

Design and Evaluation of a Wireless Magnetic-based Proximity Detection Platform for Indoor Applications

Xiaofan Jiang, Chieh-Jan Mike Liang, Kaifei Chen, Ben Zhang, Jeff Hsu
Jie Liu, Bin Cao, and Feng Zhao

Microsoft Research Asia
Beijing, P.R. China, 100080

{fxjiang, liang.mike, v-kaiche, v-benzh, v-jehs, liuj, bincao, zhao}@microsoft.com

ABSTRACT

Many indoor sensing applications leverage knowledge of relative proximity among physical objects and humans, such as the notion of “within arm’s reach”. In this paper, we quantify this notion using “proximity zone”, and propose a methodology that empirically and systematically compare the proximity zones created by various wireless technologies. We find that existing technologies such as 802.15.4, Bluetooth Low Energy (BLE), and RFID fall short on metrics such as boundary sharpness, robustness against interference, and obstacle penetration. We then present the design and evaluation of a wireless proximity detection platform based on magnetic induction - LiveSynergy. LiveSynergy provides sweet spot for indoor applications that require reliable and precise proximity detection. Finally, we present the design and evaluation of an end-to-end system, deployed inside a large food court to offer context-aware and personalized advertisements and diet suggestions at a per-counter granularity.

Categories and Subject Descriptors

B.0 [Hardware]: General; B.4 [Hardware]: Input/Output & Data Communication; H.4.m [Information Systems Applications]: Miscellaneous

General Terms

Design, Experimentation, Measurement

Keywords

Localization, Magneto-Inductive, Tracking, Virtual Zone

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IPSN’12, April 16–20, 2012, Beijing, China.

Copyright 2012 ACM 978-1-4503-1227-1/12/04 ...\$10.00.

1. INTRODUCTION

Low-power wireless technologies are often used for both communication and proximity sensing. Many applications rely on short-range radio to restrict which receivers can receive their transmissions. This property is particularly useful in *person-scope* applications, where human wearable devices can discover and interact with other devices embedded in the environment without explicit human interventions.

To make these applications intuitive to human users, the discovered objects in the environment must be within the personal interaction sphere, i.e., “within arm’s reach”. For example, a computer may automatically wake up when a user sits in front of the desk. But, it should not react to people walking past the cubicle. A department store may tag a cloth rack to beacon the sizes and material information of clothes hanging on the rack. This information only needs to reach people who can see and touch these clothes. A home appliance (e.g., refrigerator or microwave) may change its user interface when a person stands in front of it. All these applications share the same architecture: wireless beacons are attached to stationary objects, and mobile nodes carried by human communicate with them when in their proximity. To make these applications robust, beacon nodes must create predictable and stable coverage zones for mobile nodes to discover. In fact, fuzzy zone boundaries introduce uncertainty to the user location, and can lower the user experience with false positives, such as in the case of targeted advertisement. Or, it can lower security such as in the case of building door access.

Many typical low power communication technologies, such as Bluetooth 4.0 (a.k.a. Bluetooth Low Energy or BLE) and ZigBee (with 802.15.4 as the physical layer) have difficulties maintaining robust communication zones. The 2.4 GHz band is prone to external interferences. Human bodies, Wi-Fi activities, and metal objects nearby can significantly alter the signal propagation patterns and the receiving noise floor. While many previous research studies these effects from the perspective of point-to-point communication link quality, such as bit-error rates and data goodput [8, 19], little has been done to quantify the impact on proximity sensing at the human interaction scale. For example, humans react to messages in the order of seconds. As long as a node can be discovered within a few seconds, there is no perceivable difference for the applications. On the other hand, human interactions are less tolerable to obstructions, such as human body or any object blocking the view, and the reach between senders and receivers.

One of the key challenges for characterizing and comparing across wireless proximity technologies is to define comparison metrics. Practically, one can sample the waveform propagation field with only discrete receivers. In this paper, we first propose a methodology to reconstruct spatial and temporal patterns from wireless

receivers. Then, we apply a machine learning technique based on support vector machines (SVM) to infer parametric models from discrete samples. With model parameters defining the *white*, *grey*, and *black* zones, the comparison is resilient to temporary noise and disturbance.

Using this methodology, we empirically characterize and systematically compare the *proximity zones* across several technologies: BLE, 802.15.4, and 900 MHz RFID reader. Results show that none of them can produce robust proximity zones for person-scope applications. We then investigate an alternative wireless beacon based on magnetic-inductive (MI) coupling. Humans, furniture, and most construction materials do not affect magnetic propagation pattern. Thus, MI proximity zones are much more intuitive and practical for human interactions than those using existing technologies.

MI coupling components used to suffer from size, cost, and power inefficiency. However, recent development in the hardware, especially passive keyless entry (PKE) in the automotive industry, has made it feasible to integrate MI communication (MIC) onto motesized devices. To this end, we design and implement LiveSynergy, a MIC-capable platform with both MI transmitter and receiver. We evaluate the proximity zone generated by this new kind of wireless communication. The comparison to other alternatives shows significant advantages in zone consistency over time, boundary sharpness, and robustness to human obstruction. We also discuss our experience from deploying a LiveSynergy-enabled system in a cafeteria to deliver personalized nutrition and fitness information to diners at each food counter.

In summary, this paper makes the following contributions:

- We propose methodologies that empirically and systematically compare the proximity zones created by various wireless technologies.
- We design, implement, and evaluate a magnetic-induction based wireless proximity sensing platform - LiveSynergy.
- With real data traces, we show that MI solution exhibits significantly more robust zones than BLE, 802.15.4 and RFID alternatives.
- We share the experience from deploying LiveSynergy in an real-world application.

The rest of this paper is organized as follows: in Section 2, we survey related work in wireless proximity sensing. In Section 3, we give an empirical definition of proximity zones and how to compare them. After comparing the proximity zone for RF-based technologies in Section 4, we motivate and describe our design of LiveSynergy (Section 5). We evaluate LiveSynergy’s proximity zone properties in Section 6 and demonstrate a real-world application deployment at cafeteria in Section 7.

2. RELATED WORK

Proximity detection can be achieved either directly by proximity sensing, or indirectly by inferring their absolute positions. Many proximity sensing are initiated by humans, for example, using RFID scanning [5], near field communication (NFC) touching, 1-D or 2-D barcode scanning, or optical sensing such as bokode [12]. These types of technologies require explicit actions from human, adding a level of inconvenience. Proximity sensing can also be initiated by the environment, typically through computer vision techniques using conventional cameras [2, 20] or Kinect-like depth cameras.

Indirect sensing technologies are hard to achieve proximity accuracy. Positioning services such as GPS are not feasible indoors as

they require line-of-sight to satellites. Ultrasound-based localizations [16, 6] do not penetrate solid objects. Infrared-based systems, such as ActiveBadge [21], have difficulty localizing objects with fluorescent lighting or direct sunlight, and have an effective range on the order of meters, offering poor granularity for small spaces. Inertia sensors based approaches, such as double-integration of accelerometer data, are often inaccurate due to large DC/offset errors, and require frequent recalibration [3]. Wi-Fi fingerprinting typically requires a time-consuming training phase and costly recalibration [23, 1]. Wi-Fi also suffers from multi-path effect and Raleigh fading, leading to inaccurate location estimates [18]. Using pressure floor sensors for localization do not scale to large systems [13] and they do not work for lightweight objects such as those on the table.

These technology are either too fine grained, requiring explicit user actions for discovery, or too coarse grained to pinpoint objects or people within the vicinity of a few meters. As a result, current applications have to make compromises, such as requiring a person to actively swipe a security card when entering a building, or inaccurately localizing people based on association to the same Wi-Fi access point.

With increasing attention in this domain, several technologies have emerged as leading choices for proximity detection indoors. Bluetooth Low Energy (BLE) is a part of Bluetooth 4.0 specification, targeting directly at low energy application. While it suffers from the same interference problem as other technologies on the 2.4 GHz band, it has the unique advantage of being standard on newer mobile phones. Long-range RFID is another technology designed for locating objects. While it has the advantage of using passive tags, it is easily affected by obstructions, and often requires heavy post-processing for outlier reduction [22]. Devices based on 802.15.4 (PHY and MAC layers of the Zigbee specification) have also been used extensively, particularly in the sensornet community. In this paper, we evaluate these three technologies with respect to our application requirements.

LiveSynergy leverages dynamic magnetic fields for proximity sensing, which inductively couple the beacon and the receiver, alleviating much of the interference issues. The theory of operation for magnetic tracking has been explored initially by Raab et. al. [17], and has been commercialized for high-end motion tracking applications [15]. These commercial devices are typically expensive and not optimized for proximity detections at scale. In [9], Markham et. al. used magneto-induction for tracking underground beavers. Their system uses a relatively large antenna ($5m \times 1m$), which is not practical in indoor scenarios, but provided the inspiration for our work. We compare our platform against 802.15.4, BLE, and long-range RFID, and evaluate it under a real application deployment.

3. PROXIMITY ZONE

The desired length of “an arm’s reach” is application dependent. For example, in the scenario of automatic computer-login, the desirable range of detection is perhaps the size of the cubicle or the immediate space in front of the monitor; in case of targeted advertising at the shopping mall, the desired detection range could be the two-meter space in front of a particular brand of merchandise. While the detection range is moderate in these applications, the *sharpness* and *consistency* of the detection boundary are of great importance, and directly impact the application performance and user experience. For example, in the first scenario, if the region of detection has a large “grey” area where detections happen sporadically and unpredictably, the computer may not identify the correct user in a timely manner.

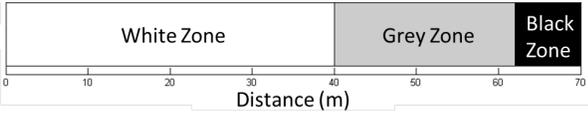


Figure 1: Example proximity zone from single dimensional dataset.

In this paper, we refer to this detection space as the **proximity zone**. Being able to accurately specify the “proximity zone” of an object or space, and reliably tell when someone enters and exits this zone are paramount for these applications. In the rest of this section, we first list a set of evaluation metrics for the proximity zones common to this class of person-scope applications. Based on these metrics, we propose an empirical definition of “proximity zone”, which allows us to better evaluate and compare various technologies.

3.1 Evaluation Metrics

The following are the metrics for evaluating proximity zones:

- **Boundary sharpness:** The boundary of the proximity zone should be as binary as possible - either in or out. The region of sporadic detections, or the “grey” zone, should be minimal.
- **Boundary consistency:** Proximity detection should be consistent over time. This requires the technology to be robust against temporal interferences such as external noises and environmental variations.
- **Obstacle penetration:** Since both the beaconing node and the listening node can be mobile, e.g., being carried by humans, the radio technology should be robust against obstructions such as the human body. Specifically, the radio reception should not change significantly regardless locations on the body.
- **Additional metrics:** A number of additional properties also impact their real-world applicability, such as the *range* and the *geometric shape* of the zones, the *beaconing frequency* achievable by the transmitters, the *power consumption* of the devices, the *form-factor* of the mobile tag, and the *cost* of the overall system.

3.2 Empirical Definition

Since waveform propagation from beacons is a physical phenomenon, it is more useful to characterize the zones from the receiver’s point of view. We propose an *empirical* definition for **proximity zones**, inferred over the packet reception ratio and the transmission distance. Unlike ranging, where RSSI or LQI is sometimes used to infer distance information, in the case of proximity detection, the event of successfully receiving a packet is the primary indicator of whether the receiver is within the proximity zone of a beacon¹.

While this definition relies on empirical data collected at the receiver, and is only a sampling of the space, it provides a practical metric for evaluation against real application requirements.

Conceptually, a **proximity region** consists of three adjacent zones: the *white zone*, the *grey zone*, and the *black zone*. It is also equivalently defined by two boundaries: the *white/grey boundary* and the *grey/black boundary*. Figure 1 shows the one-dimensional view of

¹However, we do use RSSI in deciding the primary zone in the case of overlapping zones (c.f. Section 7)

a proximity zone based on measurement data along a single axis. With additional data, the same methodologies described below can work on two-dimensional and three-dimensional proximity zones. After empirically determining proximity zones, we can then evaluate the above metrics of boundary sharpness, consistency and so on.

The classification and evaluation procedure starts from collecting packet reception statistics at various distances from the transmitter. Then, these statistics allow us to classify and determine whether each location falls in the white, grey, or black zone. Finally, we apply a machine learning technique based on support vector machines (SVM) to infer the continuous coverage of the zones from these discrete samples.

3.2.1 Classification of Points

Assume a beacon broadcasts at a fixed frequency f , for a total of α packets from time t to time t' . Let:

P = a point in space at a distance of (p_x, p_y, p_z)
from the beacon

W = detection window

PRR_τ = packet reception rate measured between τ and $\tau + W$
where $\tau \in [t, t' - W]$

We first define the indicator function $I(\tau)$ as:

$$I(\tau) = \begin{cases} 0, & \text{if } PRR_\tau \leq \epsilon \\ 1, & \text{if } PRR_\tau > \epsilon \end{cases}$$

where ϵ is the PRR threshold

Then,

$Color(P, t, t') = \text{white},$

if $\prod_{\tau} I(\tau) = 1 : \forall \tau \in [t, t' - W]$

$Color(P, t, t') = \text{grey},$

if $\sum_{\tau} I(\tau) \geq 1$ and $\prod_{\tau} I(\tau) = 0 : \forall \tau \in [t, t' - W]$

$Color(P, t, t') = \text{black},$

if $\sum_{\tau} I(\tau) = 0 : \forall \tau \in [t, t' - W]$

Here P represents the location of the receiver node. W defines the sliding window size in computing all possible PRRs over the duration of the dataset at P . We collectively term these sliding-window PRR values as $WPRR$. $I(\tau)$ labels each sliding window into 1 if the corresponding PRR is above an application specific threshold ϵ , or 0 if otherwise. Using this indicator function, we can label P as “white” if all sliding windows are 1, “grey” if some but not all are 1, and “black” if all are 0. This definition allows us to label points around a beacon into $\{\text{white}, \text{grey}, \text{black}\}$, according to how likely beacons are detected by a receiver at that point over some period of time $[t, t']$.

We note that $WPRR$ is fundamentally different from computing a single PRR over the entire dataset at P and then evaluating whether $PRR_p > \gamma$. Specifically, looking at a single PRR over the entire dataset does not consider the uniformity of successful packet receptions over time. Most applications can tolerate losing a few packets within a given time window, as long as the receiver hears at least one beacon in that time window.

Additionally, if ϵ is set to 0, it is equivalent to stating $I(\tau) = 1$ if the time between consecutive packet receptions is less than

the detection window W , which is the minimum requirement for detection to happen. In the context of the sliding window, having $\epsilon = 0$ equates to having the inter-packet arrival time to be always smaller than the detection windows W . In this specific form, one could build a histogram of the inter-packet arrival times and see if any falls beyond the detection window W . However, to allow applications to specify the degree of likelihood of detection, we leave ϵ as a user definable parameter.

3.2.2 Classification of Zones

Given a set of points P and their colors, we can then find the boundaries that divide the proximity regions into three zones of *white*, *grey*, and *black* colors. There are multiple approaches to grouping points into zones. Due to the nondeterministic nature of the boundaries of the proximity zone, we take statistical learning approaches to find these boundaries. We could consider the boundary detection problem as a classification problem. We first define two boundaries.

- The *white/grey boundary* separates the white points $\{P \mid \text{Color}(P, t, t') = \text{white}\}$ from the grey points $\{P \mid \text{Color}(P, t, t') = \text{grey}\}$. We use notation $f_{w/g}$ as the classifier that $f_{w/g}(\mathbf{x}) > 0$ for any white point \mathbf{x} and $f_{w/g}(\mathbf{x}') < 0$ for any grey point \mathbf{x}' . We use equation $f_{w/g}(\mathbf{x}) = 0$ to represent the decision boundary.
- The *grey/black boundary* separates the grey points $\{P \mid \text{Color}(P, t, t') = \text{grey}\}$ from the black points $\{P \mid \text{Color}(P, t, t') = \text{black}\}$. Similarly, we use notation $f_{g/b}$ as the classifier that $f_{g/b}(\mathbf{x}) > 0$ for any grey point \mathbf{x} and $f_{g/b}(\mathbf{x}') < 0$ for any black point \mathbf{x}' . Equation $f_{g/b}(\mathbf{x}) = 0$ represents the decision boundary.

Based on the boundaries, we infer the three proximity zones as below.

- Grey zone: $\{\mathbf{x} \mid f_{w/g}(\mathbf{x}) > 0 \text{ and } f_{g/b}(\mathbf{x}) < 0\}$. Between the white/grey boundary and grey/black boundary.
- White zone: $\{\mathbf{x} \mid f_{w/g}(\mathbf{x}) < 0\}$. Inside the white/grey boundary is the white zone.
- Black zone: $\{\mathbf{x} \mid f_{g/b}(\mathbf{x}) > 0\}$. Outside the grey/black boundary is the black zone.

We use support vector machines (SVM) as the classifier to find the boundaries of the zones in this paper. SVM seeks maximum-margin hyperplane to separate two classes. The objective of classical SVM is shown in Equation 1.

$$\begin{aligned} \mathbf{w} = \arg \min \|\mathbf{w}\|^2 \\ \text{s.t. } y_i(\mathbf{w}^T \mathbf{x}_i - b) \geq 1, \text{ for all } i \end{aligned} \quad (1)$$

where \mathbf{w} and b are the parameters to define the hyperplane to separate the two classes.

To accommodate application-specific requirements, we introduce two user-definable parameters in our classification model.

One parameter is the *error tolerance*, which indicates how clean the zones should be. There could be multiple boundary candidates sometimes. For example, we could choose one smooth boundary with some misclassified points. At the same time, we could also have a non-smooth boundary without any misclassified point. In terms of learning, it is a tradeoff between training loss and regularization. We could specify the cost parameter C for misclassified example to choose the best decision boundary for the specific scenario.

Another parameter is the *strictness* of a zone. We may expect the white zone, as well as the black zone, to contain no grey points. We should notice that although strictness is related to error tolerance, however, strictness is non-symmetry. For example, while we require no grey point be misclassified to white or black, we allow white or black points to be misclassified as grey. In terms of the classification problem, strictness could be controlled by the difference of costs for false positive and false negative. We could set the cost parameter (or the cost of false positive) C and the cost for false negative C' to be different. Formally, given C , we could define the strictness parameter γ to be $\gamma = 1 - \frac{C'}{C}$.

For our problem, we adapt the classical SVM model. By slackening on the hard constraints for error tolerance and adding the cost parameters, we formulate the problem as,

$$\begin{aligned} \mathbf{w} = \arg \min_f \{ \|\mathbf{w}\|^2 + C \left(\sum_{y_i=1}^l (1 - y_i f(\mathbf{x}_i))_+ \right) \\ + C' \left(\sum_{y_i=-1}^l (1 - y_i f(\mathbf{x}_i))_+ \right) \} \end{aligned} \quad (2)$$

where $(z)_+ = \max(z, 0)$. Parameter C and C' are the costs for misclassification of false positive and false negative, respectively.

Since our zone boundary might not be linear, we apply the kernel trick in the SVM problem to obtain nonlinear boundaries. We use the RBF kernel as the kernel function. In this case, the decision boundary could be specified as $\{\mathbf{x} \mid f_{w/g} \equiv \langle \phi(\mathbf{x}), \mathbf{w} \rangle = 0\}$ for the classifier $f_{w/g}$, where $\phi(x)$ is the feature mapping function for the RBF kernel. Similarly, for the classifier $f_{g/b}$, the decision boundary is $\{\mathbf{x} \mid f_{g/b} \equiv \langle \phi(\mathbf{x}), \mathbf{w} \rangle = 0\}$. With the boundaries defined, we could compute the sizes of each zone for our evaluation purpose. Because the black zone is not closed, it can not be computed. In fact, the zone of the greatest interest is the grey zone. However, computing the area of grey zone is non-trivial since its boundaries might be irregular. Section 6.4 presents the heuristics we use to compute the area.

The SVM also provides information on how well boundaries fit the training data in terms of *loss*. Therefore, we could further capture the *fitness* of proximity zones by the classification accuracy.

In summary, user can specify the following input parameters to the classification model:

- **Error tolerance:** Corresponds to the cost parameter C in SVM, which specifies a trade-off between smoothness of the boundary and mis-classification.
- **Strictness γ :** Computes $C' = C \cdot \gamma + C$ to specify the application's tolerance to false positives (i.e., sporadic detections occurring outside the nominal zone range) and false negatives (i.e., missing detections inside the nominal zone range). For example, for applications that cannot tolerate any missing detection or false detections, γ can be set to 1 to coerce the model to include all grey points into the grey zone.

Together with the learning model in (2), we could obtain the boundaries, as well as the following metrics,

- **Size** of the white and grey zone, which can be computed numerically based on the boundaries.
- **Boundary sharpness**, defined as the area ratio of the white zone to the sum of white zone and grey zone, $\frac{\text{size}(\text{white})}{\text{size}(\text{white}) + \text{size}(\text{grey})}$.
- **Fitness** of proximity zones, which provides a measure of how well the zone boundaries fit the data, or a confidence measure of the proximity zone classification.

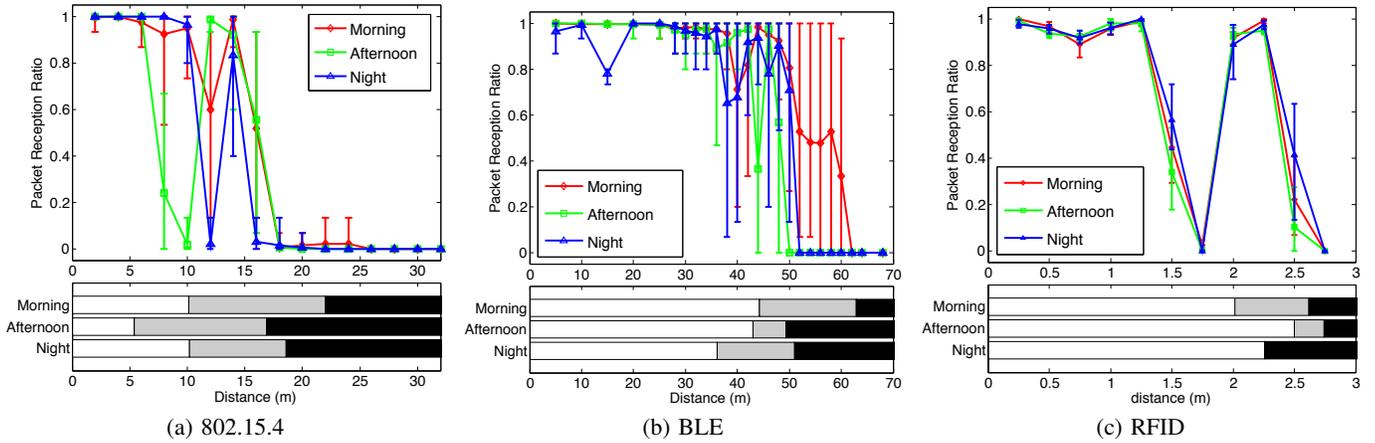


Figure 2: (top) the WPRR at different distances, and (bottom) zone boundaries, for the corresponding radio

4. EVALUATION OF EXISTING TECHNOLOGIES

Using the procedure for determining proximity zones in Section 3, we evaluate some of the common radios for beaconing: 802.15.4, Bluetooth Low Energy (BLE), and RFID.

4.1 Boundary Sharpness and Consistency

To characterize grey zones of each radio, we collected packet reception data at different positions from the transmitter as outlined below. We primarily focus on single dimensional data here since it can clearly illustrate our results. We defer the analysis of data from two dimensions to later sections.

Experiment methodology: The hardware setup consists of a pair of TI CC2540 BLE dev boards (transmitting on 2.4 GHz at 0 dBm), a pair of TelosB motes with 802.15.4-compliant TI CC24240 radio (transmitting on 2.4 GHz at 0 dBm), and a Impinj Speedway R1000 RFID reader (transmitting on 902 MHz at 8 dBm).

Experiments take place on a floor inside a typical office building, and the floor contains cubicles separated by semi-metallic walls and metallic over-head shelves. For each experiment, we position the receiver at various locations along a line from the beacon on a small hallway. The distance intervals are adaptive depending on the range of the technology. At each location, the packet reception data is collected over a period of 200 seconds. WPRR is then computed using a windows size of 3 seconds and an ϵ of 0. Using the algorithm described in Section 3, each point is then assigned a color from {white, grey, black}, and the white/grey and grey/black boundaries are computed for each experiment with a strictness parameter of 0.99.

Finally, to capture the time-variant nature of the radio due to environment changes, we repeated the experiments three times throughout the day: morning, afternoon, and night.

Results: Figure 2 shows the WPRR at each sampling location and zone boundaries for all three radio technologies at different times of the day.

From the PRR vs. distance figures on top of Figure 2, we observe that the average WPRR for 802.15.4 fluctuates significantly over a large distance between 5 and 20 meters, suggesting that the boundary is not sharp spatially. The confidence interval also varies with distances, suggesting that packet reception changes significantly over both space and time. In comparison, average WPRR for BLE is a little better than 802.15.4 spatially, leading to slightly

shorter grey zones. However, BLE has a relatively large confidence interval outside the white zone, signifying that the packet reception is not uniform over time. This suggests that BLE is more susceptible to bursty packet loss, in which case BLE requires a larger window, or human wait time, for the proximity detection to be reliable in the grey zone. However, inside the white zone, BLE has a smaller confidence interval than 802.15.4, possibly because BLE employs forward error correction rather than error detection. RFID, in comparison, has much smaller confidence intervals, which suggests that packet reception is consistent over time, and results in smaller grey zones. An interesting observation is the dip at around 1.75 m. Further experiments with different heights at the same distance reveal that this dip is isolated at that height, and is therefore an effect of the physical environment such as the multi-path effect. We summarize the boundary sharpnesses of all three in the table at the end of this subsection.

While we can see that RFID is more consistent than 802.15.4 and BLE over short periods of time from their respective WPRR confidence intervals, to better understand each radio’s tolerance to environmental changes, we consider data across three different times of the day. From Figure 2, we can see that both 802.15.4 and BLE have large variances in WPRR curves and zone boundaries, while WPRR curves for RFID almost overlap. An explanation is that both 802.15.4 and BLE occupy the unlicensed 2.4 GHz band shared with many office equipments, such as Wi-Fi, whereas RFID operates on the fairly quiet 902 MHz band. We further quantify consistency in terms of the standard deviations across zone boundaries between the three trials:

$$RMS(\sigma(\text{border}_{white/grey}), \sigma(\text{border}_{grey/black}))$$

A smaller RMS value is preferable. We summarize consistencies in the table below.

	802.15.4	BLE	RFID
Boundary sharpness	0.32	0.70	0.77
Boundary consistency	0.12	0.13	0.13

4.2 Human Obstacle Penetration

Previous studies have shown that human body can significantly affect the packet reception of 802.15.4 [11]. This section extends the experiments to include BLE and RFID.

Experiment Methodology: The user carries the receiver in the right pants pocket, which is one of the most common locations to carry mobile devices. At each distance from the transmitter, we

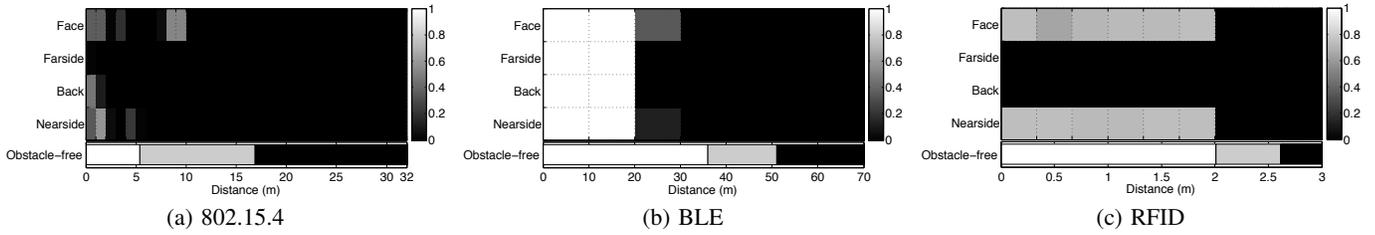


Figure 3: Packet reception ratio when the user, with the receiver in the right pants pocket, is in different body orientations

calculate PRR from 500 packets as the user changes the body orientation by 90° each round.

Results: In Figure 3, y-axis corresponds to different orientations of the user with respect to the beacon. “Face” means the user is facing the beacon at 0° ; “farside” means the user is -90° with respect to the original orientation, standing between the receiver and the beacon; “back” means the user is at 180° with his back facing the beacon; and “nearside” sideways at 90° . x-axis is distance from the beacon, and color of the block represents PRR (white=1). Proximity zones from the same time period is added as a reference. This figure shows that human body has a significant impact on all three radio technologies. PRR for 802.15.4 is almost zero beyond one or two meters, regardless of orientation. BLE has PRR of 1 in the initial 20 meters, but drops sharply to zero for the rest 16 meters in the original white zone, also independent of orientation. PRR for RFID is zero for all distance at the “farside” and “back” orientations. The other two orientation are slightly better with 50% receptions. This result suggests that none of these three technologies are suitable for proximity detection of humans.

4.3 Additional Metrics

Signal propagation and geometry: Beaconing platforms based on 802.15.4 and BLE radios usually use omni-directional antennas, with typical transmission range of about 100 m. However, as previously shown, their signals attenuate with amoeba-like propagation patterns [10], which is not desirable for applications requiring a consistent boundary.

Off-the-shelf RFID readers are usually coupled with directional antennas, with a range in the neighborhood of meters. Compared to 802.15.4 and BLE, the radio propagation attenuates more evenly in different directions. RFID antennas usually have a radiation angle less than 180 degrees.

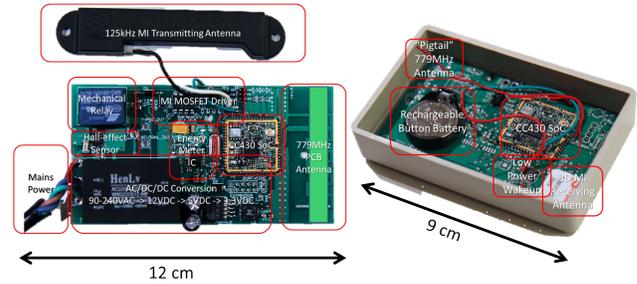
Form Factor and Costs: Form factor and costs play a huge role in realizing ubiquitous deployments. While RFID can produce a more consistent and smaller grey zone, 802.15.4 and BLE have advantages in both form factor and costs. Specifically, long-range RFID readers have a relatively large antenna due to the lower radio frequency, but the cost of long-range RFID readers range from hundreds to thousands of dollars, such as Impinj Speedway R1000.

5. LIVESYNERGY PLATFORM

Section 4 shows that existing beconing technologies are not ideal for reliable and precise proximity detection. Fundamental problems such as external interferences motivate us to explore different beconing hardware technology, rather than software tricks. In this section, we first discuss magnetic induction, then present a platform we developed for proximity detection - LiveSynergy.

5.1 Magnetic Induction Communication

Magnetic induction communication (MIC) has a different elec-



Magnetic Transmitter: Pulse generates the dynamic magnetic field in three stages. First, the MCU encodes the payload data and toggles its IO pins according to the bits. Then, a driver circuit amplifies the toggling pattern to provide sufficient power. Finally, the antenna transmits the signal as magnetic pulses in the air. The rest of the discussion focuses on each of these stages.

The magnetic transmitter software stack on MCU largely determines the data transmission rate and reliability through the choice of modulation scheme and maximum transmission unit (MTU). Preliminary results showed that simple duration-based on/off keying modulation does not achieve high reliability. We opted Manchester encoding, as the self-clocking nature reduces the chance of several types of data failure. In addition, preliminary results suggested an MTU larger than 16 bytes can significantly impact the PRR. Finally, the bit signal duration can impact the reliability or the transmission rate if the duration is too short or long respectively. In our system, a data rate of 2730 bps provides the sweet spot.

The I/O pin from the MCU, signaling at 3.3V, connects to the input of a buffer powered at 12VDC; the output of the buffer connects to the input of an inverter, also powered at 12VDC. The two ends of the MI antenna are connected to the outputs of both the buffer and the inverter. This design effectively doubles the voltage swing across the antenna. In addition, power-limiting resistors are put in series with the antenna to artificially decrease the range, if needed.

The antenna size and geometry partially determine the transmission range. In [9], a relatively large antenna ($5m \times 1m$) is used to obtain a range of 5m. However, such a large antenna is not practical in indoor scenarios. Pulse has a PCB mountable antenna specifically designed for PKE applications, with a dimension of $8cm \times 1.5cm$.

Microcontroller and Radio: To mitigate interference and noise issues that often exist in modern office buildings, and to ensure adequate RF range, we opted the sub-1 GHz band instead of the popular 2.4 GHz spectrum. Our radio stack conforms to 802.15.4c, a 779 Mhz PHY layer amendment to the 802.15.4 standard.

To simplify our RF design phase and ensure optimal RF performance, we use the SuRF core module from PeoplePower [14], which integrates TI CC430 and RF matching network (balun). CC430 is a SoC with MSP430 MCU (with 32 KB ROM and 4 KB RAM) and CC1101 low-power radio chipset. We designed a compact 1 dB monopole PCB antenna based on the TI reference design [4].

5.2.2 Link Receiver

Link (c.f. right of Figure 4) is a mobile device carried by humans to receive and decode the IDs from dynamic magnetic fields transmitted by Pulses. Link is battery-powered inside a $9.2cm \times 5.8cm \times 2.3cm$ enclosure. The three primary hardware components include MCU and radio, 3D magnetic coil, and wake up chip. Since Link and Pulse share the same MCU and radio, the discussion below focuses on the latter two.

The 3D receiver coil is tuned to 125 kHz, and sensitive in all three dimensions providing spatial freedom to the human carrying the Link. Since magnetic field is not significantly attenuated by non-ferrous materials, the signal reception is similar regardless of Link’s location on the body. The outputs from the coil are small electrical signals in x , y , and z dimensions, and they are connected to a low-power wake up chip, AS3932. AS3932 reduces analog circuitry with an integrated programmable gain amplifier (PGA), offloads computation by digitizing the analog magnetic signals, and reduces MCU power consumption by firing a wakeup signal only when a magnetic field with valid preamble is decoded.

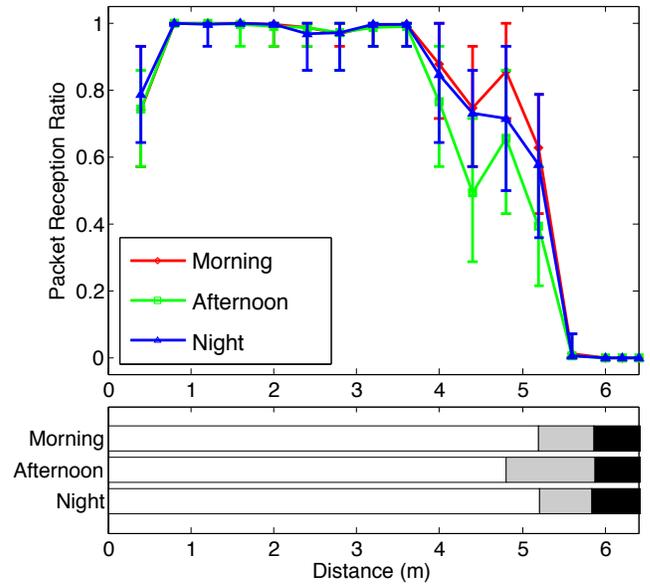


Figure 5: WPRR vs. distance for morning, afternoon, and night, and associated 1D proximity zones.

Upon being waken up by the AS3932 chip, the MCU uses its internal clock to periodically capture the value of the AS3932 output pins. After the software stack successfully verifies the 16-bit CRC, it delivers the payload data to the application.

Finally, we note that the magnetometer inside existing mobile phones use Hall elements and can theoretically detecting dynamic magnetic fields (in addition to static fields). Unfortunately, detailed inspection of popular magnetometer chips reveal that dynamic magnetic field is filtered out in hardware to reduce “noise” from the earth magnetic field. It is our hope that sensor manufacturers can lift this limit and enable MIC on future mobile phones.

6. EVALUATION OF LIVESYNERGY

The evaluation methodology is similar to that used in Section 4. We place a pair of Pulse and Link to perform the test at three different times of the day. In each trial, the Pulse is fixed in position and transmits an ID using MI at 5Hz. We position the Link at different distances from the Pulse, with variable intervals to cover more points inside and near the grey region. For each distance, Pulse transmits 1,000 packets; the Link records successful receptions, together with timestamps and RSSI. In addition, for every position, we rotate Pulse with respect to the Link, at angles from 0° to 180° in 30° increments. While this is still a small subset of all the degrees of freedom between a pair of antennas, this at least enables us to create proximity zones in both 1D and 2D space.

6.1 Boundary Sharpness and Consistency

The top part of Figure 5 shows the WPRR vs. distance graph for MI in a single dimension, collected at three different times of the day. In comparison to 802.15.4 (Figure 2(a)) and BLE (Figure 2(b)), packet reception ratio of MI is significantly better – it is more uniform, as indicated by shorter confidence intervals; and it is more consistent over time, as indicated by the similar average WPRR readings across three trials. In comparison to RFID (Figure 2(c)), the variance of WPRR readings is about the same, but MI exhibits better signal decay characteristics than RFID. There are still some fluctuations in MI’s WPRR curve since ferrous materials

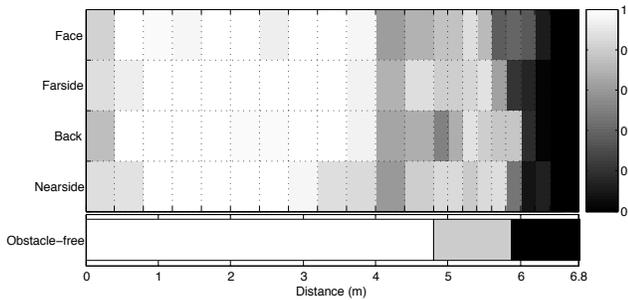


Figure 6: Body orientation vs. distance

in the environment do affect magnetic propagation, but much less severe than RFID. The bottom part of Figure 5 shows the proximity zone for MI, computed using our definition in Section 3. This figure shows that the zone boundary for MI is sharp in all three time periods of the day, with a boundary sharpness of 0.82, which is much better than 802.15.4, BLE, and RFID. From this result, we can conclude that LiveSynergy, using magnetic induction as the beaconing signal, is able to generate proximity zones with the most consistent and the sharpest boundaries among current technologies studied in this paper.

To better differentiate MI from RFID, we compare results in two dimensions, as described in more detail in Section 6.4.

6.2 Human Obstacle Penetration

Figure 6 shows the PRR vs. distance with the Link being carried in the right pant pocket at different orientations from the Pulse. This figure shows that human body has very little impact on the MI signal propagation. In contrast, 802.15.4, BLE, and RFID are easily affected by human bodies and other obstacles, as seen in Figure 3. This is because those technologies operate at much higher frequencies than MI, and also because MI is inductively coupled. This property makes MI ideal for applications where a mobile “tag” needs to be carried by a human. On the hand, MI signal is blocked by ferrous objects, which may be a problem for some applications, but desired for some others (e.g., applications using ferrous materials as natural zonal boundaries).

6.3 Additional Metrics

Geometry: From Figure 7, we can observe that the proximity zone of Pulse in two dimensions extends to all directions, covering all 360° . This same observation also applies to 3D.

Range: The maximum range (i.e., radius) is around 5m, but it can be artificially decreased by limiting the antenna output power. This range is desirable for a range of indoor applications as described in Section 1.

Beacon rate, power, size, and cost: We summarize these metrics in the table in Section 6.5.

6.4 Comparisons to RFID in 2D

Because long-range RFID exhibits the sharpest boundary among all previous technologies, as seen in Section 4, we compare MI with RFID in more detail here. Using packet reception data measured in the x-y plane, we were able to find the zone boundaries in 2D by applying our methodologies in Section 3. As Figure 7 shows, in two dimensions, the white/grey and grey/black boundaries are curves, and zones are represented as areas.

From the MI proximity zone on the left, we can see that the MI’s white zone is relatively circular, centered at the MI transmitter, with a small grey zone surrounding it. In contrast, RFID has a much nar-

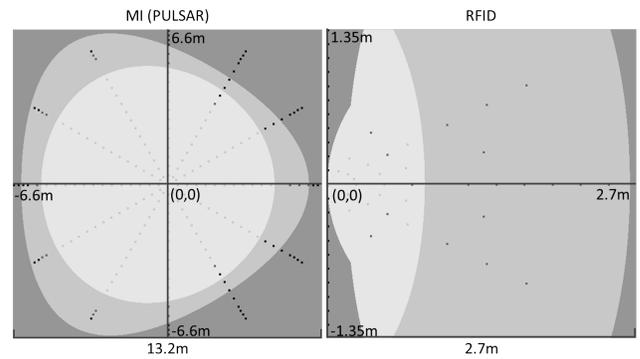


Figure 7: Proximity zones for MI and RFID in 2D space

rower proximity zone with an angle less than 180° . This geometry is not as desirable for applications that require omni-directional coverage. The grey zone of RFID is also much bigger than that of MI. We used Monte Carlo method to estimate the areas of white and grey zones for both MI and RFID, and found that the boundary sharpness ratio for MI to be 0.70, and for RFID to be 0.25. This shows that MI has a much sharper boundary than RFID. Furthermore, the range of MI is about 5m omni-directionally where RFID is only about 1m in the right half plane. Combined with the advantages MI has over RFID in terms of consistency and body penetration, MI is the more viable proximity detection solution than RFID for many indoor applications.

6.5 Summary

Summary of evaluation results for LiveSynergy/MI:

Boundary sharpness	0.82
Consistency	0.03
Penetration	Excellent
Geometry	Omnidirectional
Range	5m
Max beacon rate	50Hz
Power	19.75mA
Tag form factor	9cm x 6cm x 2cm
Cost (tx, rx)	(\$50, \$30)/pc @ 100 units

7. APPLICATION DEPLOYMENT

Results from Section 6 show that magnetic induction communication (MIC) is a viable solution for reliable and precise proximity detection. This section discusses a real-world deployment at a large cafeteria to provide personalized advertisements and diet suggestions. The discussion starts by giving the deployment overview, discussing the end-to-end system, and then presenting deployment results and experiences.

7.1 Deployment Overview

Providing personalized shopping experience has been a recurring application in the ubiquitous computing community. On one hand, shops can deliver targeted advertisements and coupons; on the other hand, shoppers can receive suggestions based on their shopping history and preferences. We highlight the potential of MIC communication in such applications by instrumenting the cafeteria inside a large company with Pulses, as seen in Figure 8. Diners carry Links, which are associated to their mobile phones running our mobile ap-



Figure 8: A Pulse is installed at the Teppanyaki counter, projecting a proximity zone of 5m radius. Links carried by diners in their pockets detect this zone as they enter.

plication. Food counters in the cafeteria are organized by types of food offered. The proximity readings are maintained at a back-end server, and used to deliver real-time targeted advertisements and coupons for the particular food counter that the customer is near. The entire user experience depends on performance of the LiveSynergy system in detecting diners, and how targeted advertisements and coupons are delivered.

Figure 10 shows Pulses at five different locations in the cafeteria: the entrance, the “Japanese” counter, the “A La Carte” counter, the “Teppanyaki” counter, and the “Vegetarian” counter. Pulses beacon twice per second, while Links report MIC beacons heard once every second. Aggregating beacon reports help lower the wireless network load, and help resolving cases where beacons from multiple Pulses are heard. Specifically, Link reports the Pulse ID with the highest signal strength in the current detection time window.

The back-end server performs a simple time-window threshold filtering on the location readings to differentiate between passing-by and staying. Using this data, the server can push context-aware information to the customer in a timely manner, through a native application running on the mobile phone.

7.2 System Architecture

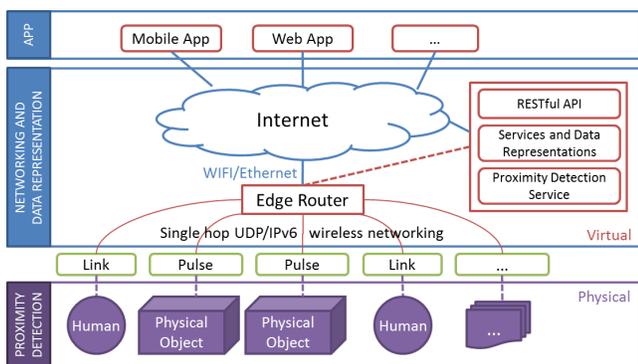


Figure 9: End-to-end LiveSynergy system architecture

Our system architecture (c.f. Figure 9) includes three main components: the LiveSynergy platform for detecting when humans enter and exit cafeteria counters; a networking and data representa-

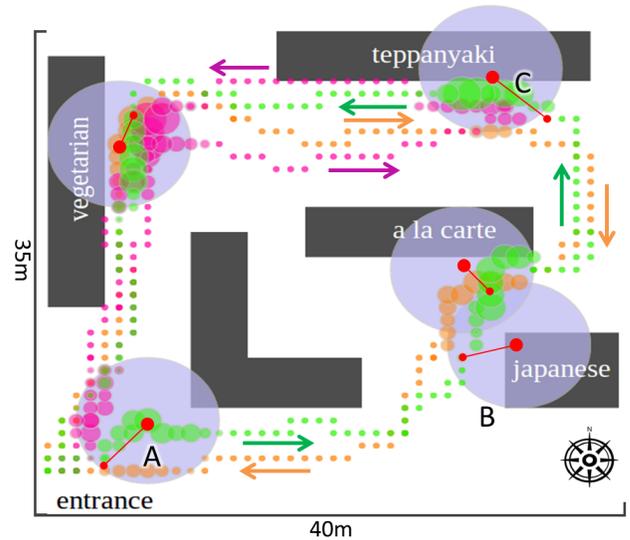


Figure 10: Movement trails from three customers in the cafeteria deployment.

tion layer enabling communication among counters, services in the cloud, and app clients; and an application layer consisting of a mobile application for providing feedback and visualization.

Networking: Our network needs to be *reliable* since interactions between humans and physical objects rely on reliable data exchange, and it should be *low latency* to support timely feedback to users. We designed our network to meet these requirements by using single-hop at the wireless sensor layer, reliable proxy gateway, and REST over IPv6. Nodes (Pulses and Links) directly connect to the edge router via single-hop, and communicate using UDP over IPv6. Applications, as well as multiple services hosted in the cloud, communicate with the edge router using HTTP over Ethernet.

Data Representation and Web Services: To enable interoperability between applications and services provide by our system, we designed a simple message format based on JSON, which is common in RESTful web services. Similar to XML, JSON is text-based, human-readable, self-describing, and language independent. However, compared to XML, JSON is light-weight with a smaller grammar and simpler data structure.

7.3 Deployment Results

Figure 10 shows the floor-map of the cafeteria, labeled with both the locations of each food counter, and the movement trails of the three diners in our study. Diners enter the cafeteria from the entrance at the lower left corner at different times. Each diner takes a different route and visits various food counters on the way. The trails are indicated by three different colors, and direction is indicated by arrows of the same color; purple-shaded circles represent proximity zones of Pulses (approximated to perfect circles in this figure), with radius equaling the mean of the actual ranges of Pulses. As our ground truth, we recorded a video as the customers walk around the cafeteria purchasing food. The video is time-stamped so that we can correlate events in the footage to the actual proximity detection data we collected.

Figure 10 also overlays our deployment data over the floor-map. Circles of varying sizes indicate locations of detections, using data reported by Links as they intersect the proximity zones of Pulses. The size is proportional to the RSSI recorded by the Link, and red lines connecting green circles to the center of the purple circles

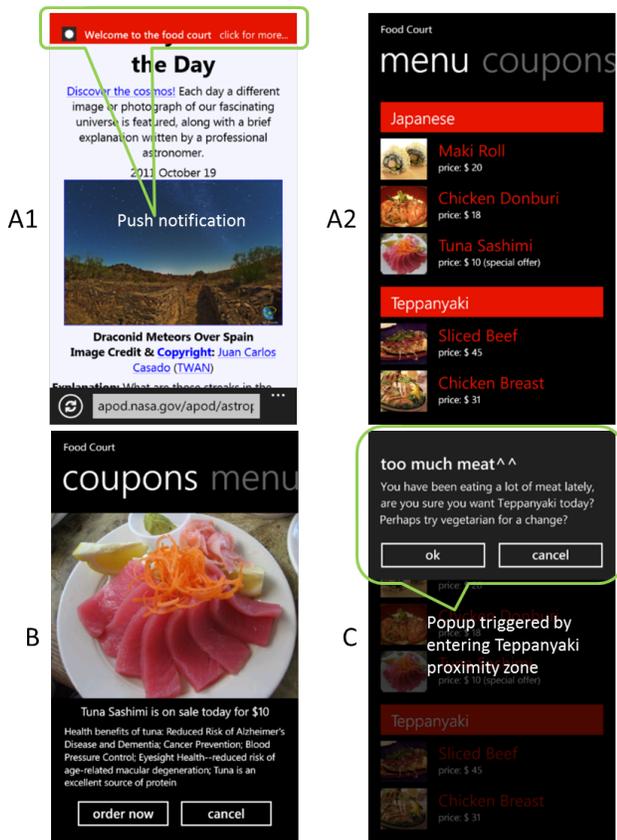


Figure 11: Windows Phone 7.5 application

indicate which Pulse that particular Link detects (i.e. being associated with). To avoid clustering the figure, we only draw red lines for the first detection after the “green” customer enters a new Pulse zone.

From this figure, one can observe a few key results. First, the boundaries of the Pulse proximity zones are sharp, with no false positives or negatives. The detections occur consecutively within each proximity zones projected by Pulses. There is a little delay between entering of a zone and the time of first detection, as seen by the slight offset between the edge of the purple circle and the first green circle. This delay is also observed at the exit of the zone. This is expected since Links run a simple window-based filtering before reporting the detection. This delay is well within the tolerance. Second, the detections are consistent for all three customers as they enter and exit the five zones, as evident from the detection locations for all three customers – the three different colors of detections circles all exist inside the same purple circles (+ the detection delay). Additionally, we observe that RSSI is indeed proportional to the distance between the Link and the Pulse. This allows the system to correctly resolve overlapping zones by “associating” to the closest proximity zone. For example, the “green” customer is first associated with the “Japanese counter”, then switched (correctly) to the “A La Carte counter” in the overlapping area of both zones.

All three diners installed our native Windows Phone 7.5 application. Figure 11 shows “green” customer’s screenshots as he walks past various food counters, as labeled by A, B, and C. First, after his Link detects the Pulse ID of the entrance and forward to the server, our web service sends a notification message to his mobile phone. Then, the phone displays a notification at the top of the

screen without interrupting the foreground application. Clicking the notification opens the main screen of our mobile application with the menu² and promotional items. At the “Japanese counter”, the web service pushes counter-specific coupons (tuna sashimi). Then, at the meat-heavy “Teppanyaki counter”, the diner receives a diet alert of excessive meat consumption based on his purchase history, with the advice of vegetarian dishes instead.

This simple mobile application demonstrates that the LiveSynergy proximity detection platform enables useful context-aware applications, without the need for heavy post-processing. We plan to build a personal energy footprint application and a cardless security entry system based on our LiveSynergy platform in the near future.

8. CONCLUSIONS

In this paper, we study the proximity zones established by wireless beacons in person-scale applications. We first propose a methodology based on sampling and classification techniques that enable us to compute zone boundaries and other metrics such as boundary sharpness, consistency, and body penetration. Using this methodology, we empirically evaluate three prominent technologies: 802.15.4, BLE, and 900MHz RFID, using data collected over a large period of time in both x and y dimensions. We show that 802.15.4 and BLE have large grey regions and are inconsistent over time. While RFID has sharp boundaries, it suffers heavily from attenuation by obstacles such as the human body. To overcome these shortcomings, we describe the design, implementation, and evaluation of a magnetic-induction based proximity sensing platform – LiveSynergy, which creates sharp and consistent boundaries that is not affected by human obstacles. These characteristics make MI significantly more robust than BLE, 802.15.4, and RFID alternative. Finally, through a real-world deployment, we demonstrate that LiveSynergy is able to successful support human interacting with a smart environment. As future work, we plan to explore the possibility of integrating MI into regular mobile devices and expanding its applications.

9. ACKNOWLEDGMENTS

We want to thank Caiquan Liu for his contributions in the early phase of the project, and Andrew Markham for his helpful feedback. And we would like to thank our shepherd, Anthony Rowe, and the anonymous reviewers for their feedback and insightful comments.

²Dish items in A2 are pre-populated and are not the actual menu. We plan to synchronize the menu with the cafeteria in the future.

10. REFERENCES

- [1] P. Bahl and V. Padmanabhan. Radar: an in-building rf-based user location and tracking system. In *INFOCOM*, volume 2, pages 775–784 vol.2, 2000.
- [2] B. Brumitt, B. Meyers, J. Krumm, A. Kern, and S. A. Shafer. Easyliving: Technologies for intelligent environments. In *HUC*, 2000.
- [3] I. Constandache, X. Bao, M. Azizyan, and R. R. Choudhury. Did you see bob?: human localization using mobile phones. In *Proceedings of the sixteenth annual international conference on Mobile computing and networking*, MobiCom '10, pages 149–160, New York, NY, USA, 2010. ACM.
- [4] Fredrik Kervel. DN023: 868 MHz, 915 MHz and 955 MHz Inverted F Antenna. <http://focus.ti.com/lit/an/swra228b/swra228b.pdf>.
- [5] D. H and D. Fox. Mapping and localization with rfid technology. In *International Conference on Robotics and Automation*, 2003.
- [6] A. Harter, A. Hopper, P. Steggles, A. Ward, and P. Webster. The anatomy of a context-aware application. *Wireless Networks*, 8:187–197, 2002. 10.1023/A:1013767926256.
- [7] X. Jiang, S. Dawson-Haggerty, P. Dutta, and D. Culler. Design and Implementation of a High-Fidelity AC Metering Network. In *Proc. IPSN/SPOTS*, 2009.
- [8] C.-J. M. Liang, B. Priyantha, J. Liu, and A. Terzis. Surviving Wi-Fi Interference in Low Power ZigBee Networks. In *Proc. SenSys*, 2010.
- [9] A. Markham, N. Trigoni, and S. Ellwood. Revealing the Hidden Lives of Underground Animals with Magneto-Inductive Tracking. In *Proc. SenSys*, 2010.
- [10] Matthew M. Holland and Ryan G. Aures and Wendi B. Heinzelman. Experimental Investigation of Radio Performance in Wireless Sensor Networks. In *SECON*, 2006.
- [11] E. Miluzzo, X. Zheng, K. Fodor, and A. T. Campbell. Radio characterization of 802.15.4 and its impact on the design of mobile sensor networks. In *EWSN*, 2008.
- [12] A. Mohan, G. Woo, S. Hiura, Q. Smithwick, and R. Raskar. Bokode: imperceptible visual tags for camera based interaction from a distance. *ACM Transactions on Graphics*, 28, 2009.
- [13] R. J. Orr and G. D. Abowd. The smart floor: a mechanism for natural user identification and tracking. In *CHI '00*, CHI EA '00, pages 275–276, New York, NY, USA, 2000. ACM.
- [14] People Power. People Power | Helping Save the Green. <http://www.peoplepowerco.com>.
- [15] Polhemus. Polhemus, Innovation in Motion. <http://http://polhemus.com/>.
- [16] N. B. Priyantha, A. Chakraborty, and H. Balakrishnan. The cricket location-support system. In *MobiCom*, 2000.
- [17] F. Raab, E. Blood, T. Steiner, and H. Jones. Magnetic position and orientation tracking system. In *IEEE Transactions on Aerospace and Electronic Systems*, 1979.
- [18] B. Sklar. Rayleigh fading channels in mobile digital communication systems .i. characterization. *Communications Magazine, IEEE*, 35(7):90–100, jul 1997.
- [19] K. Srinivasan, M. A. Kazandjieva, S. Agarwal, and P. Levis. The b-factor: measuring wireless link burstiness. In *SenSys*, 2008.
- [20] T. Teixeira, D. Jung, and A. Savvides. Tasking networked cctv cameras and mobile phones to identify and localize multiple people. In *Ubiquitous Computing/Handheld and Ubiquitous Computing*, pages 213–222, 2010.
- [21] R. Want, A. Hopper, V. Falcão, and J. Gibbons. The active badge location system. *ACM Trans. Inf. Syst.*, 10:91–102, January 1992.
- [22] E. Welbourne, K. Koscher, E. Soroush, M. Balazinska, and G. Borriello. Longitudinal study of a building-scale rfid ecosystem. In *MobiSys '09*, pages 69–82, New York, NY, USA, 2009. ACM.
- [23] M. Youssef and A. Agrawala. The horus wlan location determination system. In *MobiSys*, 2005.