SugarTrail : Indoor Navigation in Retail Environments without Surveys and Maps

Aveek Purohit, Zheng Sun, Shijia Pan, and Pei Zhang
Department of Electrical and Computer Engineering
Carnegie Mellon University
{aveek.purohit,zheng.sun,shijia.pan,pei.zhang}@sv.cmu.edu

***Abstract*—A system that helps people navigate in indoor environments on a fine-grained level can enable a variety of pervasive computing applications in retail environments. Existing indoor navigation systems rely on extensive RF tagging surveys and accurate floor plans. These prerequisites are often impractical in indoor environments.**

In this paper, we present SugarTrail, a system for indoor navigation assistance in retail environments that minimizes the need for active tagging and does not require existing maps. By leveraging the structured movement patterns of shoppers in retail store environments, the system provides higher accuracy than existing radio finger-printing approaches. With minimal setup and active user participation, the system automatically learns user movement pathways in indoor environments from radio-frequency and magnetic signatures. These pathways are clustered and used to automatically build a navigable virtual roadmap of the environment. We present results from a campus testbed and from actual radio measurements collected in an operational supermarket to show that SugarTrail system can navigate users with a success rate of $> 85\%$ and an average accuracy of $0.7m$.

I. Introduction

Fine-grained indoor navigation can enable a variety of pervasive computing applications in retail store environments. For example, using indoor navigation assistance, retail stores can guide consumers to their desired products or track movement patterns of users to fine-tune product placement. The widespread adoption of such navigation systems depends on their accuracy and ease of setup.

Indoor navigation is hard due to the difficulty in localizing people inside buildings, where GPS is not available and the environment is prone to multipath effects reducing the accuracy of range-based localization techniques relying on sound or RF signal propagation. Approaches for indoor navigation have been proposed in literature, primarily using radio-frequency (RF) based localization systems [1], [23]. RF-fingerprinting methods are popular for indoor retail or office environments because of their ability to largely use existing infrastructure of wireless access points. However, these systems have certain requirements that make them impractical, in particular:

Need for Extensive Active Tagging of Fingerprints – To overcome the highly variable and noisy indoor environment, prior radio-based localization systems involve extensive deployment efforts in collecting RF-fingerprints, and tagging each fingerprint with a location on a map. This is done manually during deployment or through crowd-sourcing [3], [15]. This approach requires people with special devices actively assigning physical locations to every RF-signature collected. Sufficient and repeated *active* user participation is impractical in many application scenarios.

Reliance on Accurate/Updated Maps – Existing indoor navigation systems rely on locating users on a map and then using the map to find walkable pathways to the users' desired destination. Thus, navigation approaches based on localization systems require updated maps of the environment. This assumption of having updated indoor maps is not valid in many applications scenarios. For example, retail stores and supermarkets have soft demarcation between spaces, such as cubicles, furniture or aisles, which are not reflected in floor plans.

In this paper, we present SugarTrail, a system for indoor navigation assistance in retail environments that minimizes the need for active tagging and does not require existing maps. By leveraging the structured movement patterns of shoppers in retail store environments, the system provides higher accuracy than existing radio finger-printing approaches.

The system uses low-cost stationary sensor nodes equipped with off-the-shelf radios, which can be casually placed in the environment by untrained personnel without location measurement. To acquire knowledge of the physical environment *without* a map, the system collects radio and compass signatures to continually record paths traversed by users (without any active participation on the part of the user such as tagging). Using this information, paths are automatically aggregated into path-clusters and a navigable *Virtual Roadmap* (VRM) of the indoor environment is built by the system. Using this roadmap, the system guides users to their desired destination.

The key technical contributions of our system are,

- An algorithm that provides fine-grained indoor navigation guidance by automatically learning a virtual roadmap from user movement traces. The algorithm leverages the structured movement patterns of retail environments to achieve higher accuracy than existing approaches.
- A novel technique for identifying user movement paths from radio and magnetometer measurements that can overcome multi-path distortions and magnetic interference common to indoor environments.
- Extensive experimental validation on campus as well as in a real supermarket during operational hours.

To fully evaluate our system in a realistic usage scenarios, we built the SugarTrail sensor node hardware and deployed our system, on campus, and in the New Wing supermarket in

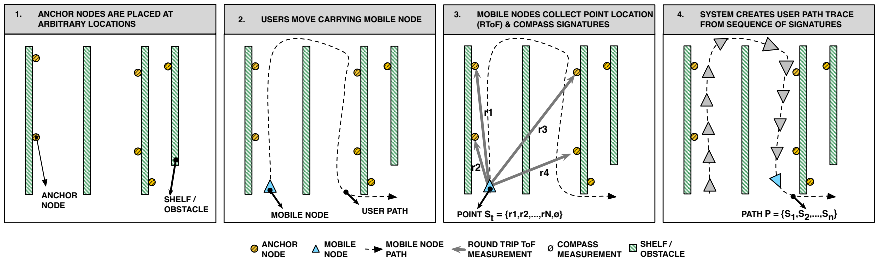

Fig. 1. The figure shows the process of recording user path traces in a retail store environment. 1) Anchor nodes are distributed in the environment without the need to measure their location. 2) Mobile nodes attached to users move through the environment and 3) Mobile nodes periodically collect magnetometer measurements and radio measurements from anchor nodes that are in range. 4) The sequence of signatures define a path and are transmitted to a base station.

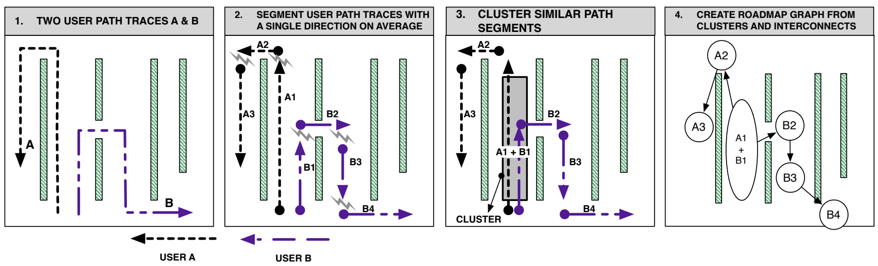

Fig. 2. An illustration of the SugarTrail algorithm using path traces from users A and B to build the VRM graph. 2) The algorithm breaks a path into path-segments, 3) clusters similar path segments, and 4) uses the aggregated information to build the VRM with path-clusters as vertices and their connections as edges. For visual clarity, the users' curvy paths are drawn as straight lines.

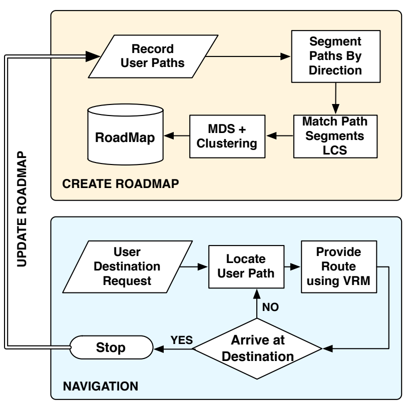

Fig. 3. The figure shows a flowchart system operation. The system has two main modes of operation. Top) Creating/updating of the "Virtual Roadmap". Bottom) Navigation mode, where the system guides users to their desired destination.

Santa Clara, California (over a 30-day period) during business hours. Based on this data, results show that SugarTrail can accurately navigate users throughout the entire supermarket, with a success rate of more than 85% and average accuracy of 0.7m.

II. System Overview

The SugarTrail system provides accurate navigation in indoor environments without access to updated maps and active RF-surveys. The system leverages the structured movement pathways of users in retail environments to enhance accuracy and minimize active user tagging.

A. *Structure of Retail Environments*

Retail environments such as grocery stores and supermarkets have structured movement pathways for users [11]. The stores typically consist of aisles with products placed in shelves along them. In addition, products may be placed on islands along the broader walkway along the perimeter called the *racetrack*. The product placement limits the trajectory of user movement to approximately unidirectional paths within aisles or along he perimeter. The SugarTrail system leverages this structure to automatically learn the graph of walkable pathways (VRM) in a store from user movement traces (measured using radio and magnetic signatures). Of course, individual users may not always walk along the aisles, may turn around, or linger around. The SugarTrail system uses magnetic signatures to discard such traces that do not contribute to extracting the VRM.

B. *System Operation*

The SugarTrail system operation can be described in four steps – 1) deployment of sensors, 2) recording of user path traces, 3) automatic creation of a Virtual Road Map (VRM), and finally 4) navigation of users. This section gives an overview of our approach.

1. Deployment of sensors: The system involves placing, without location or orientation measurement, a number of *stationary* low-cost wireless nodes equipped with off-the-shelf 802.15.4a radios in the environment. A wireless node with radio and compass is attached to *mobile* users moving in the environment. Since, no measurements are needed in placing anchor nodes, the deployment effort is low and can be performed by non-expert users. For example, in a retail store, the anchors can be placed on merchandise racks in sufficient density for mobile nodes to hear 4 or 5 stationary nodes at most locations.

2. Recording of user path: The mobile nodes attached to users, automatically record the users' path as a sequence of magnetometer readings and radio round-trip time-of-flight measurements from the stationary nodes, without participation from the user. The goal of the radio measurements is to provide approximate signatures for each location along the user's path, while the magnetometer measurements provide a unique signature for the users direction at each location on the path. Figure 1 shows this process of collecting path traces in a retail store environment.

3. Creating the VRM: Next, the system learns a "Virtual Roadmap" of the environment by comparing and combining the recorded user paths into path-clusters, without any supervised input of the physical location of the paths. This VRM is a representation of the usable pathways in the deployment region and their interconnections with each other.

SugarTrail aggregates paths from multiple users into uni-directional path-clusters to quickly build the VRM with maximum coverage of the area. To achieve this, SugarTrail analyzes the recorded paths and divides them into uni-directional segments that correspond to a distinct physical pathway, such as a hallway or a supermarket aisle, based on the modal direction of the users. This is based on the observation that users in a pathway tend to travel in a few common directions on the whole.

Figure 2 illustrates the SugarTrail algorithm for building a VRM from the path traces of two users. On the basis of the segmentation, user paths are represented as a sequence of path segments (corresponding to physical pathways) that connect to each other, such as the segments A1 to A3 and B1 to B3 in Figure 2. Subsequently, the SugarTrail system finds similar segments from paths of multiple users and clusters them together into path-clusters. The connections between path-clusters are determined from the connectivity of their constituent path segments. Finally, path-clusters and their connections are modeled as the vertices and edges of a single directed graph (VRM). For simplicity, figure 2 shows user paths as straight lines, however, the system is designed and evaluated for realistic human paths

4. Navigation: Finally, the system provides users with directions to their desired destination using the VRM. The system first locates users by matching their current path to the paths in the VRM, and then uses the connectivity information in the VRM to suggest a route. The route is conveyed to the user as a sequence of direction prompts based on compass directions that relative to the user's local magnetic field (learned from recorded user path traces).

C. Semantic Annotation of VRM

The SugarTrail system automatically creates the VRM. However, in order for users to look for a particular product such as "cereal" or "milk", the path clusters must be annotated with the names of the products located on these paths. The mechanism considered is to connect barcode or RFID scanners, already used by store employees for auditing products, with the SugarTrail mobile node. This enables the system to automatically annotate the VRM clusters identified by the mobile node with the products scanned.

Thus, a large number of users (customers) contribute to building a higher quality VRM but only a few (employees) are needed for semantic annotation, without any need for active participation from both sets of users. In contrast, RF-fingerprinting based methods require every surveyer to actively tag each rf-fingerprint to a location on a map.

III. Algorithm Design

In this section, we present the details of the SugarTrail algorithm including practical real-world considerations. A detailed flowchart of system operation is illustrated in Figure 3. As shown in the figure, the system operates in two primary modes, namely, 1) the creation or update of the VRM and 2) navigation using the VRM. Each step in the flowchart of system operation is described in detail in the following subsections.

A. Recording User Path Traces

To learn the navigable pathways of the environment without a known physical map, the SugarTrail system records user paths as a sequence of user location and direction signatures. The system employs two novel techniques for recording user paths that enable it to overcome noise inherent in indoor environments, 1) sequence of radio-based location signatures based on round-trip time-of-flight measurements, and 2) magnetometer-based local direction signatures. The two types of signatures are described below.

1) Round-Trip Time-of-Flight Signatures: The mobile SugarTrail nodes use round-trip time-of-flight (RToF) measurements to estimate location with respect to stationary anchor nodes. The round-trip time-of-flight method measures the elapsed time between the host node sending a data signal to the remote node, and receiving an acknowledgment from it.

The SugarTrail nodes use physical layer timestamps and hardware-generated acknowledgments, supported by their nano-LOC radios, to compute RToF measurements [13]. However, prior experiments analyzing the distance estimates from RToF measurements in various space configurations, show that the distance-to-measurement correlation is not high for multi-path rich indoor environments [17].

Therefore, in order to compensate for individual signature variation, SugarTrail uses a set of RToF measurements from multiple anchors, distributed around the environment, to create multi-dimensional signature vectors.

The user obtains RToF signatures while walking. Consequently, the readings from all anchors are not obtained for a single point but over a small section of the path. Each round of measurements lasts for a period of 1 second and the location of the signature can be approximated to a point on the path.

2) Magnetometer Direction Signatures: Recording a user's path involves knowing the sequence of location signatures and the spatial relationship (direction) between them. If a navigation system knows the location of the user in Cartesian coordinates, the users direction of movement can be determined directly. However, since SugarTrail navigates users without existing maps, the spatial relationship between nodes must be measured. SugartTrail uses magnetometer signatures to determine the spatial relationship.

In indoor environments, magnetometers can be significantly affected by soft-iron effects due to surrounding metallic structures such as cubicles, appliances, support beams, etc., as well as by hard-iron effects due to power lines [5]. To deal with the indoor magnetic interferences, SugarTrail uses the magnetometer as a relative and local direction signature as opposed to a global orientation measurement. Specifically, the system uses raw magnetometer direction signatures to identify two major features of the user path,

- **Points where the user changes direction** corresponding to a transition between pathways. This is used to detect path junctions where paths intersect. These junctions are later used to divide user path traces into segments that have a single direction.
- **User path segments that have similar directions** for determining co-located path segments and aligning them in the matching phase.

SugarTrail leverages two types of location-dependent signatures, namely RToF signatures and magnetometer readings. Suppose a user walks along a path P, each single location point p_i will be represented by a N-dimensional RToF signature vector

$$\overrightarrow{R_i} = \{r_1, r_2, \ldots, r_N\}, \tag{1}$$

where r_i is the RToF measurement from the ith anchor and N is the total number of anchors in the environment, as well as a direction measurement from the magnetometer ϕ_i, together forming a signature set $\overrightarrow{S_i} = \{\overrightarrow{R_i} : \phi_i\}$. We assume that the orientation of the mobile node is fixed with respect to the user. This assumption is valid for our usage scenarios, where users hold the nodes in their hands or the nodes are mounted on shopping carts.

B. Segmentation of Paths

After recording user path signatures, the system builds the VRM representing usable pathways and their interconnections with each other. To achieve this, SugarTrail analyzes recorded paths and divides them into segments that correspond to distinct physical pathways, such as a supermarket aisle.

To isolate path segments belonging to the same physical pathway, we observe that human motion along physical pathways (such as aisles) tends to be limited to a few directions on a macro-scale. For example, people tend to move forward or backward in an aisle. While, on a micro-scale people may move in diverse directions resembling a random way-point walk, by filtering out the micro-scale motion, the system can extract path segments having a single major direction.

Based on the above intuition, SugarTrail first computes the difference in direction between consecutive points in a path, then detects points at which a significant change in direction occurs. These points are marked as candidate turn points. To reduce false detection of segments caused by micro-scale directional variations, SugarTrail computes the modes m_l and m_r of two sliding windows before and after each candidate turn point along a path trace. The mode of the first window m_l is then compared to that of the second window m_r. When the difference exceeds a direction change threshold $_\triangle m_{THR}$, the point is marked as a valid turn point. Finally the paths are segmented at each of the valid turn points. The pseudo-code is described in Algorithm Box 1.

Algorithm 1 Path Segmentation

k=0, do the following for each path P_i ($i \in \{1, I\}$)
for each point p_j in path P_i **do**
    Find difference in direction between consecutive points p_j and p_{j+1}: $_\triangle\phi_{j,j+1} = \phi_{j+1} - \phi_j$
    if $_\triangle\phi_{j,j+1} >_\triangle \phi_{THR}$ **then**
        Mark the turn index t_k as $j+1$, then $k = k+1$.
    end if
end for
for each turn index t_k ($k \in \{1, K\}$) **do**
    Find the mode over sliding window before t_k: $m_l = \text{mode}(\{\phi_{t_k - win\_size:t_k}\})$
    Find the mode over sliding window after t_k: $m_r = \text{mode}(\{\phi_{t_k:t_k+win\_size}\})$
    Find difference $_\triangle m = \text{abs}(m_l - m_r)$
    if $_\triangle m >_\triangle m_{THR}$ **then**
        Segment path P_i at point p_{t_k}
    end if
end for

C. Determining Path Segment Similarity

SugarTrail measures the similarity of pair-wise path segments so that it can group physically proximate path segments into path-clusters. This process comprises of two steps. (1) The first step measures the similarity of pair-wise signatures in two path segments. (2) After acquiring the similarity of all the corresponding signatures in the two path segments, the system uses a modified longest common subsequence (LCS) algorithm to compute a similarity score between paths.

1) Signature-Wise Distance Metric: Based on the nature of obtained signatures, SugarTrail uses a hybrid metric to compute signature-wise distance consisting of a weighted sum of hamming and euclidean distance between signatures.

First, a mobile node may or may not be in communication range from a given anchor node depending on its location. From a signature point of view, an out-of-range or failed measurement leads to *nulls* in the RToF signature vector $\overrightarrow{R_i}$. To quantitatively measure this connectivity difference of RToF measurements between point p_i and p_j, we derive the Hamming distance

$$H_{i,j} = \frac{\overrightarrow{R_i} \oplus \overrightarrow{R_j}}{N}, \tag{2}$$

where the "$\oplus$" operation counts the number of non-null common elements in two vectors. The Hamming distance computes the similarity of visible anchors from the two location signatures.

Second, if two locations are physically close, the magnitude of RToF measurements from the same anchor node should be similar. Based on our analysis of location signatures, this translates to a small Euclidean distance between the N-dimensional RToF signatures.

The normalized Euclidean distance, given by,

$$d_{i,j} = \frac{\|\overrightarrow{R_i} - \overrightarrow{R_j}\|}{\text{\# of visible common anchors}} \tag{3}$$

measures the distance in RToF measurements from anchors visible in both signature vectors. To compensate for RToF measurement variance, in the implementation of the system we penalize faraway location signatures that yield small Euclidean distances by non-linearizing the Euclidean distance $d_{i,j}$ using a Sigmoidal transformation, i.e.

$$d'_{i,j} = \frac{1}{1 + e^{-d_{i,j}}}. \tag{4}$$

Finally, the signature-wise signature distance is computed as a weighted sum of both the Hamming and the Euclidean distances, i.e.

$$D_{i,j} = \frac{d'_{i,j} + \omega H_{i,j}}{1 + \omega}, \tag{5}$$

where ω is the weight to balance the Euclidean and the Hamming distance. During the evaluation, we empirically set ω to 0.5, such that the signature distance has a high correlation with the physical location distance.

2) Path-wise Similarity Metric: After computing the distance between signatures, the system needs to compute a similarity score for path-pairs that comprise of those signatures. This similarity score is used to group path segments belonging to the same physical pathways.

Since user path segments belonging to the same physical pathway should be clustered, we use the following criteria to design a suitable similarity metric for path segments,

- **High Point-Wise Similarity:** Users moving along the same physical pathway are likely to observe similar location signatures. This means the path segment similarity must consider the signature similarity of individual location points on the paths.
- **Length Independent:** Users moving on the same physical pathway may travel over different portions of the pathway. These partial path segments must be clustered together. In other words, the path segment similarity must be independent of the length of the path segment.
- **Robust Against Path Warping:** User movement varies in speed as well direction on a micro-scale. As the SugarTrail system does not use any inertial motion measurements, the path segment similarity should ignore warping of the path. In other words, skipping up to a specified maximum number of points to determine highest similarity should be possible.

To satisfy the above requirements, SugarTrail represents each user's path as a time-series of sequential signature and direction data and employs a modified version of the Longest Common Sequence (LCS) algorithm [2] to obtain similarity scores. Algorithm Box 2 describes the modified LCS algorithm.

Our modified LCS algorithm, uses the previously described signature distance metric for element-wise comparisons. Based on these comparisons, the algorithm increments the similarity score if the difference of the signature vectors $D_{i,j}$ and the index difference between the location points of the two paths is less than ϵ and δ, respectively.

ϵ and δ are thresholds that relax the correspondence requirement of signature vectors and point indices. ϵ is chosen through empirical measurements of RToF signature distance for points that are within $1m$ of each other. $delta$ is chosen according to the average walking speed of humans that limits the amount of warping that is permitted in matching two paths.

Algorithm 2 Modified LCS Subsequence Matching

for each pair of path segments P_m and P_n **do**
  Initialize the LCS matrix as an I-row, J-column all-zero matrix, where I and J are the numbers of location points, i.e. lengths, of P_m and P_n
  for each pair of points p_i^m and p_j^n in P_m and P_n **do**
    Find signature-wise signature distance $D_{i,j}$ between p_i^m and p_j^n using our hybrid distance metric.
    if $D_{i,j} < \epsilon$ and $|j - i| < \delta$ **then**
      Update the element at row i, column j in the LCS matrix as $L_{i,j} = e^{-D^2_{i,j}} + L_{i-1,j-1}$
    else
      $L_{i,j} = \max(L_{i-1,j}, L_{i,j-1})$
    end if
  end for
end for
Compute the LCS score for P_m and P_n as $\frac{L_{I,J}}{\min(I,J)}$

D. Multidimensional Scaling & Clustering

The system requires the clustering of similar paths as nodes of the VRM graph. Path signatures depend on distorted magnetometer readings that are not symmetrical for opposite directions. Consequently, the direction of traversing a path affects its signature and a dual graph is created. SugarTrail uses the K-means algorithm to cluster paths. Since, the paths are not points in Cartesian space, the system employs an intermediate step that uses the multi-dimensional scaling (MDS) algorithm to represent each path as a point in a n-dimensional Cartesian space [19].

Using the LCS similarity score matrix, the MDS algorithm creates a configuration of relative coordinates for the points, such that the Euclidean distances between them can approximately represent the original distance matrix. The K-means algorithm is then applied to the path coordinates to group them into path-clusters. The maximum number of clusters k is chosen based on the number of unidirectional pathways in the retail store.

E. Navigation

The SugarTrail system clusters all user collected path segments to create nodes of the VRM graph. It records the connection between individual path segments in a users path to create a connection matrix representing the edges. The system uses the VRM to route users from any designated point in the map to the desired destination.

1) Localization Over Paths: To navigate users, SugarTrail needs to determine their starting position, track whether they follow the navigation path, and note when they reach their destination. Algorithm 3 shows the steps involved in computing the alignment of two paths, where P_n is the user's current path and P_m is a recorded path with the best overall similarity score to P_n. Using the LCS path matching algorithm, the system

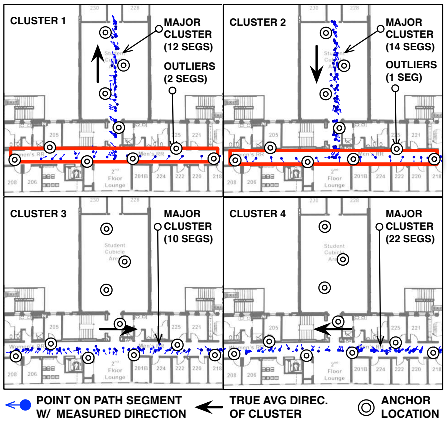

Fig. 4. Shows the clustering of path segments plotted against ground-truth locations in the T-shaped campus hallway. The segments grouped in the wrong cluster are highlighted in red.

locates the user by finding the best alignment of the current user path segment with recorded clusters of path segments. Subsequently, the system guides users to the destination by prompting them the next direction to travel in.

Algorithm 3 Modified LCS Path Alignment

```
for each path pair P_m and P_n do
  Initialize P_m's index counter to i = 1
  Initialize P_n's index counter to j = 1
  Initialize aligned path length to k = 1
  while i ≤ Length(P_m) and j ≤ Length(P_n) do
    if D_{i,j} ≤ ε and |j − i| ≤ δ then
      j = j + 1
    else if L_{i+1,j} ≥ L_{i,j+1} then
      i = i + 1
    else
      j = j + 1
    end if
  end while
end for
```

Localization of a user on a recorded path determines the accuracy of navigation. The performance of the system and it's dependence on various parameters is evaluated in the following section.

IV. Evaluation and Results

In this section, we present our evaluation methodology and experimental results. First, we validate the performance of the SugarTrail system in large-scale deployment. Second, we compare SugarTrail to an ideal system with access to accurate locations. And third, we analyze the effect of environmental conditions and amount of training data on performance.

A. Small-Scale Campus Experiment

To validate the clustering phase of SugarTrail, we created a small-scale live testbed in a single-intersection (T-shaped campus) hallway. This experiment enables us to determine if user traces can be segmented correctly and grouped into path-clusters corresponding to physical pathways.

The testbed has 10 anchors placed along the walls, distributed over the hallway, with no position measurements. The position of anchors is kept static for the period of the experiment. The mobile node in the test setup is attached to a laptop computer. The laptop computer aggregates raw compass and radio measurements, and simultaneously collects the ground truth location (recorded by clicking on a pre-loaded map of the hallway).

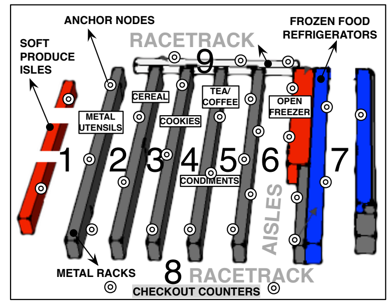

Fig. 5. Layout of the supermarket used for the experimental deployment. The open pathways are labelled by numbers 1-9.

In our preliminary tests, 3 volunteer users walked along the hallways, recording radio and magnetometer measurements at every step. In all, 20 user movement traces were obtained with 395 signatures. The obtained traces were then segmented and clustered by SugarTrail to be compared with the ground-truth paths.

The results of the segmentation and clustering are shown, superimposed on the physical map, in Figure 4. The 20 user traces were segmented into 66 path segments. The SugarTrail algorithm was able to correctly cluster the path segments into 4 uni-directional path-clusters corresponding to the 2 pathways (horizontal and vertical arms of the "T") and the 2 major directions of movement in each. 88% of path segments were clustered correctly with path segments in their ground-truth pathway and the correct direction. 3 horizontal direction segments were mis-clustered with vertical segments. This was due to the high similarity in the compass readings (indoor magnetic distortion) and small length of the path segments. 5 segments were discarded for having no dominant direction.

The small-scale testbed experiment shows that the path-cluster algorithm performs well in realistic conditions. Based on these results, we perform an end-to-end system evaluation in a large-scale supermarket scenario.

B. Large-Scale Supermarket Experiment

To test the system in a large-scale realistic application scenario, we deployed anchors in a supermarket during operational hours.

The supermarket is a challenging environment for indoor localization due to a number of factors. First, the supermarket has metallic aisles, refrigerators, and merchandise that significantly affects RF as well as magnetic characteristics. Second, the environment is dynamic due to changes in arrangement of merchandise and constant movement of people. Consequently, the experiment provides a comprehensive evaluation of the performance of the SugarTrail system at scale.

User Paths: Due to the privacy concerns of tracking real supermarket users, we modeled users based on a study of peo-

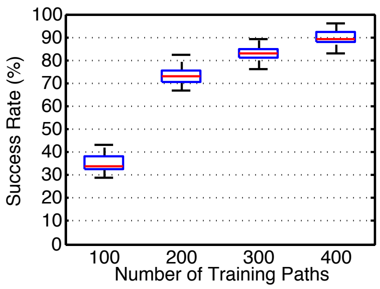

Fig. 6. The figure shows a box plot (median, 25th and 75th percentiles) of the success rate of navigation obtained by different number of paths for training data.

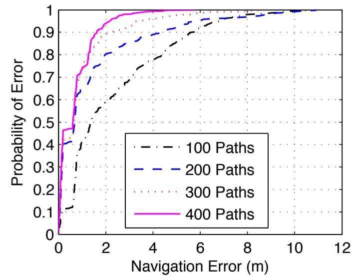

Fig. 7. Figure shows the CDF of navigation error obtained by the SugarTrail system over the entire supermarket area.

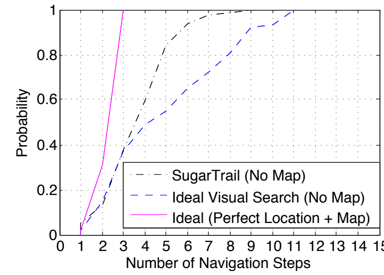

Fig. 8. The figure shows the CDF of the number of navigation steps required to navigate a user between 1000 randomly chosen start and destination point pairs.

ple movement in supermarkets [11]. The modeled users travel through the store and collect measurements corresponding to their locations.

Real Measurements: The modeled users obtain measurements from a database of previously collected high granularity sensor readings. The database of ground truth sensor readings is created by collecting real radio and magnetometer measurements over a dense square grid (each point 0.6m apart from others) in all walkable aisles in the 26m×24m supermarket. **These readings were collected in presence of shoppers and other supermarket activity**.

Having validated the ability of SugarTrail to segment and cluster actual human paths in the campus testbed, the supermarket experiment allows us to scale the system to a large number of paths without sacrificing the fidelity of the environment's RF and magnetic characteristics.

1) Experiment Setup: For our experiments, we deployed 30 stationary anchors in the 26m $\times$ 24m New Wing Yuan supermarket in Santa Clara, CA. The anchors were placed at locations on the shelves but distributed over the area of the supermarket. Figure 5 shows the layout of the supermarket and arrangement of equipment and produce.

To create a database of RToF signatures and magnetometer readings for pathways in the entire store (i.e. wherever walking is possible), with the same setup as the campus experiment, the system collected measurements on a grid of 0.6m. 1589 distinct locations were measured from the entire store. At every location, 20 distinct readings were obtained for RToF signatures as well magnetometer measurements, to capture the variance, creating a database of 31780 readings. 75% of these were used as part of the training data for SugarTrail, while 25% were used for the evaluation data set. In our experiments, we observed that signatures generally contain an average of 11 (min: 6, max: 15) anchor readings.

2) User Mobility Model: The mobility model used to generate user paths for the evaluation comprises of two components – a **macro-model**, based on a comprehensive study of over 8000 human motion paths in a supermarket by Larsen et. al. [11], which determines the macro-scale direction of users over the aisles and the perimeter of the store; and a **micro-model**, based on biased the random waypoint model (RWP), which determines the individual step size and direction when a user travels along a path. The RWP model is frequently used to simulate random human motion for Ad-hoc networks [18].

C. Performance Analysis

To exhaustively evaluate the performance of the SugarTrail system, we overlaid measured signatures in the supermarket with model generated user paths. A subset (75%) of the data is used for initial training and (25%) is used for testing.

400 user paths are used in the supermarket test environment for the learning phase of the SugarTrail system. Additional paths and signature data is used for the evaluation. This represents about half-a-day's of consumer traffic in a moderately busy supermarket and is a reasonable setup time for the SugarTrail system. Each user path traversed between 1-5 pathways in the supermarket and consisted of about 150-300 steps of varying lengths. The length of a user step was allowed to vary over the normal human stride length of 0m to 1.2m following a uniform distribution. The pause time between steps was allowed to vary from 1 sec to 2 minutes.

Every user collected RToF and magnetometer measurements every 2 seconds. The interval of 2 seconds was derived to approximately capture user steps at an average pace of $0.5m/s$, corresponding to leisurely human walking speed. However, as the stride length is variable the system does not always get measurements at every user step, as is expected in case of real human users. The algorithm makes no strict assumption about the human motion model and is therefore designed to be robust to such variations.

The user path training data set is provided as input to the SugarTrail algorithm, which in turn builds the corresponding VRM. A **new set of users (different from the training data set)** is used to query the SugarTrail system and follow navigation directions between randomly selected start and destination points. The performance of the system is evaluated based on the success rate of the navigation, the accuracy of navigation, and the number of steps in the navigation route. These metrics are further defined in the following subsections.

1) Metric – Navigation Success Rate: The SugarTrail system guides users by providing a route consisting of a sequence of path-clusters, and a set of direction prompts (based on the expected user's local raw compass reference reading) to navigate along those clusters. We define the navigation success rate as the percentage of times a user is guided correctly from a start point to the proximity of user's desired destination point.

Specifically, an instance of user navigation is designated as a success if the following conditions are satisfied – First, the user is guided into the correct physical pathways as the desired destination point. Second, the user is not prompted to take an invalid direction (obstructed path) at any point along the route. Finally, the user is not guided into any incorrect pathways along the route. Such unexpected pathways may arise due to inconsistencies in the map or inaccuracy of localization.

To evaluate the navigation success rate of the SugarTrail system we performed 20 trials, each containing 100 user navigation requests separate from the training data set. In each trial, we note the number of times the navigation was successful according to the above mentioned criterion. Figure 6 shows a box plot of the success rate of navigation obtained with varying sizes of training data. The success rate of navigation improves as the size of training data increases over time, with the system achieving 90% success rate with 400 user path training data set. This corresponds to about half-a-day's worth of data.

The trend in success rates can be explained by the fact that the SugarTrail system relies on the record of peoples' prior movement to build the navigable virtual roadmap. Consequently, the larger the number of user paths recorded the better is the expected success rate of finding a route to the destination.

2) Metric - Navigation Accuracy: We measure the fine-grained ability of the SugarTrail system to guide a user to an exact destination point (specified within the VRM). We define this as navigation accuracy, which is the distance from the actual physical coordinates of users' desired destination to the location that the SugarTrail system can guide them to.

The navigation accuracy is obtained by computing the location error of the destination for 1000 user queries. Figure 7 depicts the CDF of navigation error obtained by the SugarTrail system over the entire supermarket area. As seen from the figure, the SugarTrail system can provide a navigation accuracy better than $1.3m$, 80% of the times, with a median accuracy of under $0.6m$. Figure 7 also shows that as the amount of training data increases, the navigation accuracy also improves.

3) Comparing Efficiency of Navigation: The SugarTrail system provides a set of direction prompts to guide the user to the destination via intermediate pathways (where user changes direction). In the indoor supermarket scenario, where length of any single pathway is relatively small, the efficiency of the system depends on the number of direction prompts or *navigation steps*. The number of navigation steps can be used to compare the performance of the SugarTrail system versus other possible approaches to navigating indoors. We compare the number of navigation steps required by SugarTrail to the following approaches:

- **Ideal Oracle System:** The ideal system has knowledge of the accurate location of the user as well as access to the map of the environment. This system provides the shortest, and most efficient path possible.
- **Visual Search:** The alternative to using SugarTrail, in indoor environments without access to location or maps, is to perform a visual sequential search of the environment. For example, a human user seeking a certain product in a supermarket, can sequentially search the aisles. We assume the user does not overlook items.

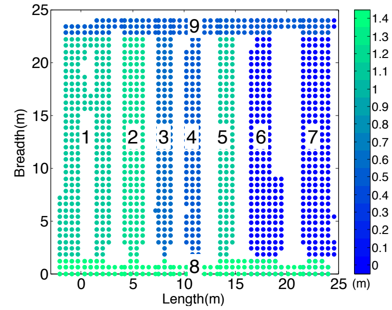

Fig. 9. The figure shows the 80th percentile navigation accuracy over different regions of the supermarket computed using a 400 path training dataset.

Figure 8 shows the CDF of number of navigation steps required to navigate a user between 1000 randomly chosen start and destination point pairs using different approaches. The figure shows that difference in navigation steps between SugarTrail and the ideal benchmark system is under 2 steps for 80 percent of the cases as compared to over 5 steps for a visual search of the environment. Thus, the SugarTrail system provides significantly more efficient navigation performance with minimal setup and maintenance effort.

4) Effect of Environment: The layout and material-content of the environment affect the nature of RToF signatures as well as magnetometer measurements. The SugarTrail algorithm is designed to minimize the effect of such variations, however, the performance of the system is naturally affected by environmental factors. We analyze the distribution of navigation accuracy obtained over different geographical regions of the supermarket to gain insight into the environmental effects.

Figure 9 shows the 80th percentile navigation accuracy over different regions of the supermarket. The accuracy is computed over 1000 destinations for each aisle of the supermarket. The figure shows that SugarTrail reduces the error attributed to environmental variations with accuracy better than 1.5 meters throughout the supermarket. The navigation accuracy is highest towards the wider refrigerator aisles (aisle 6 and 7). While, narrow aisles with metal utensils as well as many obstacles (aisles 1,2 and 3) show lower navigation accuracy.

V. Related Work

Indoor navigation has been an active area of research in the past, primarily for robotics and increasingly for mobile computing application. Systems using different modalities such as radio frequency (RF), infrared (IR) and vision have been proposed to locate users inside building and consequently provide navigation guidance.

Map-based with Dedicated Infrastructure: Map-based systems require a physical schematic map of the environment and the location of the people to navigate. Several systems that deploy a specialized localization infrastructure have been developed for indoor environments, including infrared [20], [21], ultrasound [16], and more recently RFIDs [14]. Apart from needing maps, these systems require careful positioning and measurement when placing beacons and requires a much larger deployment and maintenance effort when compared to SugarTrail.

Map-based with RF-surveys: Another research thrust has been focused on utilizing RF signals (such as WiFi access points, or Bluetooth) in office environments to provide localization. These systems fingerprint each indoor location using received signal strength (RSS) measurements from multiple access points. A location database of such fingerprints is created through manual surveys, active tagging, or user corrections [1], [3], [4], [9], [15], [23]. Every collected fingerprint must be tagged with a location and the required active user attention has been the major drawback of such approaches. Furthermore, Wi-Fi localization based navigation is generally courser-grained that required for retail store (e.g. supermarkets) applications. SugarTrail, on the other hand, does not require a pre-existing map and only requires annotation of interesting clusters of signatures for navigation, not every collected signature.

Map-based with Inertial Measurements: Another approach to localization has utilized inertial measurements with sensors attached to people to augment WiFi measurements [10], [22]. However, they depend on the existence of WiFi localization systems for initial estimate and is not fine-grained enough for aisle differentiation in stores. Furthermore, unlike SugarTrail, this approach requires known accurate maps and would be negatively affected by indoor magnetic interference.

Mapless with Advanced Sensing/Processing: Navigation approaches that do not need a preexisting map of the environment to operate have been actively explored for robot path finding applications. A large body of work on simultaneous localization and mapping (SLAM) [7], [12] algorithms allows a robot to determine its location and navigate without pre-existing maps. Recent work on WiFi-SLAM [8] uses a mobile robot to build a map in terms of RSS fingerprints. However, most SLAM implementations [6] are suited for robots with sensing (vision, laser range-finders, stereo-cameras, accurate inertial sensors etc.) capabilities unavailable on mobile devices.

VI. Conclusion

The paper presents the SugarTrail system that enables accurate indoor navigation leveraging the structured movement pathways in retail environments. By using the sequence of measurements as combined signatures (paths), the system learns traversable path-cluster to build a navigable virtual roadmap of the environment. This method when applied to a supermarket environment was able to achieve a success rate of $> 85\%$ and an average accuracy of more than $0.7m$ without the need for existing maps and extensive active radio-signature surveys. The system requires no active user participation other than carrying a mobile radio node and has no initial manual calibration requirements.

VII. Acknowledgements

Special thanks to Kaifei Chen for help with collecting experimental data and the anonymous reviewers for their insightful and constructive comments. This research was supported by Intel, NSF under grant 1121690-CNS1135874-NSF-ZHANG and DARPA under grant 1080247-D11AP00265-DOI-ZHANG.

References

[1] P. Bahl and V. Padmanabhan. RADAR: an in-building RF-based user location and tracking system. In *Proceedings of IEEE INFOCOM 2000.*, volume 2, pages 775–784. IEEE, 2000.

[2] L. Bergroth, H. Hakonen, and T. Raita. A survey of longest common subsequence algorithms. *Proceedings of SPIRE 2000.*, pages 39–48, 2000.

[3] P. Bolliger. Redpin - adaptive, zero-configuration indoor localization through user collaboration. In *Proceedings of MELT 2008*, page 55, New York, New York, USA, Sept. 2008. ACM Press.

[4] K. Chintalapudi, A. Padmanabha Iyer, and V. N. Padmanabhan. Indoor localization without the pain. In *Proceedings of Mobicom '10*, page 173, New York, New York, USA, Sept. 2010. ACM Press.

[5] J. Chung, M. Donahoe, C. Schmandt, I.-J. Kim, P. Razavai, and M. Wiseman. Indoor location sensing using geo-magnetism. In *Proceedings of the 9th international conference on Mobile systems, applications, and services*, MobiSys '11, pages 141–154, New York, NY, USA, 2011. ACM.

[6] A. J. Davision and D. W.Murray. Simultaneous Localisation and Map-Building Using Active Vision. *Transections on Pattern Analysis and Machine Intelligence*, 24(7):865–880, 2002.

[7] H. Durrant-Whyte and T. Bailey. Simultaneous localization and mapping: part I. *IEEE Robotics & Automation Magazine*, 13(2):99–110, June 2006.

[8] B. Ferris, D. Fox, and N. Lawrence. Wifi-slam using gaussian process latent variable models. In *Proceedings of IJCAI'07*, pages 2480–2485, San Francisco, CA, USA, 2007. Morgan Kaufmann Publishers Inc.

[9] A. Haeberlen, E. Flannery, A. M. Ladd, A. Rudys, D. S. Wallach, and L. E. Kavraki. Practical robust localization over large-scale 802.11 wireless networks. In *Proceedings of MobiCom '04*, page 70, New York, New York, USA, Sept. 2004. ACM Press.

[10] L. Klingbeil and T. Wark. A wireless sensor network for real-time indoor localisation and motion monitoring. In *Proceedings of IPSN '08*, pages 39–50. IEEE Computer Society Washington, DC, USA, 2008.

[11] J. S. Larson, E. T. Bradlow, and P. S. Fader. An exploratory look at supermarket shopping paths. *International Journal of Research in Marketing*, 22(4):395–414, 12 2005.

[12] J. Leonard and H. Durrant-Whyte. Simultaneous map building and localization for an autonomous mobile robot. In *Proceedings of IROS '91*, pages 1442–1447. IEEE.

[13] Nanotron Tecnologies Gmbh. Real Time Location Systems (RTLS). 2007.

[14] L. M. Ni, Y. Liu, Y. C. Lau, and A. P. Patil. LANDMARC: Indoor Location Sensing Using Active RFID. *Wireless Networks*, 10(6):701–710, Nov. 2004.

[15] J. Park, B. Charrow, D. Curtis, J. Battat, E. Minkov, J. Hicks, S. Teller, and J. Ledlie. Growing an organic indoor location system. *Proceedings of Mobisys '10*, pages 271–284, 2010.

[16] N. Priyantha, A. Chakraborty, and H. Balakrishnan. The cricket location-support system. In *Proceedings of Mobicom '00*, volume 2000, pages 32–43. ACM New York, NY, USA, 2000.

[17] A. Purohit, Z. Sun, F. Mokaya, and P. Zhang. Sensorfly: Controlled-mobile sensing platform for indoor emergency response applications. In *Proceedings of IPSN '11*, pages 223 –234, april 2011.

[18] G. Resta and P. Santi. An analysis of the node spatial distribution of the random waypoint mobility model for ad hoc networks. In *Proceedings of POMC '02*, pages 44–50, New York, NY, USA, 2002. ACM.

[19] S. Schiffman, M. L. Reynolds, and F. W. Young. *Introduction to Multidimensional Scaling: Theory, Methods, and Applications*. Emerald Group Publishing Limited, 1981.

[20] R. Want, A. Hopper, V. Falcão, and J. Gibbons. The active badge location system. *ACM Transactions on Information Systems*, 10(1):91–102, Jan. 1992.

[21] A. Ward, A. Jones, and A. Hopper. A New Location Technique for the Active Office. *IEEE Personal Communications*, 4:42–47, 1997.

[22] O. Woodman and R. Harle. Pedestrian localisation for indoor environments. In *Proceedings of Ubicomp '08*, page 114, New York, New York, USA, Sept. 2008. ACM Press.

[23] M. Youssef and A. Agrawala. The Horus WLAN location determination system. In *Proceedings of MobiSys '05*, page 205, New York, New York, USA, June 2005. ACM Press.