

Pazl: A Mobile Crowdsensing based Indoor WiFi Monitoring System

Valentin Radu, Lito Kriara and Mahesh K. Marina
The University of Edinburgh

Abstract—WiFi in indoor environments exhibits spatio-temporal variations in terms of coverage and interference in typical WLAN deployments with multiple APs, motivating the need for automated monitoring to aid network administrators to adapt the WLAN deployment in order to match the user expectations. We develop Pazl, a mobile crowdsensing based indoor WiFi monitoring system that is enabled by a novel hybrid localization mechanism to locate individual measurements taken from participant phones. The localization mechanism in Pazl integrates the best aspects of two well known localization techniques, pedestrian dead reckoning and WiFi fingerprinting; it also relies on crowdsourcing for constructing the WiFi fingerprint database. Compared to existing WiFi monitoring systems based on static sniffers, Pazl is low cost and provides a user-side perspective. Pazl is significantly more automated than wireless site survey tools such as Ekahau Mobile Survey tool by drastically reducing the manual point-and-click based measurement location determination. We implement Pazl through a combination of Android mobile app and cloud backend application on the Google App Engine. Experimental evaluation of Pazl with a trial set of users shows that it yields similar results to manual site surveys but without the tedium.

I. INTRODUCTION

Significant interest in mobile phone sensing in recent years can be attributed to several factors, including: their ubiquitous nature; rapid evolution toward smartphones with several built-in sensors; carried by humans, making them natural to be used for “mobile” sensing; and the possibility of leveraging the cloud via several available connectivity options for computing power, storage and “centralization”. Not surprisingly then, mobile phone sensing applications have been realized or envisioned in diverse domains (e.g., transportation, social networking, health monitoring) [1], [2]. When a group/community of participants (a *crowd*) is engaged with suitable incentives, mobile phone sensing becomes even more compelling for continual and fine-grained spatio-temporal monitoring of the phenomenon of interest in a *cost-effective* manner. Indeed, as Xiao et al. note in [3], the focus of mobile sensing research and applications is shifting towards *mobile crowdsensing*, which is defined as “individuals with sensing and computing devices collectively share data and extract information to measure and map phenomena of common interest” [4]. Several mobile crowdsensing applications have been developed and deployed (e.g., [5], [6]) and it remains a very active area of research.

We consider the application of the mobile crowdsensing paradigm to wireless network monitoring. Besides the many sensors, modern mobile phones feature several wireless network interfaces as connectivity options (e.g., cellular, WiFi, Bluetooth, NFC). Discussions of mobile phone sensing have been mostly centered around the use of built-in sensors and/or

specialized add-on sensors (e.g., GasMobile [5], CellScope¹, NETRA²) with connectivity options serving as a means for data sharing (see [2], for example). We expand this commonly held view to treat network interfaces also as sensors. GPS, which is an integral part of all smartphones today, presents an example of a network interface that sits at the boundary of these two views — GPS is seen as a location sensor for mobile phone sensing applications whereas it is actually a RF communication system in which GPS receiver on a phone uses signals transmitted from satellites for localization³. Another more obvious example is the use of cellular interface on smartphones for crowdsourcing based active/passive measurement of mobile networks as in [7], [8].

In this paper, we apply the mobile crowdsensing paradigm for low-cost and automated indoor WiFi monitoring. Specifically, our focus is on indoor environments with multitude of access points (APs) as is the case with WiFi deployments in enterprises and public buildings (e.g., shopping malls, hospitals). This application exploits the WiFi interface on smartphones as a measurement sensor. It is motivated by the observation that WiFi networks experience a range of coverage and interference related problems that affect users and these problems vary over space and time. Having a system for spatio-temporal WiFi monitoring at low cost allows the network administrators to better manage their wireless LANs and optimize user experience. Towards this end, we develop Pazl, a mobile crowdsensing based indoor WiFi monitoring system. The main challenge underlying the design of Pazl is the need to locate measurements taken from smartphones and the difficulty of indoor location and navigation — GPS does not typically work indoors and is known to be energy hungry even when it does work. We address this challenge via a novel hybrid localization mechanism in Pazl that also embraces crowdsourcing.

Pazl advances the state of the art on WiFi monitoring in two important ways: (1) compared to WiFi monitoring systems based on statically positioned sniffers (e.g., DIST [9]), Pazl is not only lower cost leveraging people’s smartphones and their movements but also complementary in the sense that it captures the vital user/client side perspective as opposed to the AP side or monitoring system perspective. (2) with respect to wireless site survey solutions (e.g., Ekahau Mobile Survey [10]), Pazl significantly lessens the need for manual point-and-click approach for identifying measurement location,

¹<http://cellscope.berkeley.edu/>

²<http://web.media.mit.edu/~pamplona/NETRA/>

³Technical specifications of some smartphones do acknowledge this view. See <http://www.samsung.com/global/galaxys3/specifications.html>, for example.

thereby paving the way for automated monitoring.

This paper makes the following key contributions:

- We provide evidence for spatio-temporal variability in coverage and interference characteristics of WiFi networks from a user-perspective that motivates automated monitoring. This is done via measurements obtained for a large WiFi deployment consisting of several tens of APs spread across 6 floors of the Informatics Forum building in Edinburgh, UK (section III).
- We design a hybrid indoor mobile phone localization mechanism (section IV) that combines the best aspects of two well-known localization techniques, pedestrian dead reckoning (PDR) and WiFi fingerprinting, neither of which is sufficient by itself for our purpose — location error accumulates over time with PDR especially when based on smartphone sensors and in indoor environments with complex human activities such as using the stairs, whereas WiFi fingerprinting does not work when there is no WiFi coverage which is of interest from a monitoring viewpoint. A key feature of our localization mechanism is that it exploits the locations deemed to be accurately localizable via WiFi fingerprinting for correcting PDR based location estimates. The reference fingerprint database required for WiFi fingerprinting is also constructed in a crowdsourcing manner in our proposal. We show that our localization mechanism achieves a median location accuracy below 3 meters in a building of approximately 12000 m² spread over 5 floors (section V.A).
- We develop Pazl, a mobile crowdsensing based indoor WiFi monitoring system that incorporates the above mentioned hybrid localization mechanism. The implementation of Pazl consists of two parts: (1) an Android application for collecting WiFi and sensor (accelerometer and compass) measurements from each mobile crowdsensing participant's smartphone; (2) a cloud application based on the Google App Engine to localize the measurements from different phones and to merge, store, visualize and analyze for various monitoring related aspects (e.g., coverage holes, channel usage distribution, complex interference patterns resulting from exceptionally long range of some APs as seen from certain locations). Through a user trial, we experimentally evaluate Pazl and find that it provides similar results to the state of the art Ekahau Mobile Survey tool [10] but in a significantly more automated manner by drastically reducing the manual point-and-click location determination used in the Ekahau approach (section V. B).

II. RELATED WORK

A. WiFi Monitoring and Site Surveys

Monitoring WiFi networks (802.11 wireless LANs) has received much attention from the research community (see [9] and references therein). The current state of the art approach as exemplified by DIST [9] is to deploy a number of stationary

sniffers (also called air monitors) separate from the WLAN infrastructure. Similarly, WizNet [11] employs many cheap ZigBee sensors in conjunction with digital signal processing techniques to tell apart between 802.11 and other signals. This stationary sniffer approach gives greater visibility and flexibility compared to the earlier approaches of relying on management information obtained from APs or observing the wired side of APs using SNMP, syslog and packet sniffing. However, it can be quite expensive and justified when the focus is on security which indeed is the case for most advanced WLAN monitoring systems like DIST. Also for the stationary sniffer approach to capture the user perception of WiFi networks, the density of sniffers has to be very high. Smart APs which are the norm today also cannot capture the user-side view despite the sophisticated spectrum sensing and centralized intelligence they are equipped with.

In [12], the authors argue in favor of a client-side perspective which can lead to better understanding of the actual coverage and interference conditions (due to interference within the WLAN, from other co-located WLANs and interference from non-WiFi devices using the same spectrum such as bluetooth, microwave ovens and cordless handsets). The latest 802.11 standard [13] incorporates wireless LAN radio measurements support (which were originally introduced in 2008 as part of 802.11k substandard) to get client side measurement reports to assist with seamless mobility etc. but they have two key limitations from our perspective: (1) there is no solution specified to localize measurements from clients indoors; (2) requesting client side reports from the associated AP fails in cases where the client is in an area with no coverage.

Wireless site survey is another relevant process, which is concerned with designing and planning a WiFi network by identifying the number and locations of APs prior to deployment as well as post-deployment walk-testing, analysis or diagnosis of an existing WLAN [14], the latter more related to monitoring and thus this paper. Many software and hardware tools exist to assist with wireless site surveys. Some of them are offline tools based on models such as AirTight Planner⁴ for predictive surveys to estimate coverage for a given survey of APs. A prominent example of an online site survey tool that is closely related to our work is the Ekahau Mobile Survey [10], which we use as the benchmark for our Pazl system. Different from Pazl, the Android based Ekahau Mobile Survey app targets the single user case and also requires the user to repeatedly point-and-click their location on a floor map while walk testing. In [15], a robot based spectrum survey system called Sybot is proposed where the main focus is on reducing the measurement effort. In contrast, Pazl offers a low cost, less disruptive and continual monitoring alternative by leveraging the smartphones carried around by mobile users in the monitored environment.

While monitoring and site surveys could in general involve active measurements for analyzing WLAN performance, in Pazl we limit ourselves to *passive* monitoring via WiFi scans, which are sufficient to diagnose most coverage and performance oriented WLAN problems; we do this to keep disruption and battery drain to crowdsourcing participants minimal.

⁴<http://www.airtightnetworks.com/home/products/AirTight-Planner.html>

B. Indoor Mobile Phone Localization

Common presence of sensors such as accelerometers and compasses in smartphones have made pedestrian dead reckoning (PDR) [16] an attractive technique for mobile phone localization. While some systems combine PDR with a map to avoid war driving for localization in areas away from roads and streets outdoors [17], others such as GAC combine it with occasional GPS correction for energy-efficient location tracking on roads [18]. A well-known limitation of PDR schemes is that error accrues over time unless it is corrected by a more reliable reference in between.

WiFi fingerprinting is another well-known localization technique that can exploit the presence of WiFi interfaces now common on smartphones. WiFi infrastructure is also prevalent these days in many indoor environments. Early WiFi fingerprinting systems such as RADAR [19] and Horus [20] rely on an initial training phase to construct fingerprint database for use as a reference in the positioning phase later. As the training phase can be quite time consuming and expensive, more recent WiFi fingerprinting systems make this training phase automated via crowdsourcing using mechanisms of increasing sophistication (e.g., Redpin [21], OIL [22], WiFi-SLAM [23], Zee [24]). An obvious limitation of WiFi fingerprinting is that it works only where there is WiFi coverage. Moreover, for navigation and continuous location tracking, repeated WiFi scanning can be both energy and time consuming (around 1W and ~500ms per scan).

Combining PDR with WiFi fingerprinting to overcome the above mentioned limitations of both has been considered recently in [25] and [26]. The UnLoc system [25] combines the use of inertial sensors (accelerometer, compass, gyroscope) with the notion of natural and organic landmarks that are learnt over time for indoor navigation. While the use of WiFi fingerprinting in UnLoc is limited to identifying organic landmarks based on radio environment, Pazl uses it more softly and continuously with the aid of a particle filter to opportunistically correct PDR errors. In [26] the use of WiFi fingerprinting, also to correct PDR, is limited to only those locations where maximum signal strength is seen. While both [25] and [26] use a basic PDR scheme, Pazl incorporates a more sophisticated version with activity recognition capability that would be needed in more complex environments (e.g., multi-floor buildings with elevators and stairs to move between floors). Moreover, unlike [25] and [26], for the PDR Pazl uses only accelerometer and compass for the PDR which are present in almost every smartphone, thus achieving wider applicability.

III. MOTIVATION

We observed the WiFi network behavior in the Informatics Forum building at Edinburgh University with several tens of APs deployed for campus WLAN and other APs (from testbeds and co-located wireless networks) spread across 6 floors. The dynamism of the WLAN is obvious when looking at the varying number of APs that can be observed over a period of a few days (Figure 1). There is a clear pattern showing the number of APs decreasing over the night hours and increasing during the day, matching typical usage patterns.

We used the Ekahau mobile site survey application to determine the WLAN coverage on a floor (Figure 2). We found




Fig. 1: Number of active APs detected at a single location over a long period of time.

that some areas, indicated with red color, have poor signal reception or no coverage at all. While most of these places are near metallic staircases or heavy concrete walls, to our surprise poor coverage is seen near some offices where we expect good signal reception.




Fig. 2: Coverage map obtained with the Ekahau Mobile Survey tool.




Fig. 3: AP channel distribution.

Channel usage distribution is an aspect that indicates the potential amount of interference experienced by clients operating on different channels. Figure 3(a) shows the density of APs per channel in both 2.4GHz (channels 1, 6 and 11) and 5GHz bands across an entire floor. Even though the channel usage is quite uneven and thus not ideal, we notice something even more interesting at certain locations. For example, the observed channel usage distribution in one of the offices situated on the inner ring of the building is shown in Figure 3(b). Channels 1 and 11 are heavily used compared to channel 6 in the 2.4GHz band, and the four strongest APs are in channels 1 and 11.

APs with unusually long transmission range can interfere with many other APs operating on the same channel, thus having an adverse effect on WLAN performance. This could be a result of the building environment (material, layout, etc.) as well as AP configuration settings (e.g., high transmission power). We have conducted a manual WiFi site survey for the entire building and determined the physical 3D distance each AP reaches in the building (Figure 4). We observed median coverage radius of about 45 meters. The exception was an AP sensed in all locations measured, and reaching as far as 76 meters. From closer inspection, we identified that this was an experimental AP temporarily setup for research purposes. In practice, similar effect could also result due to the nature of the radio propagation environment.




Fig. 4: CDF⁵ of coverage range of all APs.

The above observations highlight the spatio-temporal differences in coverage and interference in typical WLAN settings with multiple APs and in turn motivate the need for automated monitoring systems. We describe Pazl, our low-cost solution towards this end, in the next section.

IV. PAZL DESIGN AND IMPLEMENTATION

A. Pazl Hybrid Localization Mechanism Overview

As noted at the outset, the key challenge underlying the design of a low-cost and automated indoor WiFi monitoring system based on mobile crowdsensing is locating each measurement. We design a hybrid localization mechanism (illustrated in Figure 5) to address this challenge. Phone’s sensors (accelerometer, compass and WiFi interface) collect samples of acceleration, orientation and WiFi scans. Acceleration is used by the Activity Classification component to detect the activity performed by the user. If this activity is location specific, Map Knowledge aids in estimating the location. Acceleration and Orientation are used in the Pedestrian Dead Reckoning (PDR) component to track the continuous movement. Finally, a WiFi fingerprint is extracted from a WiFi scan and compared with those in a fingerprint database (also created via crowdsourcing) for closest matches in signal space. Locations associated with those matches in the database are used to estimate the location of the user at the moment of scan. All these different estimations are merged using a particle filter to return a single location estimation.

Next, we present the two main components of the hybrid localization mechanism in Pazl: the pedestrian dead reckoning

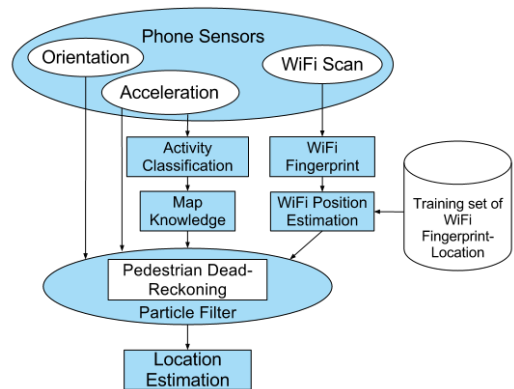


Fig. 5: Schematic of the hybrid localization mechanism.

based continuous location tracking and the WiFi fingerprinting component.

B. Pedestrian Dead Reckoning (PDR)

The PDR method tracks a pedestrian user by starting from a known location and estimating consecutive positions based on traveled distance and direction. To limit the error accumulation caused by noisy sensors and erroneous inferences about user’s activities (e.g., walking vs. going on stairs), repeated corrections to the trail are needed.

Knowing the layout of the building provides the opportunity to make those corrections. Map Knowledge can provide the information to restrict the possible directions of the movement. For instance, walking on a corridor is typically done in a straight line and minor interferences of electronic equipments with the phone’s compass can be corrected. If compass deviation suddenly gets close to a right angle, the system infers that the user has left the corridor, either to go into a room or made a turn to another corridor. The closest door or corner is then associated with the user’s location.

Moreover, the map can also provide a location estimation based on the context of the user’s movements (as in [27]). Certain activities like going up or down stairs or taking an elevator can be performed in precise locations inside a building. An Activity Classifier can identify these instances (landmarks) and assist in correcting the location estimation. In order to identify these landmarks, we first classify the user’s movement. Most activities are performed similarly every time and their acceleration patterns can help to recognize them. Different orientations of the phone can still result in the same acceleration magnitude given by:

$$a = \sqrt{a_x^2 + a_y^2 + a_z^2} - g \quad (1)$$

where g is the Earth gravity, a_x , a_y and a_z represent the acceleration detected along the three orthogonal axes.

To cover for different situations of carrying a phone, we consider two most likely cases of carrying a phone, in hand and in pocket. While walking with the phone in hand has a simpler acceleration pattern (Figure 6(a)), walking with the phone in

⁵Cumulative Distribution Function




Fig. 6: Acceleration pattern when walking (raw acceleration with blue and filtered acceleration with red).

pocket transfers a lot of vibrations to the phone when the leg with the pocket steps (Figure 6(b)). However, for both cases, a step counting method with a suitable chosen threshold is used to detect steps and thus the walking speed. Step counting is also employed to determine the change of levels when using the stairs, Figure 7.




Fig. 7: Going down the stairs acceleration pattern.

Elevator movements present a specific pattern, with significant accelerations when the elevator starts and stops (Figure 8). The number of floors traveled is obtained from the time of the elevator movement (as in [28]).




Fig. 8: Elevator acceleration pattern.

We evaluate the activity classification component using sample acceleration patterns for different activities from two participants. These samples serve as training as well as test data for different classifier algorithms implemented in the well-known machine learning toolkit Weka⁶. In our evaluation, the training set consisted of 166 instances of activities from two participants annotated with the right activity. These activities were: standing, walking, going up on stairs, going down on stairs, going up by elevator, going down by elevator, opening and closing doors. All these for both cases with the phone in pocket and with the phone in hand. Using Weka's cross-validation option, we compared two window sizes of 128 and of 256 samples, testing all the classifiers available on Weka. Naive-Bayes had the best performance on a window size of 256 samples, which we later used in our evaluation.

⁶<http://www.cs.waikato.ac.nz/ml/weka/>

With an accuracy of 85.3% for Naive-Bayes, some activities were wrongly classified, e.g., walking was confused with going down the stairs 10% of the time. In practice, it is common for multiple activities to be captured within a single 256 sample window (3.2 seconds at sensor sampling frequency of 40Hz), so the rate of bad classifications may be higher. To prevent these wrong inferences from having a significant negative effect on location tracking estimation, we need the assistance of a separate component to be resilient to such wrongly classified activities. In our system, this is the WiFi fingerprinting component which is described next.

C. WiFi Fingerprinting Component

The WiFi fingerprinting localization component in Pazl can be seen as a stand alone localization solution but because of its limitations (e.g., not being able to localize in areas with no WiFi coverage) we have used it as an alternative and complement to the PDR component.

Based on a WiFi scan measurement from the phone, the vector of the five strongest APs is taken as the fingerprint and compared to all fingerprints in the training set. The closest matching fingerprints are selected using Euclidean distance (as in [29]). A weighted mean of the first three locations corresponding to the closest matching fingerprints is chosen as the estimated location from WiFi fingerprinting component.

The frequency of WiFi scans was chosen to be one scan every 20 seconds, which is a compromise between keeping the energy consumption low as each WiFi scan increases energy consumption on the phones and gathering enough data for the WiFi coverage map generation as well as assisting with the PDR estimation sufficiently often.

Following the above outlined approach for WiFi fingerprinting, we found that location estimations for some places in the building are more accurate than others. This is illustrated in Figure 9, which shows areas with low (green) and high (red) location errors, calculated as the distance between the estimated location and the ground truth.



Fig. 9: Spatial distribution of WiFi fingerprinting based location estimation errors on the floor plan.

Upon deeper examination, we find that this spatial difference in location estimation errors is linked to the spatial distribution of locations associated with the closest matching

fingerprints from the database. Greater resemblance of fingerprints between nearby areas results in a higher accuracy for the WiFi location estimation. The correlation between the WiFi fingerprinting based location estimation error and the *perimeter* (sum of the pairwise distance between the locations corresponding to three closest matching fingerprints in the database) is illustrated in Figure 10. We can therefore use the perimeter as a metric to judge the reliability of the location estimation from the WiFi fingerprinting component — lower the perimeter, lower the error. We rely on the perimeter metric to decide when to use the WiFi location estimation to assist in correcting the PDR via the particle filter as described below.

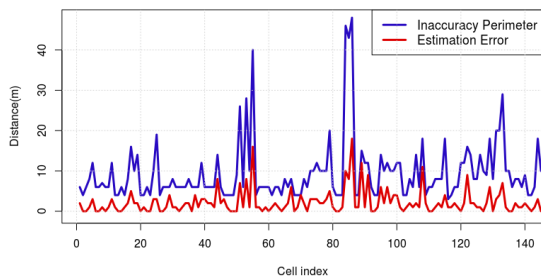


Fig. 10: Correlation between the estimation error and the perimeter of locations for the three closest matching fingerprints.

D. Particle Filter

Estimations provided by the Activity Classifier with the Map Knowledge and the WiFi component are integrated using a particle filter methodology.

Each particle is probabilistically progressed using PDR and its weight is adjusted based on the observations received from sensors, like compass indications, distance, probability of occurrence of different activities and confidence in the WiFi location estimation.

We observed that the direction (as given by compass indication) and distance travelled by the user suffer from small deviations that follow a Gaussian distribution, given by (2).

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(x-\mu)^2/2\sigma^2} \quad (2)$$

where, x is the chosen deviation and μ is the mean and σ the standard deviation of observed model.

Each particle performs its own independent PDR with its own sequence of probabilistically chosen activities, direction and distance traveled. In addition, whenever reliable WiFi fingerprinting based location estimation is available as indicated by lower perimeter metric, the weight of particles with PDR based location estimation that is close to the WiFi location estimation is increased.

E. Implementation

Pazl implementation consists of two parts: a mobile application that collects data from the phone’s sensors (including WiFi scans) and a server application that receives the

data to be processed for estimating the locations of WiFi scan measurements following the hybrid mechanism described above as well as for WiFi monitoring related visualization and analysis. The phone application is developed to run on a large variety of Android phones. To provide increased availability and concurrent access, the server application is developed to run on the cloud (Google App Engine in our implementation).

On the phone, acceleration, orientation and WiFi scans are collected only when the user is moving. When the phone is stationary, the compass and the radio interface are not used to save energy. Only the accelerometer is left on to run at a lower frequency just to sense when the user is moving again. Stored acceleration, orientation and WiFi measurements are uploaded opportunistically to the server: when the phone is charging, when WiFi access is available instead of 3G, or when upload is forced by the user.

All the WiFi samples are stored on the server together with the time and estimated location when they were obtained. For visualization of the coverage map, we use the Inverse Distance Weight based spatial interpolation [30]. Data is aggregated at a cell level of size 1 m^2 , by the median value if there are more measurements collected in the same cell. Selecting just a small set of WiFi samples based on the time when they were collected, dynamic reports can be generated, like the behavior of the network in a particular time period over several days or between different times within a day.

V. EVALUATION

A. Localization Accuracy

For our evaluation we assigned the task of collecting WiFi fingerprints for the training set of the WiFi localization component to two of our participants. They both contributed independently to collecting WiFi fingerprints from inside the building and annotating them with their exact location through a visual interface on the screen. Other systems like WiFi-SLAM [23] can automate this process, but we chose this approach to avoid the complexity of other systems and to have a higher confidence on the training set for the WiFi localization component that would serve all the other participants. Similarly, the activity classifier was trained with the sample acceleration patterns from two participants and the classifier was used to classify activities of all the other participants.

The experimental evaluation of the system involved 5 participants, all with Nexus One phones running Android 2.3. To evaluate the accuracy of the localization solution, we used the following experiment setup. A track of about 100m was chosen on the corridors with multiple (20) points, representing entrances to offices adjacent to a corridor, selected to offer the ground truth of our evaluation. Three participants walked on the track with the phone in hand and two with the phone in pocket. At the beginning of the track their time was synchronized with a clock and for every encounter of a ground truth position, the time was recorded. Location estimation errors were computed for each ground truth location as the Euclidean distance to Pazl’s location estimation.

The localization error of Pazl is presented in Figure 12. It can be seen that the accuracy for the case with the phone in pocket tends to be lower than the case with the phone in hand.

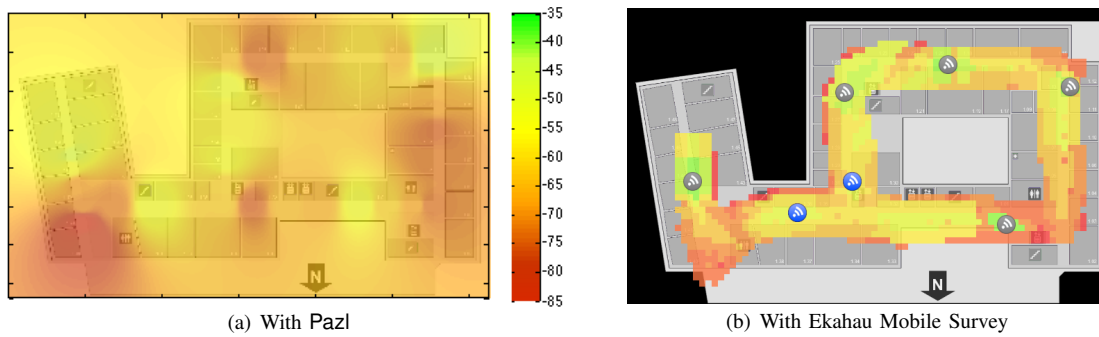


Fig. 11: WiFi coverage on a floor in dBm.




Fig. 12: CDF of location estimation errors.

This is because counting the number of steps with the phone in pocket is relatively a harder task.

B. Results Using Pazl and Comparison with Ekahau Mobile Survey

We evaluated Pazl in a small scale experiment in order to simulate a crowd-sourced site survey. The experiment was performed during a full working day (from 10am to 6pm), with 5 participants. They were asked to carry their Nexus One phones with them while moving freely during the day inside the building. No specific training was required beforehand other than just installing the Android app. We chose to focus only on a single floor in this analysis for ease of understanding, but participants were allowed to move between floors in the rest of the building using elevators and stairs as demanded by their day tasks.

The coverage maps obtained with Pazl are compared with the ones from using Ekahau Mobile Survey tool [10] in Figures 11 and 13. The Ekahau application shows the signal coverage only near the locations where measurements were collected, indicated with distinctive colors, representing different values of the Received Signal Strength (RSS). For the coverage representation using Pazl, we tried to keep the same color scheme as Ekahau to allow comparison between the two systems. Pazl estimates an extended coverage map via spatial interpolation for the entire floor plan, even for areas with no measurements. In the coloring scheme green indicates very good RSS, red indicates poor RSS, and other values of RSS are represented with a mixture of the two colors. Comparison

between the two systems can be done through color correlation or values comparisons in areas where they could both estimate the coverage, in particular on the corridors.

The coverage for the floor is shown in Figure 11 where we can observe that poor RSS was identified by both systems in the bottom left (Pazl indicated -78dBm , while Ekahau indicated -75dBm) and bottom right sides of the floor plan (Pazl indicated -66dBm , whereas Ekahau indicated -70dBm). We can also observe that both systems detect stronger RSS in almost the same places, in vicinity of APs. As for differences, Pazl estimated a region with low signal strength in the middle of the corridor, near the elevator, indicating -75dBm , whereas Ekahau recorded the signal strength in that area to be -65dBm .⁷

We also present the corresponding results for a specific AP in Figure 13, with very close match observed between the two systems. A good coverage of the AP is detected by both systems closer to where the AP is located and also on the corridor going top to bottom in the figures. Coverage is relatively worse along the other corridor going from right to left which we believe is because AP does not have a clear view of that corridor as it is occluded somewhat by the corner where the two corridors intersect.




Fig. 14: Channel usage distribution.

Channel usage distribution obtained with Pazl for a floor was compared with a manual site survey (Figure 14). The difference is between 1 and 3 APs per channel. This maybe because some of the APs located in other parts of the building

⁷Based on manual wireless site survey in that area, we observed that the max signal strength varies between -71 and -76dBm and that the nearest AP is shadowed by the corner of the wall. With only one run, this area may have witnessed direct line of sight when surveyed using the Ekahau tool.




Fig. 13: Coverage of an AP in dBm.

can be sensed only in specific areas, which might not have been reached by any of our participants over the period of the experiment.

We provided Pazl with physical location of campus WLAN APs in the building to estimate the coverage range (maximum distance of propagation) of APs that are seen from a floor (Figure 15⁸). These are not the exact maximum coverage range of APs because samples are limited to the areas traversed by participants. Still this experiment demonstrates that Pazl can detect problematic scenarios such as the APs that have unusual coverage. Analyzing a particular case, an AP located at the fifth floor was sensed at the first floor, over 55 meters away. This is due to the layout of the building which is mostly glass inside and has a large open area in the center. When we take this together with the fact that this fifth floor AP shares one of the heavily used 2.4GHz channel with other APs on the first floor, we have a scenario where channel allocation is poorly done risking interference related performance degradation from a user perspective.




Fig. 15: Coverage range of APs on a single floor obtained with Pazl.

VI. DISCUSSION

Recall that the experimental evaluation of the system involved all participants using the same type of phone. This is to avoid the problem of device calibration. With a single type of phone, all devices are expected to have similar WiFi scanning sensitivity so data was used exactly how it was sensed by the

⁸Note that this figure is different from Fig. 4 because the latter corresponds to the whole of the building

phone. There are solutions available for calibration between devices, like using kernel estimation with wide kernel widths to transform RSS from one device to another [31]. However, other factors such as the position of the phone relative to the human body may have similar or higher impact than RSS differences between different types of devices. In [32] differences larger than 10dBm between direct line of sight and human body shielding the phone are reported. Considering for all these aspects in a practical way is an open question.

Our experiment of one working day was just a proof of concept. We are planning to expand the evaluation to a longer period of time and involve a larger number of participants with their different Android phones. The goal is to identify interesting and dynamic aspects of WiFi coverage and interference as well as evaluate different incentives for people to participate. Note that there are further challenges to be addressed before a complete seamless crowdsensing based monitoring solution can be realized. For example, the localization component of a continual monitoring system needs to know when the user is making the transition from outdoors to indoors to start collecting samples and the other way around to stop collecting measurements.

VII. CONCLUSION

In this paper, we have highlighted the need for automated WiFi monitoring in indoor environments with multi-AP WLAN deployments via measurements. To address this need, we develop Pazl that leverages smartphones used by occupants of these environments and their natural mobility to realize a low cost and automated mobile crowdsensing based WiFi monitoring system. Pazl is built upon a novel hybrid localization mechanism that combines pedestrian dead reckoning and WiFi fingerprinting to locate each WiFi measurement obtained from a participating smartphone. Evaluation of Pazl showed that it yields results similar to state of the art wireless site survey tools but in a significantly more automated manner. It also captures the more valuable user-side perspective. Our future work will focus on a more thorough evaluation of Pazl over a longer period with more users capturing dynamic aspects of WiFi networks and in different environments. Other aspects for future work include incentive mechanisms to encourage user participation and broadening the monitoring beyond WiFi to consider cellular and bluetooth coverage in indoor settings.

ACKNOWLEDGMENTS

This work was supported in part by a Cisco Research Award. We thank Jiwei Li and Noelia Hernandez for their help with the data collection, and volunteers who took part in crowdsourcing experiments for their time. We also thank Richard Mortier for his contribution at the early stages of this work.

REFERENCES

- [1] N. D. Lane et al. A Survey of Mobile Phone Sensing. *IEEE Communications*, 48(9):140–150, Sep 2010.
- [2] M. Srivastava, T. Abdelzaher, and B. Szymanski. Human-centric sensing. *Philosophical Transactions of the Royal Society*, 370(1958):176–197, Jan 2012.
- [3] Y. Xiao, P. Simoens, P. Pillai, K. Ha, and M. Satyanarayanan. Lowering the Barriers to Large-Scale Mobile Crowdsensing. In *Proc. ACM HotMobile*, 2013.
- [4] R. K. Ganti, F. Ye, and H. Lei. Mobile Crowdsensing: Current State and Future Challenges. *IEEE Communications*, 49(11):32–39, Nov 2011.
- [5] D. Hasenfratz, O. Saukh, S. Sturzenegger, and L. Thiele. Participatory Air Pollution Monitoring Using Smartphones. In *Proc. 2nd International Workshop on Mobile Sensing*, 2012.
- [6] E. Kanjo. NoiseSPY: A Real-Time Mobile Phone Platform for Urban Noise Monitoring and Mapping. *Springer Mobile Networks and Applications (MONET) Journal*, 15(4):562–574, Aug 2010.
- [7] A. Gember, A. Akella, J. Pang, A. Varshavsky, and R. Caceres. Obtaining In-Context Measurements of Cellular Network Performance. In *Proc. ACM Internet Measurement Conference (IMC)*, 2012.
- [8] Open Signal Inc. <http://opensignal.com/>.
- [9] K. Tan et al. From MAP to DIST: The Evolution of a Large-Scale WLAN Monitoring System. *IEEE Transactions on Mobile Computing*, 2012.
- [10] Ekahau Mobile Survey. <http://www.ekahau.com/products/ekahau-mobile-survey/mobile-survey-overview.html>.
- [11] R. Zhou, G. Xing, X. Xu, J. Wang, and L. Gu. WizNet: A ZigBee-Based Sensor System for Distributed Wireless LAN Performance Monitoring. In *Proc. IEEE PerCom*, 2013.
- [12] P. A. Frangoudis and G. C. Polyzos. Report-based topology discovery schemes for centrally-managed Wi-Fi deployments. In *Proc. 6th EURO-NF Conference on Next Generation Internet (NGI)*, 2010.
- [13] IEEE Std 802.11 - 2012. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, 2012.
- [14] J. Geier. *Designing and Deploying 802.11n Wireless Networks*, chapter 15. Cisco Press, 2010. *Performing a Wireless Site Survey*.
- [15] K. Kim, A. W. Min, and K. G. Shin. Sybot: An Adaptive and Mobile Spectrum Survey System for WiFi Networks. In *Proc. ACM MobiCom*, 2010.
- [16] R. Harle. A Survey of Indoor Inertial Positioning Systems for Pedestrians. *IEEE Communications Surveys & Tutorials*, 2013.
- [17] I. Constandache, R. R. Choudhury, and I. Rhee. Towards Mobile Phone Localization without War-Driving. In *Proc. IEEE INFOCOM*, 2010.
- [18] M. Youssef, M. A. Yosef, and M. El-Derini. GAC: Energy-Efficient Hybrid GPS-Accelerometer-Compass GSM Localization. In *Proc. IEEE GLOBECOM*, 2010.
- [19] P. Bahl and V. N. Padmanabhan. RADAR: An In-Building RF-Based User Location and Tracking System. In *Proc. IEEE INFOCOM*, 2000.
- [20] M. Youssef and A. K. Agrawala. The Horus WLAN Location Determination System. In *Proc. MobiSys*, 2005.
- [21] P. Bolliger. Redpin – Adaptive, Zero-Configuration Indoor Localization through User Collaboration. In *Proc. ACM MobiCom MELT Workshop*, 2008.
- [22] J. Park et al. Growing an Organic Indoor Location System. In *Proc. MobiSys*, 2010.
- [23] B. Ferris, D. Fox, and N. Lawrence. WiFi-SLAM Using Gaussian Process Latent Variable Models. In *Proc. IJCAI*, 2007.
- [24] A. Rai et al. Zee: Zero-Effort Crowdsourcing for Indoor Localization. In *Proc. ACM MobiCom*, 2012.
- [25] H. Wang et al. No Need to War-Drive: Unsupervised Indoor Localization. In *Proc. MobiSys*, 2012.
- [26] Y. Kim, H. Shin, Y. Chon, and H. Cha. Smartphone-based Wi-Fi Tracking System Exploiting the RSS Peak to Overcome the RSS Variance Problem. *Elsevier Pervasive and Mobile Computing*, 9(3), Jun 2013.
- [27] D. Gusenbauer, C. Isert, and J. Krsche. Self-Contained Indoor Positioning on Off-the-Shelf Mobile Devices. In *Proc. IEEE IPIN*, 2010.
- [28] H. Ye et al. FTrack: Infrastructure-free Floor Localization via Mobile Phone Sensing. In *Proc. IEEE PerCom*, 2012.
- [29] Y. Shang, W. Ruml, Y. Zhang, and M. P. J. Fromherz. Localization from Mere Connectivity. In *Proc. ACM MobiHoc*, 2003.
- [30] D. Shepard. A Two-Dimensional Interpolation Function for Irregularly-Spaced Data. In *Proc. 23rd ACM National Conference*, 1968.
- [31] J. Park et al. Implications of Device Diversity for Organic Localization. In *Proc. IEEE INFOCOM*, 2011.
- [32] S. Sen, R. R. Choudhury, and S. Nelakuditi. SpinLoc: Spin Once to Know Your Location. In *Proc. ACM HotMobile*, 2012.