SpinLoc: Spin Once to Know Your Location

Souvik Sen
Duke University
souvik.sen@duke.edu

Romit Roy Choudhury
Duke University
romit.rc@duke.edu

Srihari Nelakuditi
University of South Carolina
srihari@cse.sc.edu

ABSTRACT

The rapid growth of location-based applications has spurred extensive research on localization. Nonetheless, indoor localization remains an elusive problem mostly because the accurate techniques come at the expense of cumbersome war-driving or additional infrastructure. Towards a solution that is easier to adopt, we propose *SpinLoc* that is free from these requirements. Instead, SpinLoc levies a little bit of the localization burden on the humans, expecting them to rotate around once to estimate their locations. Our main observation is that wireless signals attenuate differently, based on how the human body is blocking the signal. We find that this attenuation can reveal the directions of the APs in indoor environments, ultimately leading to localization. This paper studies the feasibility of *SpinLoc* in real-world indoor environments using off-the-shelf WiFi hardware. Our preliminary evaluation demonstrates accuracies comparable to schemes that rely on expensive war-driving.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Wireless communication

General Terms

Design, Experimentation, Performance

Keywords

Wireless, Localization, Cross-Layer, Application

1. INTRODUCTION

Despite numerous research efforts [1–8], indoor localization is still not a mainstream technology. We believe that the main hurdle lies in most of them requiring careful war-driving. Crowdsourcing this operation [9] is an attractive option, but unlikely to be adopted broadly since many users may not be willing to report their signal strength measurements to a localization server. Moreover, due to the lack of GPS in indoor settings, gathering ground-truth is also hard. SecureAngle [10], an innovative PHY layer technique proposed recently, computes the direction of a WiFi-capable device with respect to an AP. While certainly appealing, the timeframe for ubiquitously installing special APs – with 8 or more antennas – is not a short-term proposition. Nonetheless, we believe this is the right direction to approach indoor localization, and propose to approximate it today, with off-the-shelf hardware, zero infrastructure, and absolutely no war-driving. The only tradeoff is that the user needs to make a slight effort – a spin – every time she needs her location. We explain this through an example, followed by the technical underpinnings.

Consider a shopping mall where the user intends to localize herself. With SpinLoc installed on her phone, she turns on the application, and makes a 360° rotation at her current location. Using the signals recorded during the rotation, and the already-known AP locations, SpinLoc computes the location of the user (detailed later). The location is marked on the floorplan of the mall and displayed on the phone screen. Other than making the floorplan and AP locations available to an Internet database, the mall authorities are not expected to make any investment – no infrastructure installation; no war-driving. We show that in such settings, SpinLoc can offer localization accuracies in the order of $6.5m$ with 4 audible APs in the vicinity. With more APs, the median accuracy can improve upto $5m$.

The technical underpinning of SpinLoc is actually quite simple. Without loss of generality, consider an AP located in the westward direction. When the user spins at her location, at some point the phone is between the AP and the user's body, and at a different instance, the body lies between the AP and the phone (Figure 1). Given that the human body is a significant attenuator of WiFi signals (in the 2.4 and $5GHz$ frequencies), the signals arriving at the phone differ significantly between these two configurations. In particular, when the phone is between the AP and the user (Figure 1(a)), the direct path from the AP to the phone is strong. However, when the user's body lies between the AP and the phone (Figure 1(b)), the direct path is severely blocked, resulting in large attenuation. By recording the compass direction at which this attenuation is maximum, it is feasible to infer the AP's direction. With multiple APs in the vicinity, the direction to each AP can be computed from the same spin. The knowledge of all AP directions permits triangulation, ultimately yielding an estimate of the user's location.

Importantly, SpinLoc does not use received signal strength (RSSI) as reported by WiFi cards. While RSSI may be somewhat appli-

HotMobile'12, February 28–29, 2012, San Diego, California, USA.

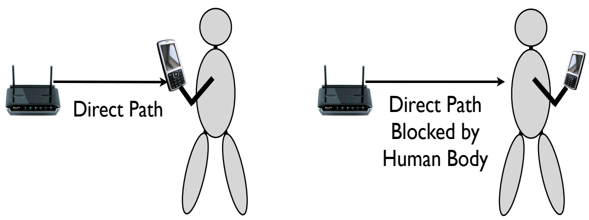

Figure 1: User orientation w.r.t AP: (a) facing (b) blocked

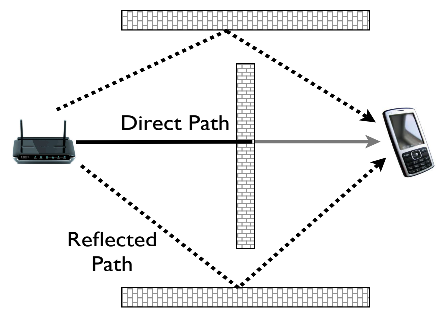

Figure 2: Transmitted signal travels through multiple paths before reaching the receiver

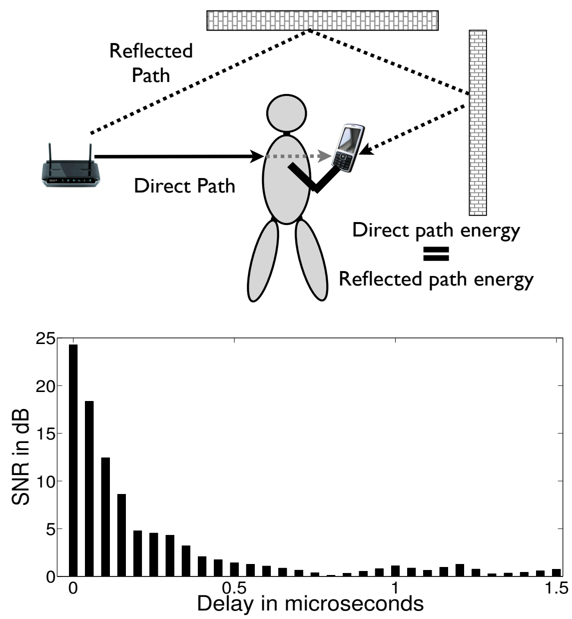

Figure 3: (a) Spinning will not offer the AP direction if energy on both signal paths are added. (b) Power delay profile of an indoor transmission.

cable in outdoor environments [11], rich multipath in indoor environments derails RSSI-based approaches. To circumvent this problem, SpinLoc relies only on the signal strength of the direct signal path, i.e., the signal component that traverses along the straight line joining the AP and the mobile device. Fortunately, this information can be extracted from the *power-delay profile* of a link, a physical layer information that is exported by the Intel 5300 card. Thus, if designed well, SpinLoc can be a candidate for near-future deployment using off-the-shelf hardware.

Of course, building such a system to cope with real-world scenarios entails a range of technical and social challenges: (1) The phone's compass error may be significant – how does that affect SpinLoc's accuracy? (2) The attenuation of the direct signal path may vary with other humans in the environment – how can SpinLoc cope with such variations? (3) Will the idea work even when the direct signal path is weak? This paper addresses these questions and presents promising evidence to justify deeper investigation. The proof-of-concept is built on Dell laptops, Android phones, and Cisco APs, and achieves localization accuracy of 5 to 12m in a university cafe and an engineering building. Even when the compass errors are large (upto 30°), the accuracy does not degrade more than 12m, so long as there are 5 APs within communication range. Finally, the energy footprint of the system is small, suggesting real-world viability.

2. INTUITION AND MEASUREMENTS

We begin this section with a brief background on wireless multipath propagation, followed by our key hypothesis and initial measurement-based verification.

2.1 Background

Wireless signal propagation is similar to light. A transmitted signal scatters in all radial directions and reflects on different surfaces, including walls, furnitures, etc. Hence, in addition to a direct path from the transmitter to the receiver, copies of the same signal arrive through many reflected paths, each with a different delay and attenuation. The wireless radio combines these multipath copies, and ultimately extracts the information embedded in the signal. Figure 2 illustrates 3 example signal paths from the transmitter to the receiver.

Among all the multipath copies, we define the *direct path* as the straight line joining the transmitter and the receiver. To understand when the human is precisely between the AP and the phone, SpinLoc must track *only* the direct path signal. Otherwise, if SpinLoc uses the union of all signal components (as is the case with RSSI), it would be difficult to identify when the human has blocked the signal. Figure 3(a) explains this next with an example.

In Figure 3(a), assume that the two signal components have been equally attenuated, one due to absorption by the human, and the other due to multiple reflections. Now in both configurations (i.e., the human is on the left of the phone and blocking the direct path, or on the right of the phone and blocking the reflected path), the sum of the incident energies will be identical, making it difficult to infer the AP's direction. However, if only the direct path signal is used, one might expect a drop when the human is on the left of the phone, but not when she is on the right. This motivates the need to *only* use the direct path signal. Unfortunately, today's WiFi interfaces do not provide the individual signal components from which we can pick the desired signal component. RSSI, readily available from almost all interfaces, is the sum of energies over all signal components, and thereby, unreliable in multipath-rich indoor environments.

In search of a mechanism to extract the direct path signal, we learnt that the Intel 5300 WiFi card exports some physical layer information, that can be translated to the *power-delay profile (PDP)*. Loosely, the PDP captures the amount of energy incident on the receiver at different delays. Since the direct path arrives quicker at the receiver than all other reflected paths, we find that pick-

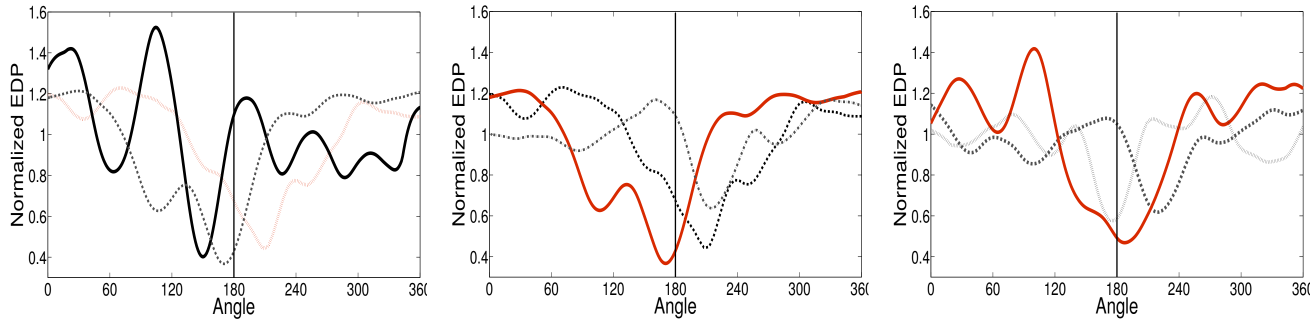

Figure 4: Measured EDP across user orientation when (a) AP is visible to phone (b) AP is behind a wall but close to phone (c) AP is behind a wall but far away from phone. Different curves correspond to different experiments.

ing the least-delay value of the PDP essentially provides us with the energy of the direct path. Figure 3(b) shows the PDP of an indoor transmission, where the AP was visible to the laptop. Since the direct path does not pass through obstructions in this case (and thus gets less attenuated), it yields the strongest signal component. While this is typical, it is certainly possible that a wall obstructs the direct path, making it weaker than other reflected paths. Importantly, SpinLoc is not sensitive to the relative energy of the direct path against that of other paths; instead it focuses only on the absolute *energy of the direct path,* denoted EDP. Comparing EDP across different configurations (during a spin) will help in revealing the AP's direction.

2.2 SpinLoc: Hypothesis

We summarize the SpinLoc intuition as follows. When a human is present between an AP and her own mobile device (as in Figure 1(b)), her body attenuates the direct signal path from the AP. This is because the human body with high water content has been shown to be a significant absorber of (2.4 and 5 GHz) WiFi signals [11]. Now, when the human turns 180° from this orientation (Figure 1(a)), her phone is located between her body and the AP, and is not subject to the attenuation. As a generalization of this, we present the following hypothesis: if a user rotates 360° at her own position, the direction that exhibits minimum energy for the direct path (EDP) is the direction opposite to the AP. If such directions can be computed for at least 3 APs, then triangulation is feasible, ultimately yielding the user's location.

We verify our hypothesis using measurements from off-the-shelf Intel 5300 cards. This card exposes per-subcarrier channel frequency response (CFR) to the user – an inverse fast fourier transform (IFFT) of the CFR outputs the power delay profile (PDP). We obtain the energy of the direct path (EDP) from the PDP, and track its variation as the user spins in her location. Three important questions are of interest. (1) Does minimum EDP accurately yield the AP's direction. (2) Does the presence of additional humans in the vicinity affect our hypothesis? (3) Can RSSI be used to also infer the direction of the AP? The following measurements are designed to answer these questions.

2.3 Measurement and Verification

Our experiments are performed in a relatively busy engineering building, with faculty offices and classrooms. To simultaneously measure the PDP and the user's compass orientation, we taped a Google NexusOne phone to a laptop (this is because we did not find any device that has the Intel 5300 card and a compass). While holding this laptop-phone device, we ask a user to rotate 360° at her location. On average, a rotation lasts around 10 seconds. The device is made to receive approximately 100 packets per second and record the *energy of the direct path (EDP)* for each received packet. We average the EDP over all packets received in a given orientation. Then, to cope with fast fading, we smoothen the series of per-orientation EDP by using a simple moving average (discussed later).

We begin with an experiment where the AP is visible to the user – this implies the existence of a strong direct path when the user faces the AP. Figure 4(a) plots the EDP as a function of the user's orientation with respect to this AP (the 3 curves are from 3 distinct locations). Assuming the AP's direction to be the 0° reference, the EDP should ideally be minimum at 180°. Evident from the graph in Figure 4(a), the EDP dip is indeed close to 180. A pertinent question is whether this technique holds even when the direct path is not as strong (such as when it passes through an obstruction). To this end, we place the AP behind a wall that blocks the direct path between the AP and the mobile device. Figure 4(b) shows a consistent behavior even in this scenario – the maximum EDP dip is still close to 180°. In a subsequent experiment, we keep the AP behind the wall and move the user far away from the AP, forcing the WiFi signal to be weaker. Still, the EDP dips around 180 although the dip is less sharp (Figure 4(c)). We find consistent results over multiple other experiments, suggesting promise with SpinLoc.

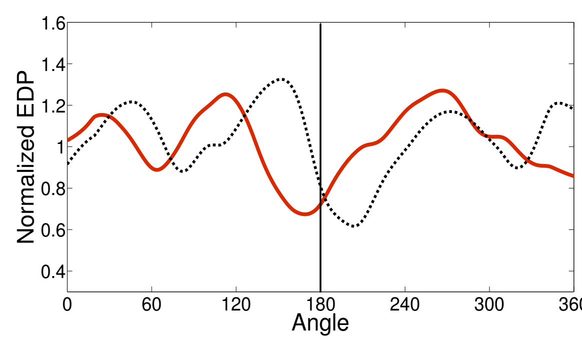

Figure 5: Measured direct path energy across user orientation in presence of another blocking human.

Effect of other humans: We next investigate if the presence of other humans in the vicinity derails SpinLoc. For this, we perform a controlled experiment. We position a second human on the direct path between the AP and the device user – the gap between the two humans is $2m$. This is expected to reduce EDP even when the user is facing the AP. Figure 5 plots the variation of EDP when the user rotates. Observe that although the EDP dip is less sharp, the minimum value is still near 180°. This suggests SpinLoc's robustness to humans in the environment.

Why not use received signal strength (RSSI)? Previous work has shown that humans can attenuate the RSSI of a signal by blocking it [11], and this can be used in *outdoor* environments to estimate the AP's direction. However, this observation does not extend to indoor environments, where wireless propagation is heavily dominated by multipath. This is because RSSI can be approximated as the sum total of energy over all the signal paths. As explained in Section 2.1, the amount of energy blocked by the human in different orientations can be the same, resulting in no clear dip (or multiple dips). Figure 6 captures this behavior – when the user spins, the RSSI dips are non-unique, and often happen far away from the ideal 180° orientation.

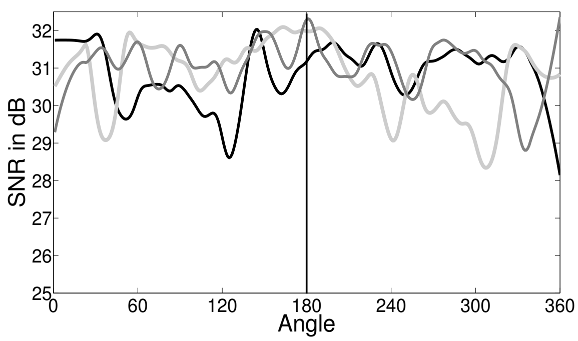

Figure 6: Measured RSSI across user orientation.

3. SYSTEM DESIGN

Translating SpinLoc's high level idea into a functional prototype entails two tasks: (1) How to find the direction to an AP as precisely as possible? (2) How to localize a mobile device with imprecise direction information? Of course, the simplicity of design is vital for SpinLoc because the entire operation must be executed on the mobile device – we assume no reliance on any localization server.

3.1 Finding AP direction with SpinLoc

We observe that as the user spins, her body gradually blocks and subsequently unblocks the direct signal path from the AP to the device. Even when the user is at 90° from the AP, the direct path signals may still be partially blocked, perhaps by the user's arms or shoulders. Consequently, the energy on the direct path (EDP) will decrease and increase, forming crests and troughs, as shown in Figures 4. We exploit the troughs to correctly identify the AP direction. A naive approach might be to find the angle corresponding to the minimum EDP and declare the *opposite angle* as the correct AP direction. But this approach may not be robust in the presence of fast fading and measurement noise. The direct path signal may combine with other signals in the air (from other interfering transmissions), causing its energy to fluctuate instantaneously even without the human obstacle. However, averages over multiple packets can be expected to eliminate these fluctuations. Therefore, we perform a moving average on the sequence of EDPs, much like a low pass filter. Figure 7 shows the effect – the dashed curve shows the raw EDP variations while the solid line shows the same variation after filtering. Clearly, filtering makes the blocking effects easier to recognize. SpinLoc now declares the angle corresponding to the minimum EDP as the angle opposite to the AP.

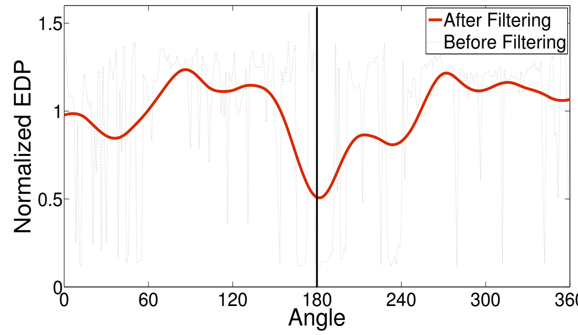

Figure 7: (Un)Filtered EDP w.r.t. user's orientation.

3.2 Localization using Angle Information

Once the angle to each of the APs is estimated – *note that SpinLoc estimates all these angles in one spin* – SpinLoc determines the user's location using triangulation. Let us denote the estimated angle from the phone to AP_i as θ_i. For triangulation, we draw a line from AP_i along the direction of $((\theta_i + 180)\ mod\ 360)$; this is the opposite direction of θ_i. Denote this line as L_i (Figure 8). SpinLoc then computes the intersection points of all pairs of $< L_i, L_j >, i \neq j$. The centrioid of these intersecting points is declared as the estimated location of the device.

We tune this method as follows. We find that SpinLoc's angle estimation accuracy reduces at weaker signal strengths, and hence, we choose only relatively strong APs (20dB or stronger) for localization. Furthermore, if two APs are located at nearly the same direction, their intersection point is likely to be far away from the mobile's actual location (in the extreme case, if the two APs are aligned, their intersection point will be located at infinity). To remove such outliers, SpinLoc uses the stronger of the two APs for localization when their estimated angles differ by less than 20°. Figure 8 illustrates the overall process.

Leveraging RSSI information: We explore if using the RSSI information can benefit SpinLoc. Although RSSI is a crude indicator of distance, our hypothesis is that it may be beneficial in conjunction with reasonably good angular information. Thus, based on the recorded RSSI, we estimate the distance between the device and AP_i as D_i[1]. Now for each AP_i, we plot a point that is located D_i distance away in the direction of $(\theta_i + 180)\ mod\ 360$. SpinLoc then computes the centroid of these points as the estimate of the device's location. We evaluate the performance in the next section.

3.3 Points of Discussion

The locations of APs within a building (such as a mall or museum) have to be made available to SpinLoc – is this realistic? We believe that any indoor localization system will need the floorplan to provide a semantic meaning to the computed location. If the

[1]We use standard pathloss equations, with pathloss exponent of 3 for indoor environments.

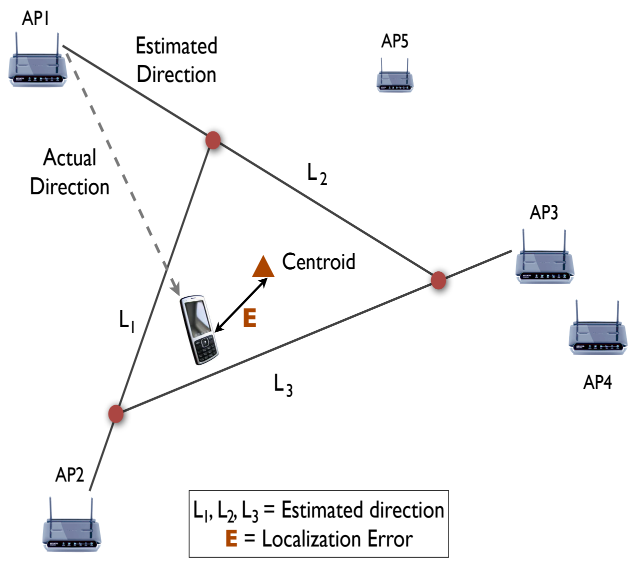

Figure 8: Illustration of SpinLoc's localization procedure: Only AP1, AP2, AP3 used for triangulation. AP4 is eliminated because it has a similar direction as AP3. AP5 is a weak link and can result in a large estimation error.

mall administration is willing to extend the floorplan, the AP locations may be easy to add.

SpinLoc is reactive because the user invokes localization – can SpinLoc be proactive? With the phone compass always on, it might be possible to track naturally occurring rotations of the user, when she turns corners or makes about turns. SpinLoc may then be able to deduce the direction of a subset of APs, and combine with some degree of dead-reckoning to estimate the user's location. The viability of such a scheme is one of our main topics of investigation.

To get beacons from all the APs, does SpinLoc require the APs to be on the same channel? This is not necessary because every time the user invokes SpinLoc, the WiFi interface in the device can perform a channel scan. This will permit the device to receive beacons from all APs. Of course, the indoor space would need to have at least 3 APs in the audible range, which we believe is quite common.

4. PERFORMANCE EVALUATION

Prototype and Experimental Setup: As mentioned earlier, we implement SpinLoc using a laptop with an Intel 5300 wireless card and a Google Nexus One phone. The phone is time synchronized and physically attached to the laptop – it records the compass orientation and sends it to the laptop. The laptop receives 100 small beacon packets per second from Cisco E4200 APs, operating at $40MHz$ on the $2.4GHz$ band. The user spins carrying the laptop-phone module in her hand. We evaluate SpinLoc across 55 locations in two environments: (1) engineering building with offices and classrooms and (2) a university cafe. In the engineering building, we experiment with $6APs$ at 30 locations. The university cafeteria is relatively smaller; we deploy 4 APs and report results from 25 locations. We covered approximate areas of 1000 m^2 and 800 m^2 respectively in these buildings.

Angle Determination Accuracy: At each of the 55 locations, SpinLoc estimates the angle of every audible AP[2]. Since we know the ground truth for each of these APs, we plot SpinLoc's angle estimation error as a CDF in Figure 9(a). Evidently, the mean is less than 20°, but for around 20% of the cases, the errors can be as high as 40 – 60°. We postulate that the high errors are due to weaker links. To investigate this further, we plot the average angle error as a function of link SNR in Figure 9(b). The figure shows that the angular error indeed decreases with increasing SNR. The reason is that weak links may not have a significant direct signal path and hence less likely to exhibit a sharp EDP dip, even when the user blocks the signal. This led us to exclude weak APs in the design of SpinLoc.

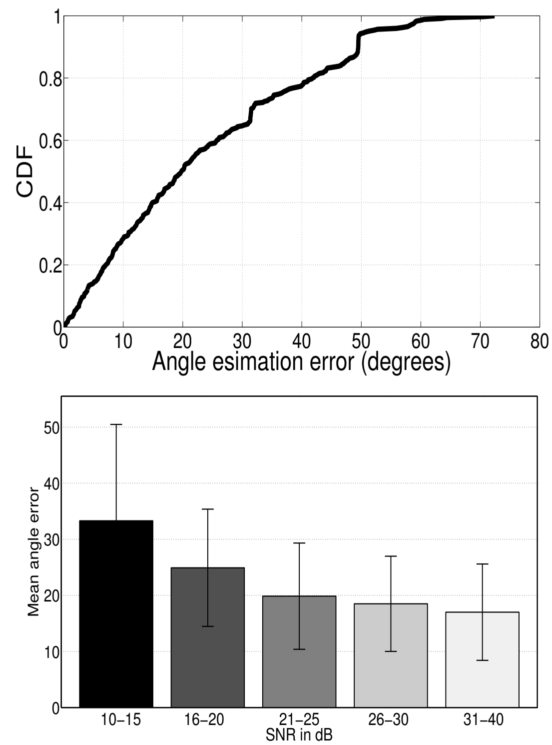

Figure 9: (a) SpinLoc angle estimation error (b) Error decreases with increasing SNR, stronger the AP better the accuracy.

Localization Accuracy: The above results suggest that links stronger than $20dB$ on average, have less than 20° angle error. Hence, for better accuracy, SpinLoc only uses APs that meet this criteria. Figure 10(a) plots the CDF of localization error across 55 locations. The median localization accuracy is 7.2 meters. Figure 10(a) also shows the benefit of leveraging RSSI – the median accuracy improves to within 5 meters. Both the approaches outperform simple RSSI based triangulation which has a median accuracy of 14.7 meters (Figure 10(a)). Figure 10(b) shows that the accuracy improves with increasing number of (strong) APs at a given location. Considering that SpinLoc does not need anything else other than the APs' location, we believe these results may be deemed promising. Of course, conclusive results about SpinLoc's accuracy will require far more extensive evaluation.

[2] We use the term angle and direction interchangeably.

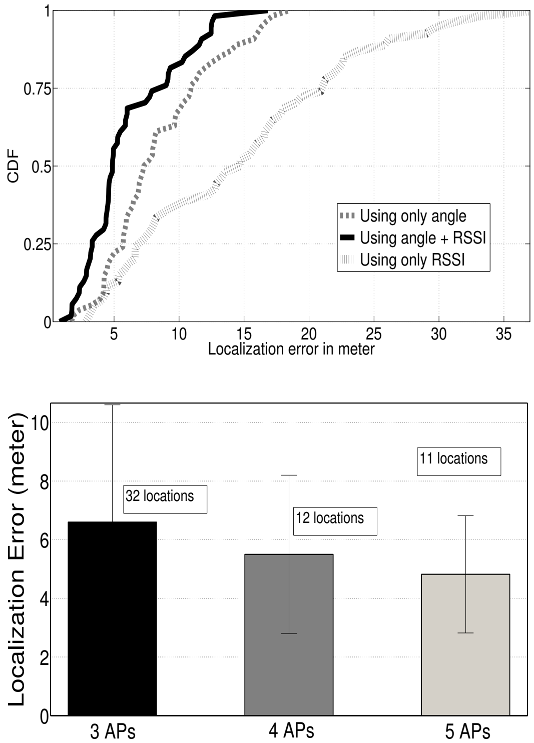

Figure 10: (a) SpinLoc localization accuracy across 55 spots (b) Accuracy w.r.t number of APs per location.

5. LIMITATIONS AND NEXT STEPS

We discuss a few concerns with SpinLoc, and potential ways to alleviate them.

Will SpinLoc need frequent spins for navigation? SpinLoc trades off wardriving and infrastructure for some user involvement. Whether this is acceptable to users is likely to depend on how frequently they need to spin, say within a mall. We believe it is possible for SpinLoc to combine naturally occurring turns and spins of users, with other direction and distance estimation methods, to reduce the need for frequent spins. Our ongoing work is directed towards a *spin-in-the-worst-case* type of approach.

Will SpinLoc consume substantial energy? SpinLoc neither needs to download signal maps, nor does it require CPU-intensive matching operations. In this regard, the energy consumption is likely to be quite low. However, if the region is sparse in WiFi APs, the WiFi channel scanning operation may consume some energy. However, SpinLoc could stop scanning once it has discovered the requisite number of APs. We plan to investigate the energy implications in greater detail in future work.

6. RELATED WORK

A wide variety of approaches have been proposed for indoor localization each incurring a different form of overhead. RF signal strength-based localization schemes such as RADAR [1], Horus [2] and PinLoc [12] perform detailed site surveys a priori to generate WiFi based location fingerprints. Place Lab [3] and Active Campus [4] attempt to reduce the overhead of calibration, coupling information from WiFi and GSM base stations. Time-based techniques such as PinPoint [5], and TPS [6] utilize time delays in signal propagation to estimate distances between wireless transmit-receiver pairs. The Cricket system [7, 8] utilizes ultrasound and RF signals, requiring ultrasound detectors on mobile devices for localization, limiting its applicability.

Angle-of-arrival based techniques utilize multiple antennas to estimate the angle at which signals are received, and then geometrically localize devices [10, 13]. These techniques require quite sophisticated systems of 4 to 8 antennas and non-trivial signal processing capabilities, unlikely on mobile devices in the near future. Borealis [11] attempts to find the direction of a rogue AP by rotating a smartphone around a signal blocking obstacle. They rely on RSSI only and hence are limited to outdoor environments. SpinLoc's ability to utilize PHY layer information from off-the-shelf WiFi cards for effective indoor localization, makes it a candidate for immediate adoption.

7. CONCLUSION

This paper explores the feasibility of localizing a device by deliberately inserting blockages in wireless signal reception. If a user spins at her current location, we find that the direction of the AP can be determined with a median error of 20°. When combined with RSSI information, the location accuracy can reach almost $5m$ in dense WiFi conditions. While today's best indoor localization schemes may be comparable (or slightly better), they come with the overheads of war-driving, additional infrastructure, or heavy computation. We believe SpinLoc may be a simple and alternative approach, perhaps more suited to near-term deployment in indoor environments.

8. REFERENCES

[1] V. Bahl et al. RADAR: An in-building rf-based user location and tracking system. In *INFOCOM*, 2000.

[2] M. Youssef and A. Agrawala. The horus WLAN location determination system. In *MobiSys*, 2005.

[3] Yu-Chung et al. Accuracy characterization for metropolitan-scale wi-fi localization. In *MobiSys*, 2005.

[4] William G. Griswold et al. Activecampus: Experiments in community-oriented ubiquitous computing. *Computer*, 2004.

[5] M. Youssef et al. Pinpoint: An asynchronous time-based location determination system. In *ACM Mobisys*, June 2006.

[6] X. Cheng et al. TPS: A time-based positioning scheme for outdoor sensor networks. In *IEEE Infocom*, March 2004.

[7] N. Priyantha, A. Chakraborty, and H. Balakrishnan. The cricket location-support system. In *MOBICOM*, 2000.

[8] A. Smith, Balakrishnan, et al. Tracking moving devices with the cricket location system. In *ACM Mobisys*, June 2004.

[9] K. Chintalapudi, A. Iyer, and V. Padmanabhan. Indoor localization without the pain. In *MOBICOM*, 2010.

[10] J. Xiong and K. Jamieson. SecureAngle: improving wireless security using angle-of-arrival information. In *ACM HotNets*, 2010.

[11] Z. Zhang et al. I am the antenna: Accurate outdoor ap location using smartphones. *ACM MOBICOM*, 2011.

[12] S. Sen et al. Precise Indoor Localization using PHY Layer Information. In *ACM HotNets*, 2011.

[13] D. Niculescu and B. Nath. VOR base stations for indoor 802.11 positioning. In *ACM MobiCom*, September 2004.