PowerLine Positioning: A Practical Sub-Room-Level Indoor Location System for Domestic Use

Shwetak N. Patel[1], Khai N. Truong[2], and Gregory D. Abowd[1]

[1] College of Computing & GVU Center
Georgia Institute of Technology
801 Atlantic Drive, Atlanta GA 30332-0280 USA
{shwetak, abowd}@cc.gatech.edu
[2] Department of Computer Science
University of Toronto
10 King's College Road, Toronto ON M5S 3G4 Canada
khai@cs.toronto.edu

Abstract. Using existing communications infrastructure, such as 802.11 and GSM, researchers have demonstrated effective indoor localization. Inspired by these previous approaches, and recognizing some limitations of relying on infrastructure users do not control, we present an indoor location system that uses an even more ubiquitous domestic infrastructure—the residential powerline. PowerLine Positioning (PLP) is an inexpensive technique that uses fingerprinting of multiple tones transmitted along the powerline to achieve sub-room-level localization. We describe the basics behind PLP and demonstrate how it compares favorably to other fingerprinting techniques.

1 Introduction

Recent advances in indoor location systems use existing wireless communication infrastructure (*e.g.*, 802.11 and GSM) to provide a value-added location service. The major advantage to these approaches is that a consumer does not have to purchase any specialized equipment and can still benefit from location-aware computing. Leveraging public infrastructure has many advantages, but one major drawback is that users have very little control of the infrastructure itself. Service providers adjust the operational parameters of WiFi access points and cellular towers with little warning. These changes require recalibration of the location system. An alternative is to introduce new infrastructure in the home by distributing many low-cost, short-range beacons. The time required for installation and the possible impact to home aesthetics, however, may limit adoption.

Inspired by this strategy of leveraging existing infrastructure, and recognizing that there are drawbacks to relying on infrastructure not controlled by an individual household, we were motivated to invent a solution for indoor localization that would work in nearly every household. This paper presents such a solution, with the significant insight being the use of the residential powerline as the signaling infrastructure. We describe the first example of an affordable, whole-house indoor localization system that works in the vast majority of households, scales cost-effectively to support the tracking

P. Dourish and A. Friday (Eds.): Ubicomp 2006, LNCS 4206, pp. 441 – 458, 2006.

of multiple objects simultaneously, and does not require the installation of any new infrastructure. The solution requires the installation of two small plug-in modules at the extreme ends of the home. These modules inject a low frequency, attenuated signal throughout the electrical system of the home. Simple receivers, or positioning tags, listen for these signals and wirelessly transmit their positioning readings back to the environment. This solution, henceforth referred to as PowerLine Positioning, or PLP, provides sub-room-level positioning for multiple regions of a room and tracks multiple tags simultaneously. PLP has a localization accuracy of 87–95% for classifying regions at 3 meters and 67% at 1 meter resolution. We have installed and tested this system in a variety of homes and compare the performance against previous 802.11 and GSM solutions.

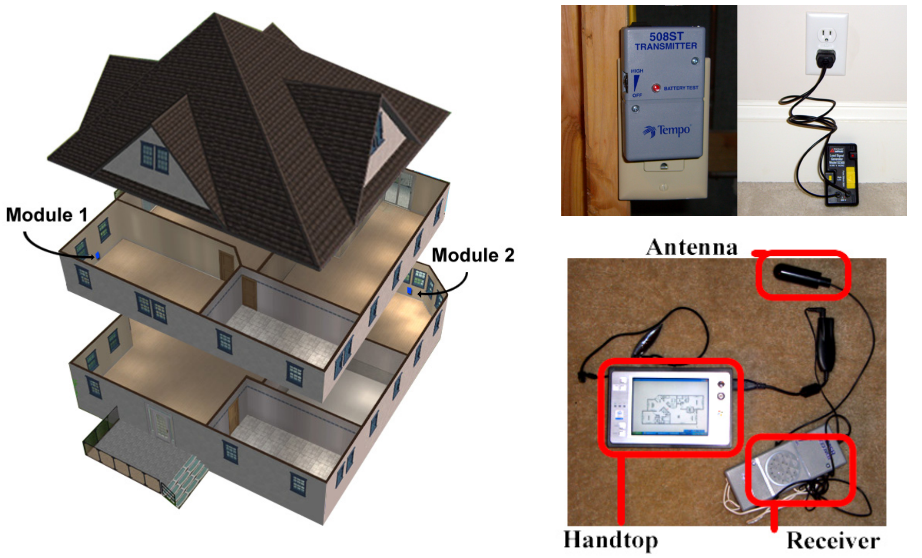

Fig. 1. Left: Placement of two signal-generating modules at extreme ends of a house. Right: The PLP system components. The top shows two examples of off-the-shelf, plug-in tone generator modules. The bottom shows a working prototype of the location tag, consisting of a receiver and antenna hooked to a handtop computer for analysis.

2 Related Work

Indoor positioning has been a very active research problem in the ubicomp community in the preceding half decade [5]. Several characteristics distinguish the different solutions, such as the underlying signaling technology (*e.g.*, IR, RF, load sensing, computer vision or audition), line-of-sight requirements, accuracy, and cost of scaling the solution over space and over number of items. Although we do not intend to provide a complete survey of this topic, we highlight projects with characteristics most relevant to the motivation for powerline positioning, namely the requirements for additional infrastructure and algorithmic approach.

The earliest indoor solutions introduced new infrastructure to support localization [1, 12, 14, 16]. Despite some success, as indicated by commercialized products [15], the cost and effort of installation is a major drawback to wide-scale deployment, particularly in domestic settings. Thus, new projects in location-based systems research reuse existing infrastructure to ease the burden of deployment and lower the cost. The earliest demonstrations use 802.11 access points [2, 3, 8], and more recent examples explore Bluetooth [9] and wireless telephony infrastructure, such as GSM [13] or FM transmission towers [7]. Concerns about system resolution eliminate the FM solution for domestic use. Another concern we highlighted in the introduction is that individuals and households may not be able to control the characteristics of these infrastructures, resulting in the need to recalibrate should parameters change. The desire to control the infrastructure and to scale inexpensively to track a large number of objects inspired the search for a solution like the powerline system presented here.

Traditional wireless signal triangulation, such as that using 802.11 access points, uses Received Signal Strength Indicator (RSSI) information to estimate distance and determine a location based on its geometric calculations Other techniques include the use of Time of Arrival, as in the case of ultrasound, or Angle of Arrival, such as with Ultra-wideband positioning [15]. Ultrasonic solutions, such as Cricket [14] and Active Bat [1], provide precise centimeter resolution, but require line-of-sight operation indoors and thus require considerable sensor installations for full coverage. Technologies, such as 802.11 triangulation, avoid issues of occlusion but suffer from multipath problems caused by reflections in the environment.

Fingerprinting of received signals can help overcome the multipath problem. Fingerprinting improves upon other means of estimation by taking into account the effects that buildings, solid objects, or people may have on a wireless or RF signal, such as reflection and attenuation. Fingerprinting works by recording the characteristics of wireless signals at a given position and later inferring that position when the same signature is seen again. A survey of signals over some space allow for the creation of a map that can be used to relate a signal fingerprint to a location.

Our location system relies on the space's powerline infrastructure. Powerlines are already in place in most homes and the power network reaches more homes than either cable systems or telephone lines. Thus, for many years, people have been using powerlines in homes to deliver more than just electricity. Several home technologies use the powerline for communications and control. The most popular example is the X10 control protocol for home automation, a standard that is more than 30 years old and is a very popular, low-cost solution for homeowners. Over the past decade, there have been a number of efforts to produce powerline communications capabilities, driven by industrial consortia like HomePlug Powerline Alliance [6] and efforts such as Broadband over Powerline (BPL).

3 System Overview

In this section, we present the theory behind the operation of PLP, discuss the two-phase localization algorithm based on signal fingerprinting, and describe the details of our prototype system that we used to evaluate the operation of PLP in real homes.

3.1 Theory of Operation

We developed the PLP system based on a popular wire-finding technique employed by many electricians and utility workers to locate or trace hidden wires behind a wall or underground. In this technique, an electrician connects an exposed end of the wire to a tone generator, which can range from 10–500 kHz, and locates the hidden wire using a handheld, inductive tone detector. Some detectors use LEDs to indicate the tone strength and others play an audible sound. In either case, the electrician scans the area for the loudest tone, indicating the approximate location of the wire. Following the presence of the tone reveals the path of the wire.

We use the following properties of the wire-finding technique to produce a viable solution for a location system:

- it is easy and inexpensive to propagate a signal or tone throughout the entire electrical system in a home without any electrical interference;
- it is possible to set the power of the tone so that it attenuates as it reaches the periphery of the home, and the electrical wiring appears in varying densities throughout the home, creating a time-independent spatial variation of the signal throughout the home; and
- the tone detectors or receivers are fairly simple, cheap to construct, and have low power requirements.

In the PLP system, we extend the wire-finding technique to include two plug-in signal generator modules. We connect the modules directly into electrical outlets, and their respective signals emanate from those outlets to the rest of the home. We install one of the two modules into an outlet close to the main electrical panel or circuit breaker and plug the other into an outlet that is located along the powerline infrastructure furthest from the first module (see Figure 1). In most cases, physical distance is a good estimate for electrical distance. In the case of a two-story house with a basement, for example, one module would be placed at the west end of the house in the basement (where the main panel is located) and the other in the east end on the second floor. Each module emits a different frequency tone throughout the powerline. As part of the installation, the signal strength must be adjusted such that significant attenuation occurs and the tone signal still reaches the opposite end of the home. Both modules continually emit their respective signals over the powerline and portable tags equipped with specially tuned tone detectors sense these signals in the home and relay them wirelessly to a receiver in the home. Depending on the location of the portable tag, the detected signal levels provide a distinctive signature, or fingerprint, resulting from the density of electrical wiring present at the given location. A receiving base station in the home (*e.g.*, a wireless receiver connected to a PC) analyzes the fingerprint and maps the signal signature to its associated location based on a site survey.

We currently focus on amplitude of the tones only, which has shown good results on its own. However, phase difference between tones is another feature characteristic that can further assist in localization and is the basis of some of our future work.

When the modules are active, the tone detector or receiver tag detects the presence and amplitude of the attenuated signals throughout the home. Because electrical wiring typically branches inside the walls, ceiling, and floors, signal will be present throughout much of the main living areas of the home. Some factors that contribute to the amplitude of the received signal at any given location:

- the distance between the receiver and electrical wiring,
- the density of electrical wiring in an area, and
- the length of electrical wiring from the modules to the receiver's location.

Figure 2 shows a signal map of a bedroom (left) and of a kitchen (right) from the same house. In the bedroom, the strength of both signals increases near the walls where there is the greatest concentration of electrical wiring and outlets. The strength of signal A (left value in each cell of Figure 2) is weaker than the strength of signal B (right value in each cell) in the kitchen, and the opposite is true for the bedroom. Because the two rooms are on different floors and at opposing ends of the house, each room is closer to a different module.

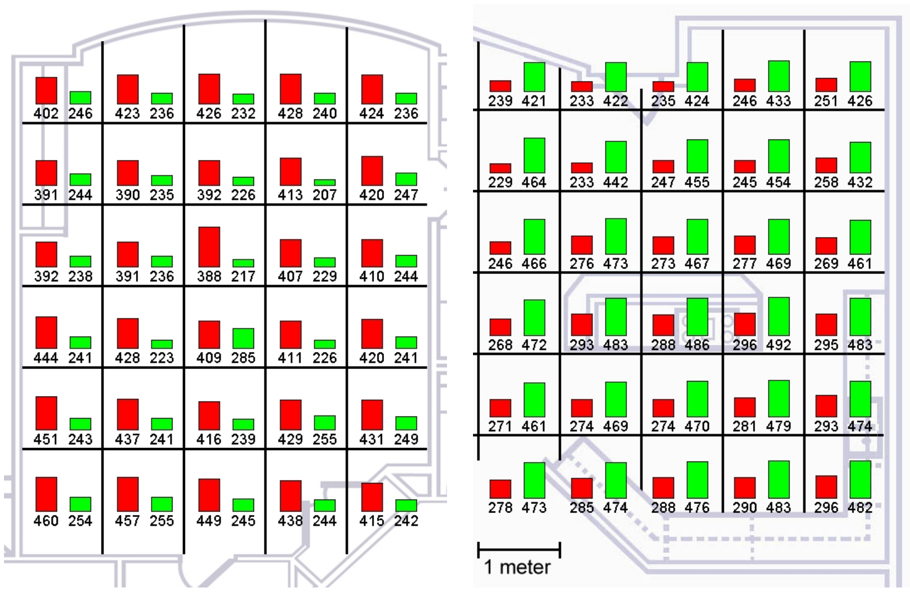

Fig. 2. Left: Signal map of a bedroom. In each 1 meter cell, the left-hand number corresponds to signal strength from one tone generator and the right-hand number corresponds to the signal strength of the other tone generator. Right: A similar signal map of the kitchen in the same house.

Most residential houses and apartments in North America and many parts of Asia have a single phase or a split single phase electrical system, which enables any signal generated on a given outlet to reach the entire electrical system. Larger buildings and even some homes in Europe have two and three phase electrical systems, in which the electrical system may split into separate legs for lower voltage applications. For multi-phase electrical systems, the signal can be coupled between the phases using a simple capacitor. In a home, this would typically be plugged-in in a 240 V outlet, such as that used for clothes dryer. We currently focus on common residential single or split single-phase electrical systems operating at 60 Hz. However, the system can be extended to accommodate other electrical systems.

3.2 PLP Localization Algorithm

The PLP system relies on a fingerprinting technique for localization. Although this technique often provides more detailed and reliable data, it requires the generation of a signal topology via a manual site survey. The granularity of the survey dictates the final accuracy of the positioning system. For PLP in the home, the site survey is a one-time task provided the modules stay fixed and the electrical characteristics of the home remain the same.

Effective application of fingerprinting requires the signals to have low temporal variations, but high spatial variation. As discussed above, the propagation of signals transmitted via the powerline exhibits both of these properties, because the detected signals vary little unless the modules have been moved or the electrical system has been significantly remodeled. The use of two different signals and the variability in the electrical wire density throughout the homes provides this spatial variation.

The localization algorithm used in PLP proceeds in two steps. The first step predicts the room, and the second predicts the sub-regions within that room. Both use *k*-Nearest Neighbor (KNN) classification.

3.2.1 *k*-Nearest Neighbor (KNN) Classification

The room and sub-room localizers use a *k*-Nearest Neighbor (KNN) [11] classification to determine the receiver's room location. KNN-is a memory-based model defined by a set of objects known as learned points, or samples, for which the outcomes are known. Each sample consists of a data case having a set of independent values labeled by a set of dependent outcomes. Given a new case of dependent values (the query point or unknown value), we estimate the outcome based on the KNN instances. KNN achieves this by finding *k* examples that are closest in distance to the query point. For KNN classification problems, as in our case, a majority vote determines the query point's class. For this task, given an unlabeled sample, χ, we find the *k* closest labeled room samples in our surveyed data and assign χ to the room that appears most frequently within the *k*-subset. For our distance measure *d*, we use the Euclidean distance,

$$d\ (x, y) = \sqrt{(\sum_{i=1}^{2}(x_i - y_i)^2)}\,,$$

in which tuples x = <Signal A_{x1}, Signal B_{x2}> and y = <Signal A_{y1}, Signal B_{y2}>. The tuple *x* refers to the labeled signal point and tuple *y* refers to the unlabeled query point sensed by the receiver tag. For more modules, we increase the dimension to match the number of modules.

3.2.2 Room and Sub-room Localization

The key differences between the room and sub-room localizers are the labels assigned to the data points and the value for *k* used in the localization. For the room level classification, we assign room labels to samples from the site survey. In the sub-room classification, we further subdivide the same samples and assign sub-room labels to them. For each home, there is an optimal value of *k* for the room level localizer. Within the same home, there is an optimal value for the sub-room level localizer for each room. Thus, for localization, we first execute the KNN classification using the room labeled samples and its optimal *k* value. After determining the room, we execute KNN on the sub-room labeled samples from that room and its optimal *k* value to determine the sub-room.

3.2.3 Training the System and Determining *k* in KNN

The choice of *k* is essential in building the KNN model and strongly influences the quality of predictions, for both room-level and sub-room-level localization. For any given problem, a small value of *k* will lead to a large variance in predictions. Alternatively, setting *k* to a large value may lead to a skewed model. Thus, *k* should be set to a value large enough to minimize the probability of misclassification and small enough (with respect to the number of cases in the example sample) so that the *k* nearest points are close enough to the query point. Thus, an optimal value for *k* achieves the right balance between the bias and the variance of the model. KNN can provide an estimate of *k* using a cross-validation technique [11].

Splitting the localization into two steps can help control the cluster sizes. In localizing the room, we want to use a larger value of *k* so that we consider a larger region when trying to find where the unknown signal potentially maps. To localize within a room, we consider smaller values of *k* so that we match finer clusters and because of the smaller data sets within a room than the whole home.

The training interface allows end users to build a signal map of the home (see Figure 4). The user loads a pre-made or hand-drawn floor plan of the residence into the application. The interface displays the floor plan, and we physically travel to different locations in the home and choose the approximate location on the floor plan. When a location is selected, the application stores the fingerprint for that location, which is a one-second average of the two detected signals. The same process continues throughout different points in the home. Surveying at a granularity of approximately 2-3 meters in each room produces more than sufficient accuracy for the test cases presented in Section 4. The interface allows the user to assign meaningful labels to different room and sub-room areas, such as “kitchen” and “center of master bedroom.”

For optimal performance in sub-room level localization, we typically segment each room into five regions: the center of the room and areas near the four walls of the room. The user is free to select the location granularity (assuming sufficient training sets) of their choice for important regions. However, the desired segmentation may not reflect the actual segmentation the underlying set of signals can provide. For example, a user may want to segment the middle part of a bedroom into four regions, but there might not be enough signal disparity among those regions for the KNN classifier to work well. We provide some assistance in overcoming those limitations by automatically clustering the room into potential sub-regions that are likely to be accurately classified based on the room’s signal map. We employ a *k*-means clustering algorithm [4, 10, 11] to provide graphical suggestions on where to segment for a desired number of sub-regions.

After the signal map has been constructed and all data has been labeled, the algorithm cross-validates model data to find suitable k values for the room and sub-room classifiers. Cross-validation involves the division of the data samples into a number of v folds (randomly drawn, disjoint sub-samples or segments). For a fixed value of k, we apply the KNN model on each fold and evaluate the average error. The system repeats these steps for various k values. The system selects the value for k achieving the lowest error (or the highest classification accuracy) as the optimal value for k. This value for k depends on the home and the number of sample points. Generally, we see optimal k values near 10 for the room localizer and k values near 3-5 for the sub-room localizer.

3.3 Implementation Details

3.3.1 Module Design

For rapid development and investigation, we modified commercially available tone generators and tone detectors used by electricians for wire finding. We used the Textron Tempo 508S and the Pasar Amprobe 2000 tone generator modules. These modules produce a 447 kHz and 33 kHz tone, respectively, on an energized 120 V AC powerline without causing any interference to household appliances. Additionally, the modules are powerful enough to transmit a tone up to 500 meters over the electrical wire (both hot and ground) and can be adjusted to emit at a lower signal strength. For the PLP prototype in this paper, we manually adjusted the signal strength depending on the size of the residence. We collected samples with the receiver near the module and samples near the opposite side of the home where the second module is located. We then tuned the signal strength such that we produced a large signal difference between the two locations without turning it down so much that the tone did not reach the far end. It was important to turn down the output level and use the middle of the receiver's dynamic range, because very high signal strengths would overwhelm the receiver and would not produce as large of a signal difference. Although we manually performed the steps described above, it is possible to build the modules to self-calibrate its output level during the installation and surveying steps.

Based on the cost of the commercial wire-finder that inspired the PLP system, the cost for each module would be approximately US$50.

3.3.2 Tag Design

We modified a Textron Tempo 508R passive wideband tone detector to act as a prototype tag that would send sensed signals to a portable computer for analysis (see Figures 1 and 3). The toner has a built in frequency divider that maps a range of high frequency tones to audible sounds while still preserving the amplitude of the original signal. The receiver's internal frequency divider translated the 447 kHz signal to about 1000 Hz and 33 kHz signal to about 80 Hz. We altered the tone detector to interface with the audio line-in jack of a portable computer to capture the signals. The tone detector also has an integrated omnidirectional antenna. We found the antenna worked best when held vertically (perpendicular to the ground). When placed in this position, the azimuth orientation did not affect the received signal levels.

For experiments reported in this paper, we used a rather large tag prototype that was easier for us to build. There are a variety of ways to construct a small and inexpensive version of this tag. One way is to feed the radio transducer or antenna through a series of op-amps and into a DsPIC microcontroller. A low-power Ming or Linx RF transmitter would transmit the readings back to a receiving computer. Alternatively, we could bypass the need for a microcontroller by using multiple tone decoder ICs, similar to the NE567 IC, which supports signal power output. Powered by a small lithium cell, the tag could easily be the size of a small key fob and run for a significant period of time using a mechanical motion switch. We believe the tags could be constructed at US$20 each, based on current retail hobbyist prices.

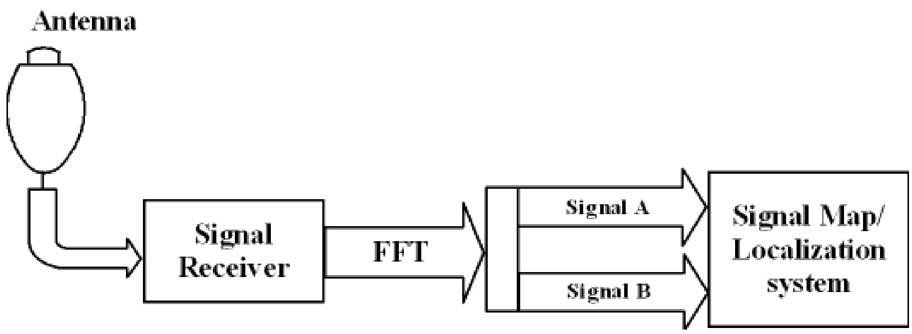

Fig. 3. Block diagram of the overall tagging system of the PLP System

3.3.3 Software

In our experimental set-up, we wrote an application in C++ to sample the signal from the sound card's line-in jack where the prototype receiver tag is connected. The application acquires 16-bit samples at a rate of up to 44 kHz and performs a Fast Fourier Transform (FFT) on the incoming signal to separate component frequencies for our analysis. The application performs this analysis in very close to real-time and makes the raw signal strengths for the two frequencies of interest (447 kHz and 33 kHz) available through a TCP connection for other parts of the PLP system to access (see Figure 3).

A second application, written in Java, performs the machine learning and provides the user interface for the system (see Figure 4). The Java application connects to the FFT application and reads the raw signal values. The application provides the user interface for surveying the home and an interface that shows the current location after it has been calibrated. The Weka toolkit [17] allows for real-time programmatic execution of KNN queries to our location model. We also use Weka for *post hoc* analysis, such as cross-validating the model when determining optimal k values and performance testing.

The experimental prototype used for empirical validation consisted of a Sony Vaio-U handheld computer with all software applications (signal receiver, learner, and the user interfaces) loaded and the receiver hardware connected (see Figures 1 and 4). Using this small but powerful device provided us with an easy method for surveying homes.

Fig. 4. User interface used for mapping and localizing the position of the connected receiver

4 Performance Evaluation

We evaluated the performance of the PLP system in 8 different homes of varying styles, age, sizes, and locations within the same metropolitan city. We evaluated new homes and older homes, both with and without remodeled and updated electrical systems (see Table 1 for specifications of the homes). In addition to evaluating our system, we simultaneously conducted infrastructure tests of WiFi and GSM availability to provide some comparison with other indoor localization results. The infrastructure tests only involved logging the availability of wireless 802.11 access points and multiple GSM towers in the home. A WiFi spotter application running on the Sony Vaio-U logged the wireless access points, and an application written on the Audiovox SMT-5600 GSM mobile phone the multiple cell towers.

In each home analyzed, we first installed the PLP system, calibrated the two tone modules, and created a signal map by surveying the whole home. When creating the signal map, we took at least two signal readings every 2-3 meters throughout the home to ensure we gathered enough training and test data (Table 2 shows the number of sample points for each home). Each reading was taken for 3 seconds with an individual holding the receiver in hand (about 1.5 meters from the ground). After creating the signal map, we used the interface on the handheld to assign the appropriate room and sub-room labels to the data.

We reported the classification accuracy of the room and sub-room predictors. The sub-room accuracy was calculated independent of the room-level predictor. We use 3 meter regions for the sub-room-level tests. To obtain room-level accuracy, we conducted a 10-fold cross-validation test on the room localizer using the collected data samples. We repeated this test for various k values to find the best accuracy measure, which also served as our reported accuracy value. To determine the sub-room level accuracy, we took the data samples for each room and performed a 10-fold cross-validation using the sub-room localizer, again for different values of k. Similar to the room-level tests, we looked for the *k* value that provided the highest accuracy for predicting regions in a room. After testing each room, we averaged all the sub-room localization accuracies to produce a single accuracy value for each home.

Table 1. Details of the homes where the PLP system was deployed and evaluated

Home	Year Built	Electrical Remodel Year	Floors/ Total Size (Sq Ft)/ (Sq M)	Style	Bedrooms/ Bathrooms/ Total Rms.	Population Density
1	2003	2003	3/4000/371	1 Family House	4/4/13	Suburb
2	2001	2001	3/5000/464	1 Family House	5/5/17	Suburb
3	1992	1992	1/1300/120	2 Bed Apartment	2/2/6	Downtown
4	2002	2002	3/2600/241	1 Family House	3/3/12	Suburb
5	1967	2001	2/2600/241	1 Family House	3/3/11	Suburb
6	1950	1970	1/1000/93	1 Family House	2/2/5	Suburb
7	1926	1990	1/800/74	1 Bed Loft	1/1/5	Downtown
8	1935	1991	1/1100/102	1 Family House	2/1/7	Suburb

4.1 PLP Accuracy

4.1.1 Between Homes Comparison

In Table 2, we report the results of the PLP room-level and sub-room level accuracies for various homes. Room accuracy ranged between 78–100% and sub-room accuracy ranged between 87–95%. The modern homes and the older homes with updated electrical infrastructure resulted in similar performance results. The updated electrical systems in these homes were accompanied with an overall remodel of the home, which tends to include the addition of electrical outlets and lighting. The single family home that exhibited a significantly lower accuracy (Home 8) was an older home with an updated electrical system. However, that home had a two-phase electrical system, which we only learned after installing the PLP system. Because it is a smaller house and Phase 1 drives a small number of outlets, we simply placed the modules on Phase 2 to produce acceptable (though not optimal) coverage throughout the house. However, installing a simple phase coupler would have improved its performance.

The condominium and apartment test cases also produced promising results. The condominium was converted from an office building, but the electrical system was completely remodeled to a residential style system. Although one wall of the condominium used a metal conduit to run its electrical wire, PLP still worked because the room with the conduit was small and the receiver was never too far from the wall. The apartment also featured a similar residential style electrical system. Because of the small size of the living spaces, we had to turn down the power of the modules significantly in the two cases, unlike the larger homes we tested.

The older homes without an updated electrical system exhibited lower results for two reasons. First, these homes lack a proper electrical ground, resulting in one less path for the signal to propagate, because we send the signal both on the hot and ground wires. Homes with an updated electrical system have an extra electrical ground wire running through the home, which is usually grounded to the copper water

Table 2. Accuracy results by home. For each home, we report the accuracy of room-level prediction and the average sub-room-level prediction across all rooms (Note the room-level and sub-room accuracy values are independent of each other). The sub-room-level regions were defined to be approximately a 3 meter square. The WiFi and GSM measurements indicate the maximum number of access points or towers seen at all times during the surveying and the total number of unique access points or towers seen during the whole surveying period.

Home	Size Sq Ft/ Sq M	Sample points	Rooms surveyed	Room Accuracy	Sub-Room Accuracy at 3 M	WiFi Always/ Max	GSM Always/ Max
1	4000/371	194	13	89%	92%	3/12	3/5
2	5000/464	206	15	95%	93%	1/3	2/4
3	1300/120	95	6	90%	90%	3/7	4/12
4	2600/241	183	11	88%	87%	1/3	3/5
5	2600/241	192	10	92%	93%	2/4	3/6
6	1000/93	76	5	100%	94%	0/2	4/6
7	800/74	65	5	93%	95%	2/11	3/9
8	1100/102	80	7	78%	88%	2/6	3/7

pipes. This grounding enables additional signal propagations to certain areas of the home. Second, these homes tended to have fewer electrical outlets than the modern or remodeled ones, resulting in poor detection in some areas.

4.1.2 Understanding Classification Errors

To understand the types of classification errors encountered by the PLP system, we analyzed the confusion matrices for each home. For some homes, most of the classification errors resulted from misclassifying rooms as one of the adjacent rooms. The adjacency errors appeared when trying to localize very near the boundary or the wall of a room. These errors were more prevalent in larger houses near common walls between two adjacent rooms of similar size. Open spaces that were divided into multiple rooms also resulted in errors. Other homes, however, exhibited more random

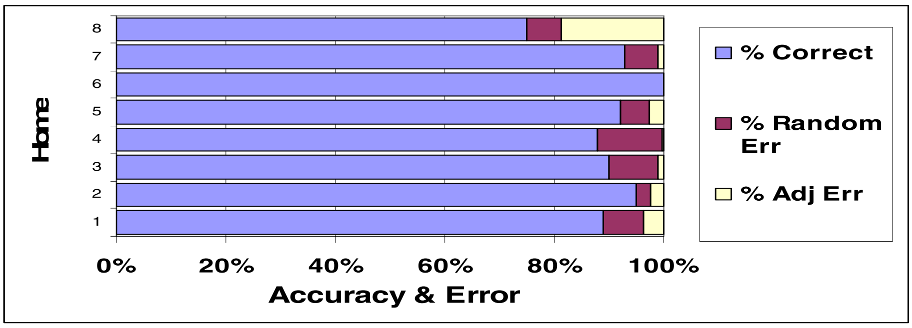

Fig. 5. The table shows the percentage of incorrect room predictions identifying a room that is adjacent to the correct room

classification errors possibly due to errors in the survey map, sparse sampling, or in error readings coming from the receiver at that time. One possible solution to guard against misclassifications is to use hysteresis to compare against certain classifications and see if those classifications follow a valid trail. Some homes could benefit from hysteresis, especially those with significant random error (see Figure 5).

4.2 Number of Modules and Performance

We conducted accuracy tests using a varying number of modules. Although our goal was to minimize the additional hardware the user must install in a home, there might be cases in which higher accuracy is more desirable. Adding additional modules is the main way to increase overall accuracy. Figure 6 shows both room-level and sub-room level accuracies for an increasing number of modules for a particular home as an example. Additional modules do increase the accuracy for both predictions, but there is a point of diminishing returns. For this home (Home 1), two or three modules are the best number. We observed similar trends in other homes we tested and generally, two modules were sufficient.

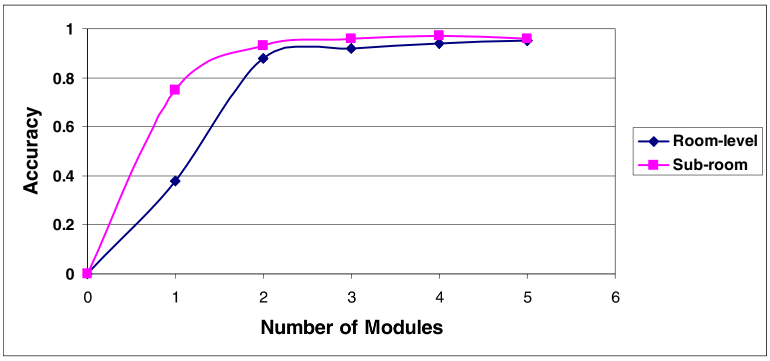

Fig. 6. The effect of number of modules on room-level and sub-room-level classification accuracies. Tests were conducted on Home 1.

4.3 Resolution

In our initial evaluation, we sub-divided rooms into approximately 3 meter regions. This resolution yielded high accuracies around 90%. Higher resolution, or smaller subdivisions of each room, is possible, but at the cost of accuracy. In addition, higher resolution also requires dense mapping of an area. To investigate the specific accuracy to resolution tradeoff, we performed a fine-grain survey (sampling down to every 0.5 meter for a total of 96 samples) of a room (6m X 6m) in Home 1. With our current implementation, the lowest obtainable practical resolution is 1 meter. The accuracy falls below 70% for 1 meter regions (see Table 3), because there is a theoretical limit to the detectable differences between small movements in the space and the signal amplitude. However, finer granularity may be possible by considering the phase difference between the two signals. From our observation, the maximum amplitude differential is about 20 units when moved 1 meter for a modern home.

Table 3. The sub-room-level accuracies for smaller sub-regions for a particular room in Home 1. A total of 96 points were surveyed.

Sub-room region size	4 m	3 m	2 m	1 m	0.5 m
% Accuracy	94%	91%	74%	67%	42%

4.4 Temporal Signal Stability

Fingerprinting works best with a signal that is time-independent but spatially diverse. The data presented so far only considered results over relatively short periods of time, usually around 1 hour worth of data collected at a particular home. To test the stability of the signals over time, we conducted two separate tests. First, in Home 1, we conducted separate surveys over the course of several weeks. We trained the system on data from one survey and checked its accuracy against data collected from different surveys. Room prediction was correct 88% of the time (compared with the value of 89% for Home 1 in Table 3) and sub-room level prediction was correct 89% of the time (compared with the value of 90% in Table 3). Second, in Home 2, we collected 45 hours of data over a three-day period (Saturday through Monday) in a single location (the kitchen). The kitchen was an interesting test because it contained a large number of features that could affect the tone signals (*e.g.,* plentiful overhead lighting, appliances being turned on and off throughout the day, talking on a cordless phone, people gathering around the tag). Figure 7 depicts the stability of the signal for four different 3-hour intervals. The results suggest there is deviation (17 units on average), but it is not significant enough over the full dynamic range to cause major classification errors.

Modifications to the electrical infrastructure can contribute to accuracy errors and require recalibration, which was a problem we noted for other infrastructure solutions (802.11 and GSM). However, most situations, such as turning on a light switch, only energize a portion of the electrical line and do not affect significantly the accuracy in our experience. More studies are needed to empirically study this. Construction of a “day” and “night” map using a richer data set can allay some of these concerns. The addition of an extension cord may impact the accuracy, depending on location and length. PLP could be designed to recognize potential changes in the infrastructure from past data to notify the user that re-surveying of a particular area is necessary.

Although we did not observe any problems with electrical interference with our continuous logging, during our site tests we did often observe electrical interference caused by home electronics and appliances, such as computers, televisions, and stereos. When we held the receiver next to some of these electronic devices, its broadband electrical noise often overwhelmed the receiver and caused spurious readings. This problem only existed when the receiver was very close (within a few centimeters) from such devices. To guard against learning or localizing incorrect fingerprints, one solution is to look for these signal interferences and filter out those readings, indicated by a clear broadband signature, before using the data in analysis.

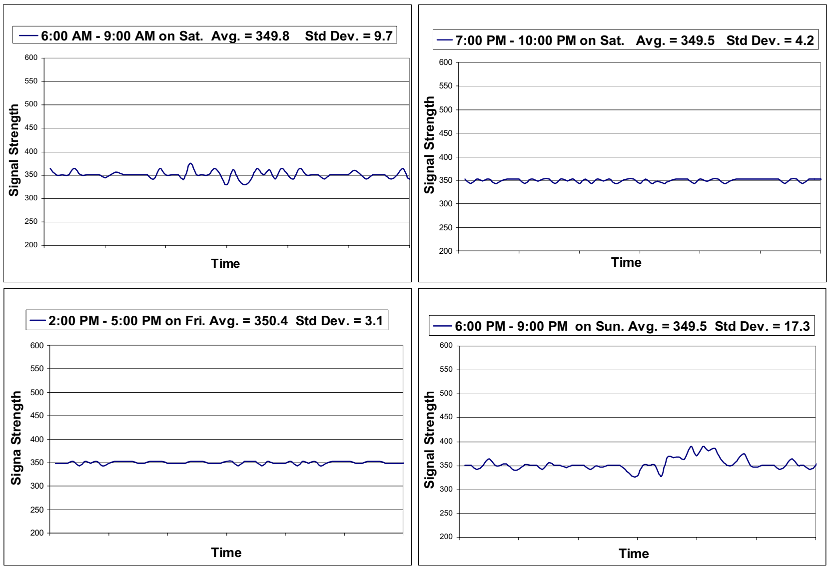

Fig. 7. Temporal signal stability in the kitchen area of Home 2. The graphs show the signal values for the two toner modules (combined using the Euclidean distance) over various intervals during four days of continuous recording. The average signal values and the standard deviations are shown above each graph. The full dynamic range of the vertical axis is 0-1000.

5 Discussion

PLP is very promising as an inexpensive and reliable sub-room-level indoor positioning service. In this section, we investigate the viability of this system and offer some comparison to previous solutions.

5.1 Infrastructure and Cost Comparison Against WiFi and GSM

The cost of infrastructure for WiFi is distributed across a community and assuming dense enough living conditions, it is a reasonable expectation that a single residence will be able to access other WiFi access points nearby. This is less likely in sparser housing, in which case users would be required to purchase multiple WiFi access points. Various cellular telephony service providers cover the cost of the infrastructure for GSM. The coverage is fairly dense in most metropolitan areas and will only get better over time. However, coverage is still fairly sparse in rural settings and many homes do not get very good cellular service in some rooms (see Table 2). Almost every home in the U.S. has electrical power, and it is an assumed cost of the homeowner to maintain this infrastructure over the lifetime of the home. Thus, the infrastructure is already available and usually well maintained.

One key advantage of leveraging the powerline infrastructure is user control of the infrastructure. Users have very little control of the parameters of GSM cellular towers or a neighbor's WiFi access point, thus changes can happen unexpectedly. In contrast, users have control of the powerline infrastructure. Furthermore, as we showed in Section 4.4, there is stability in signal propagation over this infrastructure.

The cost and power requirements of the location tags favor that of the PLP system because of its simple sensing requirements, as opposed to the more sophisticated chipset associated with GSM and WiFi reception. In addition, the cost of the tone generating modules would also be cheaper than buying additional access points if one were investing in a location system for the home.

5.2 The Powerline Infrastructure

In the United States, modern homes now follow a strict electrical code called the National Electronic Code (NEC). Electrical codes only became widely enforced in the 1980s, although many homes before that already followed similar guidelines. Although the specific regulations may change depending on state and city ordinances, each follows the same general requirements. These regulations ensure the electrical systems are consistent across homes of different sizes and styles. Specifically, the requirements outlined in the NEC favor the infrastructure requirements needed for the PLP system to work in modern homes. These requirements include regulations for certain "homerun" circuits through the house, a minimum number of outlets in a given space, and minimum lighting requirements throughout the house. Although PLP already performed reasonably well in older homes, it consistently achieved very good results in the new or remodeled homes that follow these requirements (see Table 3).

We specifically developed PLP to provide an affordable location system for home environments. However, commercial buildings must comply with strict electrical codes for which the PLP design must be altered to support. First, commercial wiring typically uses a two or three phase electrical system that prevents the signals from propagating throughout the entire electrical system. This problem is solved by installing an inexpensive phase coupler. Second, most commercial electrical wiring runs through a metal conduit, which blocks significant portions of the tone emanating from the wire (PVC conduits do not cause a problem). One solution to this problem is to increase greatly the signal strength and the other is to send the signal through both the electrical wiring and the metallic conduit itself. This problem also applies to homes that have been converted from commercial buildings without remodeling the electrical system.

5.3 General Comparison of PLP Against 802.11 and GSM

The significant advantage of PLP when compared against two popular fingerprinting techniques using WiFi/802.11 [2] and GSM [13] lies in the better resolution, control of the infrastructure, and power requirements (see Table 4).

Table 4. An overall comparison of PLP against two popular location systems that also use fingerprinting

	PLP	GSM	WiFi
Output Type	symbolic	symbolic	symbolic (geometric using triangulation)
Resolution and Accuracy	3 m – 90% 1 m – 67%	20 m – 90% 2-5 m – 50% [13]	6 m – 90% 2-3 m – 50 % [2]
Infrastructure Requirements.	2 plug-in signal modules	Located within GSM cellular service range	3 – 4 WiFi access points
Infrastructure Control	Full	None	Partial (dependent on ownership of access points)
Cost	US$20 for tag and US$50 per module	US$25 for tag	US$25 for tag and US$50 per access point
Spectral Requirements	10 kHz – 500 kHz	900 MHz and 1800 MHz	2.4 GHz
Update Rate	> 20 Hz	> 20 Hz	> 20 Hz
Tag power Req.	~50 mA (Pic + op-amp + antenna)	~200 mA (GSM receiver module)	~100 mA (microcontroller operated WiFi detector)
Simultaneous Tracking	Theoretically no limit	Theoretically no limit	Theoretically no limit

6 Conclusions and Future Work

PLP is a promising indoor positioning system for the home that uses its powerline infrastructure and requires only the addition of two plug-in modules to the home infrastructure and the use of simple location tags. The system is capable of localizing to sub-room level precision using a fingerprinting technique on the amplitude of tones produced by the two modules installed in extreme locations of the home. The density of electrical wiring at different locations throughout the home provides a time-independent spatial variation of signal propagation.

Our critical analysis of PLP, and the experimental validation in eight different homes, suggests the following advantages over current indoor location solutions:

- PLP leverages a truly ubiquitous resource, the powerline infrastructure, available in almost all homes.
- PLP requires very minimal addition to the infrastructure (two plug-in modules).
- PLP achieves superior sub-room-level classification, with an accuracy of 93% on average at a resolution of 3 meters.
- PLP does not detract from the appearance of the home.

Our next step is to build smaller, less expensive, and lower powered tags for practical deployments of PLP. In addition, we plan to incorporate other spatially varying signal features, such as phase differences between the tones in addition to the amplitude to increase the accuracy and resolution of PLP in the fingerprinting process. Further stability analysis is also planned to determine the full viability of PLP.

References

1. Active Bat. The BAT Ultrasonic Location System. http://www.uk.research.att.com/bat/.2006.
2. Bahl, P. and Padmanabhan, V. RADAR: An In-Building RF-Based User Location and Tracking System. In the proceedings of *IEEE Infocom.* Los Alamitos. pp. 775-784. 2000.
3. Castro, P., Chiu, P., Kremenek, T., and Muntz, R.R. A Probabilistic Room Location Service for Wireless Networked Environments. In proceedings of *Ubicomp* 2001. pp. 18-34. 2001.
4. Hamerly, G. and Elkan, C. Learning the k in k-means. In the proceedings of *The Seventeenth Annual Conference Neural Information Processing Systems (NIPS)*. 2003.
5. Hightower, J. and Borriello, G. A Survey and Taxonomy of Location Systems for Ubiquitous Computing, University of Washington Tech Report CSC-01-08-03. 2001.
6. HomePlug Powerline Alliance. http://www.homeplug.org. March 2006.
7. Krumm, J., Cermak, G., and Horvitz, E. RightSPOT: A Novel Sense of Location for a Smart Personal Object. In the proceedings of *Ubicomp 2003,* Seattle, WA, pp. 36-43. 2003.
8. LaMarca, A., Chawathe, Y., Consolvo, S., Hightower, J., Smith, I., Scott, I., Sohn, T., Howard, J., Hughes, J., Potter, F., Tabert, J., Powledge, R., Borriello, G., and Schilit, B. Place Lab: Device Positioning Using Radio Beacons in the Wild. In the proceedings of *Pervasive 2005*, Munich, Germany. pp. 116 – 133. 2005.
9. Madhavapeddy, A. and Tse, T. Study of Bluetooth Propagation Using Accurate Indoor Location Mapping. *The Seventh International Conference on Ubiquitous Computing (UbiComp 2005).* Tokyo, Japan. pp 105-122. September 2005.
10. MacQueen, J.B. Some Methods for Classification and Analysis of Multivariate Observations. In the proceedings of *The Fifth Berkeley Symposium on Mathematical Statistics and Probability*. pp 281-297. 1967.
11. Mitchell, T. *Machine Learning*, McGraw Hill. ISBN 0070428077. 1997.
12. O'Connell, T., Jensen, P., Dey, A.K., and Abowd, G.D. Location in the Aware Home. Position paper for Workshop on Location Modeling for Ubiquitous Computing at Ubicomp 2001. September 30, Atlanta, GA. 2001.
13. Otsason, V., Varshavsky, A., LaMarca A., and de Lara, E. Accurate GSM Indoor Localization. In the proceedings of *The Seventh International Conference on Ubiquitous Computing (UbiComp 2005)*. Tokyo, Japan. September 2005.
14. Priyantha, N. B., Chakraborty, A., and Balakrishnan, H. The Cricket Location-Support System. In the proceedings of *The International Conference on Mobile Computing and Networking (Mobicom 2000)*. Boston, MA. August 2000.
15. Ubisense. http://www.ubisense.net. 2006.
16. Want, R., Hopper, A., Falcao, V., and Gibbons, J. The active badge location system. *ACM Transactions on Information Systems*. Volume 10. pp. 91-102. January 1992.
17. Weka 3: Data Mining Software in Java. http://www.cs.waikato.ac.nz/ml/weka/. March 2006.