promotions. We examine the accuracy and robustness of ShopMiner in various testing scenarios. Results show that ShopMiner achieves high accuracy in customer shopping behavior identification and holds potential for practical deployment.

REFERENCES

- [1] L. Shangguan *et al.*, "Shopminer: Mining customer shopping behavior in physical clothing stores with cots rfid devices," in *Proc. ACM SenSys*, Nov. 2015, pp. 113–125.
- [2] D. R. Bell and J. M. Lattin, "Shopping behavior and consumer preference for store price format: Why 'large basket' shoppers prefer EDLP," *Marketing Sci.*, vol. 17, no. 1, pp. 66–88, Feb. 1998.
- [3] Y. H. Cho, J. K. Kim, and S. H. Kim, "A personalized recommender system based on Web usage mining and decision tree induction," *Exp. Syst. Appl.*, vol. 23, no. 3, pp. 329–342, Oct. 2002.
- [4] R. Kohavi, N. J. Rothleder, and E. Simoudis, "Emerging trends in business analytics," *Commun. ACM*, vol. 45, no. 8, pp. 45–48, Aug. 2002.
- [5] M. Popa, A. K. Koc, L. J. M. Rothkrantz, C. Shan, and P. Wiggers, *Kinect Sensing of Shopping Related Actions*. Berlin, Germany: Springer, pp. 91–100, 2011.
- [6] M. Popa *et al.*, "Analysis of shopping behavior based on surveillance system," in *Proc. SMC*, Oct. 2010, pp. 2512–2519.
- [7] Y. Liu, Y. Zhao, L. Chen, J. Pei, and J. Han, "Mining frequent trajectory patterns for activity monitoring using radio frequency tag arrays," *IEEE Trans. Parallel Distrib. Syst.*, vol. 23, no. 11, pp. 2138–2149, Nov. 2012.
- [8] L.-A. Tang *et al.*, "On discovery of traveling companions from streaming trajectories," in *Proc. ICDE*, Apr. 2012, pp. 186–197.
- [9] T. Staake, F. Thiesse, and E. Fleisch, "Extending the EPC network: The potential of RFID in anti-counterfeiting," in *Proc. ACM SAC*, Mar. 2005, pp. 1607–1612.
- [10] J. Han *et al.*, "Cbaid: A customer behavior identification system using passive tags," *IEEE/ACM Trans. Netw.*, vol. 24, no. 5, pp. 2885–2898, Oct. 2016.
- [11] T. Liu, L. Yang, X.-Y. Li, H. Huang, and Y. Liu, "Tagbooth: Deep shopping data acquisition powered by rfid tags," in *Proc. INFOCOM*, May 2015, pp. 1670–1678.
- [12] T. Bouwmans, F. El Baf, and B. Vachon, "Background modeling using mixture of Gaussians for foreground detection—A survey," *Recent Patents Comput. Sci.*, vol. 1, no. 3, pp. 219–237, Nov. 2008.
- [13] P. Zhang, J. Gummesson, and D. Ganesan, "Blink: A high throughput link layer for backscatter communication," in *Proc. ACM MobiSys*, Jun. 2012, pp. 99–112.
- [14] N. Fet, M. Handte, and P. J. Marrón, "A model for wlan signal attenuation of the human body," in *Proc. ACM UbiComput*, Sep. 2013, pp. 499–508.
- [15] T. K. Vintsyuk, "Speech discrimination by dynamic programming," *Cybern. Syst. Anal.*, vol. 4, no. 1, pp. 52–57, Jan. 1968.
- [16] E. J. Keogh and M. J. Pazzani, "Derivative dynamic time warping," in *Proc. SIAM SDM*, Apr. 2001, pp. 1–11.
- [17] T. Wei and X. Zhang, "Gyro in the air: Tracking 3D orientation of batteryless Internet-of-Things," in *Proc. ACM MobiCom*, Oct. 2016, pp. 55–68.
- [18] S. Basagni, "Finding a maximal weighted independent set in wireless networks," *Telecommun. Syst.*, vol. 18, no. 1, pp. 155–168, Sep. 2001.
- [19] Q. Pu, S. Gupta, S. Gollakota, and S. Patel, "Whole-home gesture recognition using wireless signals," in *Proc. ACM MobiCom*, Sep. 2013, pp. 27–38.
- [20] W. Wang, A. X. Liu, M. Shahzad, K. Ling, and S. Lu, "Understanding and modeling of WiFi signal based human activity recognition," in *Proc. ACM MobiCom*, Sep. 2015, pp. 65–76.
- [21] X. Zheng, J. Wang, L. Shangguan, Z. Zhou, and Y. Liu, "Smokey: Ubiquitous smoking detection with commercial wifi infrastructures," in *Proc. INFOCOM*, Apr. 2016, pp. 1–9.
- [22] D. Zhang, H. Wang, and D. Wu, "Toward centimeter-scale human activity sensing with Wi-Fi signals," *Computer*, vol. 50, no. 1, pp. 48–57, Jan. 2017.
- [23] G. Wang, Y. Zou, Z. Zhou, K. Wu, and L. M. Ni, "We can hear you with Wi-Fi!" *IEEE Trans. Mobile Comput.*, vol. 15, no. 11, pp. 2907–2920, Nov. 2016.
- [24] C.-W. You, C.-C. Wei, Y.-L. Chen, H.-H. Chu, and M.-S. Chen, "Using mobile phones to monitor shopping time at physical stores," *IEEE Pervasive Comput.*, vol. 10, no. 2, pp. 37–43, Jun. 2011.
- [25] T. Kanda, D. F. Glas, M. Shiomi, H. Ishiguro, and N. Hagita, "Who will be the customer?: A social robot that anticipates people's behavior from their trajectories," in *Proc. ACM UbiComp*, Sep. 2008, pp. 380–389.
- [26] S. Lee, C. Min, C. Yoo, and J. Song, "Understanding customer malling behavior in an urban shopping mall using smartphones," in *Proc. ACM UbiComp*, Sep. 2013, pp. 901–910.
- [27] S. Sen *et al.*, "Accommodating user diversity for in-store shopping behavior recognition," in *Proc. ACM ISWC*, Sep. 2014, pp. 11–14.
- [28] S. Rallapalli, A. Ganesan, K. Chintalapudi, V. N. Padmanabhan, and L. Qiu, "Enabling physical analytics in retail stores using smart glasses," in *Proc. ACM MobiCom*, Sep. 2014, pp. 115–126.
- [29] M. Radhakrishnan, S. Eswaran, A. Misra, D. Chander, and K. Dasgupta, "Iris: Tapping wearable sensing to capture in-store retail insights on shoppers," in *Proc. PerCom*, Mar. 2016, pp. 1–8.
- [30] L. Shangguan *et al.*, "OTrack: Towards order tracking for tags in mobile RFID systems," *IEEE Trans. Parallel Distrib. Syst.*, vol. 25, no. 8, pp. 2114–2125, Aug. 2014.
- [31] L. Shangguan, Z. Yang, A. X. Liu, Z. Zhou, and Y. Liu, "Relative localization of RFID tags using spatial-temporal phase profiling," in *Proc. USENIX NSDI*, May 2015, pp. 251–263.
- [32] L. Yang *et al.*, "Tagoram: Real-time tracking of mobile RFID tags to high precision using cots devices," in *Proc. ACM MobiCom*, Sep. 2014, pp. 237–248.
- [33] J. M. Seguí *et al.*, "Enhancing the shopping experience through RFID in an actual retail store," in *Proc. ACM UbiComp Adjunct*, Sep. 2013, pp. 1029–1036.
- [34] H. Li, C. Ye, and A. P. Sample, "IDSense: A human object interaction detection system based on passive UHF RFID," in *Proc. ACM CHI*, Apr. 2015, pp. 2555–2564.




Zimu Zhou (S'13) received the B.E. degree from the Department of Electronic Engineering, Tsinghua University, in 2011, and the Ph.D. degree from the Department of Computer Science and Engineering, The Hong Kong University of Science and Technology, in 2015. He is currently a Post-Doctoral Researcher with the Computer Engineering and Networks Laboratory, ETH Zurich.



Longfei Shangguan (S'11) received the B.S. degree in software engineering from Xidian University in 2011, and the Ph.D. degree from HKUST in 2015. He is currently a Post-Doctoral Research Associate with the Computer Science Department, Princeton University. His research interests include pervasive computing and RFID system. He is a Student Member of the ACM.



Lei Yang received the B.S. and Ph.D. degrees from the School of Software and the Department of Computer Science and Engineering, Xian Jiaotong University. He was a Post-Doctoral Fellow with the School of Software, Tsinghua University. He is currently a Research Assistant Professor with the Department of Computing, The Hong Kong Polytechnic University. His research interests include RFID and backscatters, wireless and mobile computing, pervasive computing, and smart home. He was a recipient of the Best Paper Awards of MobiCom'14 and MobiHoc'14, Best Video Award of MobiCom'16, and the ACM China Doctoral Dissertation Award.



Xiaolong Zheng (S'13–M'15) received the B.E. degree from the School of Software Technology, Dalian University of Technology, in 2011, and the Ph.D. degree from the Department of Computer Science and Engineering, The Hong Kong University of Science and Technology, in 2015. He is currently a Post-Doctoral Researcher with the School of Software and TNLIST, Tsinghua University. His research interests include wireless sensor networks and pervasive computing. He is a member of the ACM.



Yunhao Liu (S'03–M'04–SM'06–F'15) received the B.S. degree in automation from Tsinghua University, China, in 1995, and the M.S. and Ph.D. degrees in computer science and engineering from Michigan State University in 2003 and 2004, respectively. He is currently a Chang Jiang Professor with the School of Software and TNLIST, Tsinghua University. His research interests include wireless sensor network and RFID, peer-to-peer computing, and pervasive computing.