

I Am the Antenna: Accurate Outdoor AP Location using Smartphones

Zengbin Zhang[†], Xia Zhou[†], Weile Zhang^{†§}, Yuanyang Zhang[†], Gang Wang[†],
Ben Y. Zhao[†] and Haitao Zheng[†]

[†]Department of Computer Science, UC Santa Barbara, CA 93106, USA

[§]School of Electronic and Information Engineering, Xian Jiaotong University, P. R. China
{zengbin, xiazhou, yuanyang, gangw, ravenben, htzheng}@cs.ucsb.edu, zwllwz@stu.xjtu.edu.cn

ABSTRACT

Today’s WiFi access points (APs) are ubiquitous, and provide critical connectivity for a wide range of mobile networking devices. Many management tasks, *e.g.* optimizing AP placement and detecting rogue APs, require a user to efficiently determine the location of wireless APs. Unlike prior localization techniques that require either specialized equipment or extensive outdoor measurements, we propose a way to locate APs in real-time using commodity smartphones. Our insight is that by rotating a wireless receiver (smartphone) around a signal-blocking obstacle (the user’s body), we can effectively *emulate the sensitivity and functionality of a directional antenna*. Our measurements show that we can detect these signal strength artifacts on multiple smartphone platforms for a variety of outdoor environments. We develop a model for detecting signal dips caused by blocking obstacles, and use it to produce a *directional analysis* technique that accurately predicts the direction of the AP, along with an associated confidence value. The result is Borealis, a system that provides accurate directional guidance and leads users to a desired AP after a few measurements. Detailed measurements show that Borealis is significantly more accurate than other real-time localization systems, and is nearly as accurate as offline approaches using extensive wireless measurements.

Categories and Subject Descriptors

C.2 [Computer-Communication Networks]: Network Architecture and Design

General Terms

Design, Experimentation, Performance

Keywords

Access point location, WiFi, smartphones

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

MobiCom’11, September 19–23, 2011, Las Vegas, Nevada, USA.

Copyright 2011 ACM 978-1-4503-0492-4/11/09 ...\$10.00.

1. INTRODUCTION

WiFi networks today are ubiquitous in our daily lives. WiFi access points (APs) extend the reach of wired networks in indoor environments such as homes and offices, and enable mobile connectivity in outdoor environments such as sports stadiums, parks, schools and shopping centers [1, 2]. Even cellular service providers are relying on outdoor WiFi APs to offload their 3G traffic [2, 3]. As Internet users become reliant on these APs to connect their smartphones, tablets and laptops, the availability and performance of tomorrow’s networks will depend on well tuned and managed access points.

A critical part of managing access points is the ability to locate individual access points based on their signal [16, 19, 23]. Doing so allows network administrators to identify APs causing excessive interference to others, or unauthorized APs that may provide easy entry for malicious attackers [22]. For individual users, it allows them to locate and get closer to neighborhood WiFi hotspots, build war-driving databases, or pinpoint neighboring APs so they can better position their own basestations.

Unfortunately, much of this is not possible today, because current techniques to locate outdoor WiFi access points require either extensive wardriving measurements, followed by significant offline computation [11], or complex hardware components such as steerable directional phased array antennas [28]. For the wardriving scenario, locating each AP requires measurements from a large number (35+) of locations, making the process extremely time and energy intensive. For the directional antenna solution, the specialized hardware components cost several thousand dollars each, which clearly limits their availability to only a small portion of system administrators for large enterprises.

In this work, we ask the question: “*is there a cost- and time-efficient alternative to perform accurate outdoor location of WiFi access points?*” A potential solution using common-off-the-shelf (COTS) hardware would break access point location out of a small niche market of enterprise administrators, and make it available to home users and small businesses managing their own local hotspots.

Our solution is derived from a key insight: “*by rotating a standard wireless receiver around a blocking object, we can effectively emulate the sensitivity and functionality of a directional antenna.*” We exploit the property that the signal strength observed by a wireless receiver drops most significantly when there is a large obstacle directly between it and the transmitter. Such a drop in signal strength is strongest when the receiver is directly adjacent to the obstacle, and

should be observable as long as the obstacle is large enough to block a significant portion of the reception angle. Therefore, by “rotating” the receiver’s position with respect to the obstacle (and the signal “void” it produces), and observing the received signal strength, we can determine the approximate direction of the transmitter. We refer to this process as *directional analysis*.

While this is a general observation potentially applicable to a variety of wireless transmissions on different frequencies and hardware, we focus our attention in this paper on a single use case: outdoor location of WiFi access points using smartphones. Applying our insight, we hypothesize that a user can accurately locate WiFi APs using common-off-the-shelf smartphones as receivers, and her own body as the blocking obstacle. To perform a directional analysis operation, she slowly rotates her body around 360 degrees, while keeping the smartphone in front of her and performing periodic received signal strength (RSS) measurements. The observed RSS should be at its lowest point when the user’s body is directly between the smartphone and the wireless AP. By walking towards the predicted direction of the AP, and periodically repeating the directional analysis, a user should be able to zone in and locate a specific AP.

To validate our hypothesis, we first perform detailed experiments to see if such dips in signal strength can be observed using today’s smartphones and standard WiFi access points. We run tests on several different environments on the UC Santa Barbara campus, target both 802.11b/g and 802.11n (MIMO) APs, and use three different smartphone platforms, Android, WindowsMobile, and Apple iPhone4. Using a large number of real measurements, we do in fact observe the expected dip in signal strength when the tester’s body blocks the smartphone from the access point. We consistently observe the artifact across both platforms, environments with different levels of obstructions, and for users of different heights and weights.

Borealis. These measurement results lead us to develop *Borealis*, a system for locating outdoor WiFi access points using software on commodity smartphones. Borealis users can perform robust directional analysis by turning their bodies on a 360° axis, and use this technique to locate a transmitting access point. We address several challenges in the process of building Borealis, most significant of which is that errors in directional analysis are impacted by environmental conditions, and particularly exacerbated by multipath propagation around areas with multiple buildings. Our solution is to build a model that predicts the impact of blocking obstacles on signal strength at the receiver, use it to identify large blocking effects, and in doing so, produce more accurate predictions of the AP’s direction. Combined with prior techniques on direction-guided user navigation [10, 14, 15], we produce a system that efficiently guides users to an AP. In scenarios where the AP is housed indoors, our system guides the user to an outdoor location closest to the AP.

We implement and deploy Borealis as a user application and a set of kernel modifications on the Android platform. We modify the WiFi driver to focus scans on specific channels, thus allowing more frequent RSS measurements while reducing energy usage. Our user application assists the user in directional analysis by logging each RSS measurement along with data from the orientation sensor, and using it to compute the most likely direction of the access point along

with a prediction confidence. The user can repeat the process with success until she is within 2 meters of the AP.

We perform detailed experimental evaluation on Motorola Droid and HTC G1 phones. Our results show that Borealis produces significantly more accurate AP direction values than GUIDE [10], the recent system using triangle gradients to compute AP direction. We also find that running Borealis in real-time produces similar accuracy compared to an offline version using learning techniques on a large number of RSS measurements. Finally, our results show that using our direction predictions, Borealis leads the user on a path that reasonably approximates the shortest path to the AP.

Our work shows that with small software modifications, today’s smartphones can effectively replace directional antennas to locate outdoor WiFi access points. Borealis is the first example of a potential class of systems that approximate directional antenna systems using rotation around signal-blocking obstacles, and its underlying principle can potentially be applied to build location systems for other wireless transmissions.

2. PRELIMINARIES

In this section, we briefly describe the problem scenario and assumptions. We then summarize existing work that is most relevant to our target scenario.

Our focus is to accurately locate outdoor WiFi AP using common off-the-shelf smartphones as receivers. This functionality is a critical part of managing WiFi access networks for both small business and home users. The problem scenario is simple – a user, holding a smartphone, seeks to find the physical location of a WiFi AP from its BSSID. Note that our AP location problem is different from the general wireless localization problem [4, 6, 12]. The user seeks to locate a transmitting AP rather than determining her own location.

2.1 Related Work

In general, a receiver locates a transmitter by examining received signal strength (RSS), time of arrival (TOA) [7], time difference of arrival (TDOA) [13], or angle of arrival (AOA) [20]. The latter three methods all require simultaneous measurements at multiple receivers or antenna array. They are commonly used in cellular networks where neighboring basestations collaborate to locate a mobile device. They are, however, not feasible in our target scenario. Thus we focus our discussion on the RSS based method.

We categorize the RSS based solutions into three groups.

Model-based. Solutions in this category consider absolute values of RSS. Using RSS values measured at one [26] or multiple [8] locations, existing designs seek to derive the physical distance between the tester and the transmitter or the exact transmitter location, based on a radio propagation model. This solution, however, is fundamentally limited by the inaccuracy of propagation models for practical environments [12].

Gradient-based. This type of solution compares the RSS values obtained from different locations. Based on the assumption that a location closer to the AP will have a higher RSS value, existing designs estimate the AP direction by computing the gradient of the RSS value across different locations. This is done either online [10, 29] using a small set of local measurements, or offline [11] by integrating the results




Figure 1: Illustrations of users facing the AP, with back to the AP, and rotating while holding a smartphone.

of a large number of measurements. The problem with this line of solution is its idealized assumption that RSS values increase as the receiver moves closer to the transmitter. Our own measurements show that this assumption often breaks in practice, leading to large errors.

Directional Antennas. The use of directional antenna, either at transmitter [17, 27] or receiver [14, 15, 18, 28], can significantly boost localization performance. For example, by rotating the beam of its antenna, a receiver can pinpoint the direction of the AP as the direction that provides the highest receive signal strength [28]. The drawback is that it requires specialized hardware.

Among the above solutions, the directional antenna based solution achieves most of our system requirements: it is accurate, uses a single radio, and operates online with a small measurement overhead. However, directional antennas are prohibitively expensive for home users and not portable enough for personal use.

3. THE BLOCKING OBSTACLE EFFECT

The key insight that enables our approach to locating access points is surprisingly simple. We hypothesize that when placed next to a large obstacle, the signal strength perceived by a wireless receiver will be highly dependent on where the obstacle is relative to the position of the receiver. The closer the obstacle is to blocking line-of-sight between the AP and the receiver, the more of the signal is blocked, and the weaker the signal strength seen by the receiver.

We apply this hypothesis to our context of smartphone based AP location. The body of a user holding a smartphone will block a portion of the incoming WiFi signal. The closer the user is to being on the straight line between the smartphone and the AP, the weaker the signal perceived by the phone. This blocking effect of the human body has been observed on a variety of frequencies and radio hardware [9, 24, 30], even in indoor environments [4]. Consider Figure 1(a)-(b). When her back is towards the AP, the user’s body becomes an obstacle and blocks the signals from reaching the smartphone. When facing the AP, the body is no longer an obstacle. Therefore, we expect that as a user rotates herself in place, as in Figure 1(c), the phone’s received signal strength will display an interesting artifact: a peak when she faces the AP and a dip when her back faces the AP. Thus by measuring signal strength at different rotational angles, a user can gain a hint of which direction points towards the AP.

In this section, we verify our hypothesis using detailed smartphone experiments. Our experiments seek to study the impact the body as an obstacle has on signal strength

as the user rotates herself. We seek to understand the factors that cause this artifact, and the impact on this artifact by factors such as propagation environment, phone hardware and WiFi standards.

3.1 Smartphone Experiments

Our experiments use four popular smartphones: Motorola Droid, HTC Dream/G1, Apple iPhone 4, and LG Fathom (WindowsMobile 6.5). The first two phones support 802.11b/g and the last two support 802.11b/g/n. We program each smartphone to poll its RSS reports, and record the received WiFi signal strength from a given AP. Using its built-in compass, each phone records its relative orientation¹ as the user rotates in a clockwise direction. To diminish the impact of fast fading, we smooth the RSS values using a moving average window of 20° along the measured angle.

Our experiments will answer five key questions:

- a) *Is there a clear observable artifact?*
- b) *What is the source of this artifact?*
- c) *Is it a function of the propagation environment?*
- d) *Is it sensitive to different phone hardware or AP configurations, i.e. 802.11b/g vs. 802.11n?*
- e) *Is it sensitive to the height/weight of users, or how they position the phone with respect to their body?*

Experiment 1: Signal Patterns during Rotation. Our first experiments examine the signal strength artifacts during user rotation. At each test location, a user holds a smartphone in the hand and rotates her body by 360° at a speed of 6° per second. Each experiment records the instantaneous RSS values and the phone/user orientations. We set up the AP in an open space area where a direct line-of-sight path exists between the AP and the phone. In the rest of the paper, we refer to this environment as Simple LOS. This basic configuration allows us to study the impact of body position with minimum impact from other factors.

The results confirm our hypothesis. Figure 2(a) and (b) plot instances of the signal strength distribution as a function of the user orientation, for both Droid and G1 phones. For easy illustration, we mark the direction of user facing the AP as 60° and the direction with back to the AP as 240°. The results, marked as “w/ human shield”, show a clear dip pattern. When the user has her back facing (BF) the AP, the received signal strength is about 10-15dB lower than that when the user faces the AP. We observed similar patterns as we varied the distance between the phone and the AP. We observed up to 18dB difference in RSS values even when the user was only ≈2 meters away from the AP. We also varied the height of the AP between 2 and 18 me-

¹Relative to the initial direction the user was facing.

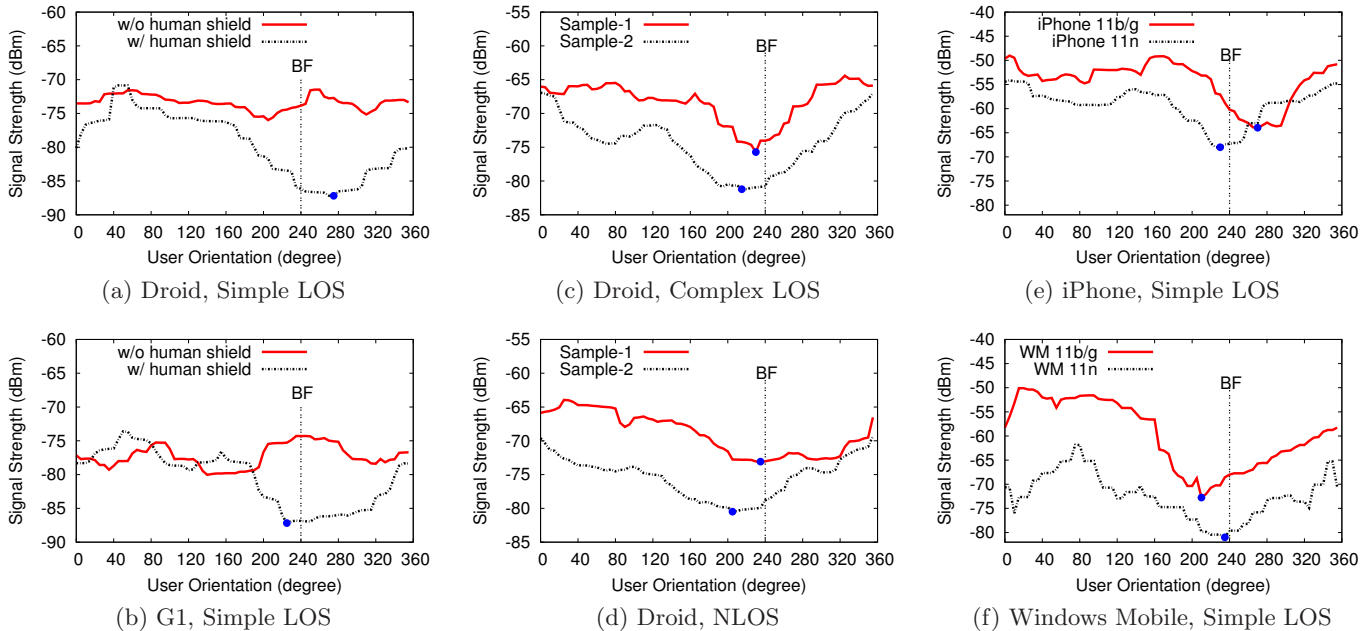


Figure 2: Observed signal strength profiles with user rotation. We mark the user direction facing the AP as 60° , the direction with user’s back to the AP as 240° (marked by BF). (a)-(b) When a user holds the phone with Droid and G1 phones in a simple LOS environment, the signal profile displays a clear low signal artifact. (c)-(d) The artifact is consistently visible in the complex LOS and NLOS environments. (e)-(f) The artifact is easily visible using iPhone 4 and Windows Mobile phones with both 802.11b/g and 802.11n APs.

ters, and observed no difference in results. Finally, in each graph, we mark the angle representing back-facing the AP with a vertical line (BF), and mark the point of lowest signal strength with a blue dot.

Experiment 2: Source of the Artifact. Our next question is to identify whether the signal strength artifacts are caused by the position of the user’s body, or the orientation of the phone antenna [25]. First, we repeat Experiment 1 multiple times, each time changing the orientation of the phone antenna, and observe no change in the results. Next, we repeat the same experiment, but remove the human body as an obstacle by placing the phone flat on top of a small chest-high table. We shift the orientation of the phone in 20° increments, and measure the signal strength for 30 seconds with the phone in each orientation. The results are plotted in Figure 2(a) and (b) as “w/o human shield,” and show that the low-signal artifact is clearly gone. Thus we conclude that the low signal artifact is mainly caused by position of the human body, and its impact in blocking the AP’s signal.

Experiment 3: Dependency on Environments. We also repeat the above experiments under different propagation environments. In addition to Simple LOS, we include an open area with surrounding buildings and moving objects (Complex LOS), and an environment with multiple buildings and objects and thus no line-of-sight paths between the phone and the AP (NLOS). We show illustrations of these three measurement areas in Figure 3. We perform multiple measurements at different positions in each environment, and plot two sample results for Complex LOS and NLOS in Figure 2(c) and (d). We see that the low signal

artifact still exists, while the magnitude of the dip may vary between different positions. This is expected, since reflections from surrounding objects can create multipath signals that reduce the impact of the user body blocking signals.

Experiment 4: Impact of Phone and AP Configuration. These experiments measure the low signal artifact under different phone hardware and AP configurations. Figure 2(e) and (f) show the standard experiment results using both iPhone4 and WindowMobile phones in a Simple LOS environment. We see that both phones exhibit the same artifacts for both 802.11b/g and 802.11n APs. The AP hardware results are notable, since it means that even connecting to 802.11n APs who use multiple antennas, the low signal artifact still comes across clearly in measurement results.

Experiment 5: Impact of User Posture and Body Shape. We repeat the above experiments with 6 additional users, with varying heights and weights, ranging from 5’4”, 100lbs to 6’, 160lbs. Each user holds the phone using a slightly different posture. The same artifact consistently appears in all results, thus demonstrating that the low signal effect is prevalent across users.

3.2 Key Observations

Overall, our experiments lead to two key findings. First, our experiments confirm that the position of the user’s body can significantly affect the smartphone’s received WiFi signal strength. When the user has her back facing the AP, her body becomes an obstacle and significantly degrades the received signal strength.

Second, for each of our result graphs, we use BF (for back facing AP) to mark the opposite AP direction where the




Figure 3: Graphical illustration of the three propagation environments used in our experiments, all located on the UC Santa Barbara campus. For both Complex LOS and NLOS, the AP was mounted right below the roof of an office trailer, and there are many static and moving obstacles nearby, including trees, cars, bikers and pedestrians. For NLOS, the path between the AP and the measured region is blocked by office trailers.

blocking effect should be at its strongest. We also mark the point of lowest signal strength with a large blue dot. It is clear that across all of our experiments with different phones and different environments, the angle with the lowest signal strength does not capture the predicted direction away from the AP. The “error” between the angle showing the lowest signal and the ideal angle can be as low as near zero (WindowsMobile, 11n), or as high as $\approx 40^\circ$ (Sample 2, Droid NLOS). Clearly, an accurate AP location system cannot simply rely on finding the angle with lowest signal strength, and must use more sophisticated techniques to determine the AP direction.

4. ACCURATE ACCESS POINT LOCATION

Motivated by our findings, we propose *Borealis*, a new AP localization system for smartphones that leverages signal strength artifacts to compute the direction of an AP. Unlike conventional solutions that either require sophisticated radio hardware (*i.e.* directional antennas) [28], or extensive war-driving measurements [11], our solution uses off-the-shelf smartphones and produces real-time results with a small number of measurements.

4.1 Borealis Overview and Challenges

The concept behind our design is simple. At a fixed location, a user rotates herself and records the signal strength profile of the target AP. Based on the observed signal artifact, she estimates the direction of the AP from her current location. She moves towards the AP, occasionally repeating the measurement step, until she arrives at the AP.

Borealis has two key requirements. First, given the limited resources on smartphones, Borealis must use minimal energy and computational resource in its directional analysis. Second, Borealis must produce results in real-time, and its directional measurements should be minimized to conserve user effort. We can meet both of these goals by

designing a system that minimizes computation while producing accurate results. The more accurate the result, the fewer number of direction analysis measurements are necessary, thereby saving both device battery power and user effort.

Our biggest challenge is how to reliably determine the AP direction based on the measured signal strength profile. Our initial experiments in Section 3 show that simply using the angle with lowest signal strength to estimate AP direction can generate large estimation errors (as indicated in Figure 2 by the distance on the x-axis between BF and the dot representing minimum signal strength). Across roughly 40% of all our experiments, using minimum RSS to estimate the AP direction leads to an error between 40° and 120° .

This error comes from two factors. First, as the user rotates, we measure the signal strength of various directions sequentially, not simultaneously. Thus time dynamics in signal propagation produces variance in signal strength values across different directions. Second, the measurement time duration at each angle is limited, which leads to measurement noise. Extending the per-angle measurement time can help reduce noise, but also increases user effort, power consumption, and gap between sequential measurements.

4.2 Modeling the Body as a Signal Obstacle

Borealis addresses the significant error introduced by time dynamics and propagation effects using a model-driven approach. We first study the smartphone measurement traces to understand the relationships between signal strength profile, user orientation, and AP direction. We build a model to capture these relationships and predict the impact of blocking obstacles on signal strength. We use this model to develop an accurate AP direction prediction mechanism, along with a way to estimate the confidence of each estimate.

Our first observation from our experiments is that as the user rotates, not only does her body block the wireless signal and reduce its strength, but also that this signal degradation occurs *at a range of angles*, not just when the user is facing away from the AP.

To better understand this phenomenon, we introduce a simplified object-blocking model. Consider a simple propagation environment without any obstacles blocking the signal; a line-of-sight (LOS) path is the dominant signal path between the AP and the phone. We model the human body and the phone in a diagram in Figure 4(a), where the human body has width b and the distance between the phone and the body is p . Next, Figure 4(b) illustrates the signal propagation condition as the user rotates herself. We see that when the orientation of the user and phone is within an angular sector between \vec{d}_1 and \vec{d}_2 , no LOS path will reach the phone. We refer to this angular sector as the *blocking sector*. Because the dominating path is blocked, the signal strength observed in this sector will be significantly lower than that of the other orientations.

Modeling the Blocking Sector. Using simple geometry, we can further characterize the blocking sector, particularly its size in degrees.

THEOREM 1. *When line-of-sight is the dominant propagation path between the AP and the smartphone, the angular size β of the blocking sector is defined by*

$$\beta = 180^\circ - 2(\arctan \frac{2p}{b} - \arcsin \frac{bp}{d\sqrt{4p^2 + b^2}}). \quad (1)$$




Figure 4: An abstract model of the human shield effect. (a) A simple model of a user holding a smartphone, with body width b and phone/body distance p . (b) The signal propagation condition as the user rotates. When the user orientation is between \vec{d}_1 and \vec{d}_2 , the LOS path will be blocked by the human body. This range of orientation is referred to as the blocking sector. (c) A geometric representation of the blocking sector.

where p is the distance between the human body and the phone, b is the body width, and d is the distance between the AP and the user.

PROOF. As shown in Figure 4(c), line l passes the AP (denoted by A) and the center of human body (denoted by B). Assume the rightmost point of human body is C and consider the triangle $\triangle ABC$. By the sine law, we have $\frac{b}{2 \sin \theta} = \frac{d}{\sin \gamma}$, where θ and γ are marked in Figure 4(c). Similarly, we have $\gamma = \theta + \alpha$. Given that $\gamma = \arctan \frac{2p}{b}$, we arrive at $\alpha = 2(\arctan \frac{2p}{b} - \arcsin \frac{bp}{d\sqrt{4p^2+b^2}})$. Since $\beta = 180^\circ - \alpha$, we prove that Equation (1) holds. \square

In practice, the distance between the AP and the user is much longer than the human body width ($\approx 0.4\text{m}$), namely $d \gg b, d \gg p$. Furthermore, when a user holds the smartphone in a natural posture, the smartphone is roughly half of the body width away, namely $b \approx 2p$. Then we can reduce (1) to

$$\beta \approx 180^\circ - 2 \arctan \frac{2p}{b} = 90^\circ. \quad (2)$$

Summary of Findings. The objective of our analysis is not to model the exact impact of the human shield effect, but to capture the large-scale trend of the signal strength profile, and its relationship with the user orientation and the AP direction. Along these lines, our analysis leads to two key insights:

- During user rotation, the signal strength degrades heavily when the user orientation is within a range, defined by the blocking sector. The angular size of the sector is roughly 90° under general configurations.
- We can derive an estimate of the direction to the AP, by taking the opposite direction of the center angle of the blocking sector β , as in Figure 4(c).

4.3 Directional Analysis via Blocking Sector

Motivated by the insights from our model, we propose to estimate AP direction by locating the blocking sector within the signal strength profile. In essence, we organize the observed signal strength profile into overlapping sectors of size β , and locate the candidate sector that displays the largest relative signal degradation. The opposite direction from the center of the blocking sector is then the AP direction.

We now describe the detailed procedure of our proposed directional analysis. The input to our analysis is the *raw*

RSS value (in dBm) corresponding to each measured user orientation. Here we do not apply any data smoothing, unlike the results shown in Section 3. We represent the measured RSS profile during a user rotation by

$$R = \{(RSS(\theta_i), \theta_i) | i \in [1, N]\} \quad (3)$$

where N is the number of measurement points, $RSS(\theta_i)$ represents the raw RSS reading (dBm) at point i , and θ_i is the user orientation at point i (the clockwise angle from the North orientation).

Organizing Candidate Sectors. We first form candidate sectors of width β . This is done by applying a sliding window of width β on the cyclic version of the signal strength profile, and shifting the window by Δ after forming a sector. The j^{th} sector contains measurement points whose θ satisfies: $(j-1)\Delta \leq \theta < (j-1)\Delta + \beta$. As a result, we create a total of $K = 360/\Delta$ overlapping candidate sectors. The value of Δ directly affects the accuracy and computation overhead. From our experiments, we did not observe noticeable gain of using $\Delta < 20^\circ$. Thus we chose $\Delta = 20^\circ$ in this work.

Locating the Blocking Sector. For each candidate sector S_j , we compute the *relative signal degradation* by subtracting the average signal strength of the sector from the average signal strength outside the sector:

$$\text{Diff}(j) = \frac{\sum_{\theta \notin S_j} RSS(\theta)}{N - |S_j|} - \frac{\sum_{\theta \in S_j} RSS(\theta)}{|S_j|} \quad (4)$$

where $|S_j|$ is the number of measurement points within the sector S_j and N is the total measurement points in the signal profile. By considering the signal strength distribution within and outside the sector, this relative degradation metric helps to mitigate the negative impact of non-uniform distribution of measurement points within the signal profile. While (4) may not be an optimal function for identifying the blocking sector, it works sufficiently well for our purposes.

Using the Diff function, the blocking sector S^* is then

$$S^* = \underset{S_j}{\text{argmax}} \text{Diff}(j), \quad (5)$$

the sector that suffers the heaviest signal strength degradation. After locating S^* , we mark the opposite direction of the center orientation of S^* as the AP direction.

Figure 5(a) demonstrates our proposed direction analysis in terms of the RSS signal profile, the function Diff, the center angle of the detected blocking sector (marked by the red




Figure 5: Borealis' directional analysis. (a) Deriving the AP direction based on the blocking sector. (b) Choosing the blocking sector size. (c) Predicting the confidence of our direction analysis by computing the cross-correlation between the measured and ideal signal strength profiles.

triangle), and the actual direction opposite to the AP direction (marked by BF). In this example figure, we smooth the RSS profile using a 20° sliding window average to show the general trend. We see that Borealis obtains an estimate within 10° from the actual direction. Yet if we use the MinR method, the estimated direction will be 30° away from the ideal result. The gain of our solution comes from the fact that we examine the signal strength distribution at the sector level, which not only captures the key effect of body blocking, but is also robust against low levels of statistical variance in the measured data.

Impact of β . The above discussion also leads to another question: *How important is the chosen value of β ?* While this is difficult to study analytically, we performed experiments to verify the impact of β on the direction estimation. Figure 5(b) plots the percentage of measurement locations whose estimation angular error is less than 30° , as a function of β , for all three environmental settings in Figure 3. We see that the value of β that produces the least amount of error is between $[80^\circ, 100^\circ]$ for all three settings. For simplicity, we use $\beta = 90^\circ$ in subsequent experiments.

4.4 Confidence of Directional Analysis

Borealis' directional analysis provides another result, by providing a confidence level associated with each estimate of AP direction. Since no direction estimate is perfectly accurate, an associated confidence value gives the user additional useful information. As we show in Section 4.5, Borealis uses this confidence value of each estimate to control how often a user needs to repeat the direction estimate. Therefore, Borealis can bound the impact of directional analysis errors during user navigation.

We compute the confidence by comparing the signal strength profile to an idealized profile derived from our abstract model on body blocking. That is, ideally, the measured signal strength profile should display a dip pattern, where the signal strength is significantly lower for a range of angles. Based on this insight, we build an ideal profile using a square-waveform-like curve, shown in Figure 5(c). The width of the dip is β ($\beta = 90^\circ$ in this paper), and the center of the dip is set to the opposite of the (estimated) AP direction. The amplitude of the square is A_p for the peak and A_d for the dip, $A_p > A_d$.

We use the confidence to capture the similarity between the measured signal strength profile and the ideal profile.

It can be computed as the cross-correlation coefficient of the two profiles [21], a widely used similarity metric for any two waveforms [5]. The correlation coefficient seeks to capture the similarity between the measured and ideal dip patterns, assuming both patterns produce the same AP direction estimation. Hence before computing the correlation, we first align the center of the dip in the ideal profile to the opposite direction of the estimated AP direction, as shown in Figure 5(c). Next, let $\mathbf{T} = [t_0, \dots, t_{m-1}]$ and $\mathbf{R} = [r_0, \dots, r_{m-1}]$ denote the vectors of RSS values for the idealized (and aligned) profile and the measured profile, respectively. The confidence value ρ is:

$$\rho = \frac{1}{m} \cdot \frac{\sum_{i \in [0, m-1]} (t_i - \bar{t})(r_i - \bar{r})}{\sigma_{\mathbf{T}} \cdot \sigma_{\mathbf{R}}}, \quad (6)$$

where \bar{t}, \bar{r} are the mean values of \mathbf{T}, \mathbf{R} respectively, and $\sigma_{\mathbf{T}}, \sigma_{\mathbf{R}}$ are the standard deviations of \mathbf{T}, \mathbf{R} . The larger ρ is, the more confident Borealis is about its estimate. It is easy to prove that the values of A_p and A_d do not affect ρ , as long as $A_p > A_d$. Thus we set $A_p=1$ and $A_d=0$.

4.5 Direction-Guided User Navigation

While direction analysis provides a good estimate of the AP's direction relative to a location, the goal of Borealis is to allow users to determine the AP's physical location. We choose to adopt methodology similar to prior work [10, 14, 15], where a user moves towards the AP and performs periodic direction estimates to tune its direction. Navigation ends when she reaches the AP.

We modify prior approaches by leveraging our prediction confidence results to determine how frequently a user should update her direction of movement. Doing so less frequently, *i.e.* longer distances between measurements, reduces the number of measurements required. But this means error from a single measurement will have a greater impact on the efficiency of the path taken. We dynamically select the step size based on the confidence value predicted for each direction estimate. A high confidence implies a reliable estimate and means the user can travel a longer distance before repeating the estimate. In contrast, a direction estimate with low confidence means the user is likely in a location with complex propagation conditions, and should repeat the estimate after moving a short distance along the predicted direction. We use detailed experiments to evaluate this in Section 6.




Figure 6: Borealis architecture overview.

5. A BOREALIS PROTOTYPE

We implemented a prototype of Borealis on both Motorola Droid and HTC Dream/G1 phones, running the Android 2.2 Froyo OS. In the application layer, we implemented Borealis’ directional analysis (with confidence prediction) and user navigation using Java and Android SDK. In the OS layer, we modified the WiFi driver to perform faster and more efficient RSS measurements. The architecture of our prototype is shown in Figure 6.

Application Layer. Our program allows the user to identify a target AP based on its SSID, and directs the user to start the rotation. During each rotation, our program continues to collect user orientation and RSS values by polling the compass sensor and WiFi RSS reports. Because the smartphones report orientation data at a finer granularity than RSS, we pair each RSS reading with an orientation reading based on their time stamps. Our program then analyzes the measured signal profile, and computes the AP direction and the confidence value of the current estimation. The direction is then displayed on the phone to guide the user navigation. During navigation, our program also determines online whether the next measurement point is reached, and informs the user accordingly.

OS Layer. While our main development effort lies in the application layer, in the OS layer we modified the Android WiFi driver to boost the speed and efficiency of WiFi signal measurements. Android offers a native “Scan” function to collect RSS reports of a target AP. However, this operation scans all 13 WiFi channels sequentially, wasting both time and energy. To fix this, we modified the WiFi driver to only scan and report RSS on requested channels. This change reduces the RSS report time by a factor of 10.

For our current prototype, each user rotation takes 1 minute to obtain a smooth RSS profile. We are currently working on understanding the minimum duration that produces an accurate result. Our preliminary results show that rotation duration of 10 seconds produces similar results to those of 1 minute durations. In addition, we are also investigating techniques that would allow us to measure RSS signal profiles opportunistically using users’ natural movements, rather than forcing them to perform in-place rotation.

6. EVALUATION

In this section, we evaluate our Borealis prototype using experiments on five Motorola Droid and HTC G1 phones. We use a Linksys WRT54GL 802.11b/g router as the WiFi access point, with a transmit power of 200mW. Our experiments were conducted over multiple days by seven users with different body shapes and ways of holding phones. Be-

cause G1 and Droid display similar results, we only show the Droid results for brevity.

We evaluate the impact of radio propagation by experimenting in three representative environments, as illustrated by Figure 3. In *Simple LOS*, we placed the AP on top of a shelf 2 meters in height, on a 100m×200m lawn. All the experiments were on the lawn and away from large buildings. In *Complex LOS*, we placed the AP on top of a trailer building (5-meter in height) and experimented in the nearby parking lot surrounded by trailer buildings of similar height. In these locations, we could still see the AP. For *NLOS*, we used the same configuration of Complex LOS but experimented along the hidden walking paths where the AP was no longer in sight. For each environment, we experimented with at least 400 locations for each phone. There were random human movements throughout our experiments, e.g. people walking or biking.

To evaluate Borealis, we compare four systems for deriving AP direction via signal measurements.

- **MinR** – The baseline algorithm for our proposed directional analysis. It treats the opposite direction of the weakest signal as the AP direction.
- **Borealis** – Our proposed directional analysis.
- **Offline Analysis** – An offline version of our directional analysis using a clustering-based learning algorithm. It first collects the signal strength profile from roughly 360 locations in each environment, along with the actual direction of the AP in each case, and uses clustering techniques to build an optimized model, which is then applied to the remaining 60 locations for each environment to generate accuracy results.
- **GUIDE** [10] – A prior work that measures signal strength at three locations (forming a triangle) and computes the signal strength gradient to determine the AP direction. This method is the most comparable to Borealis since it is online and requires a very small set of measurements.

In the following, we evaluate Borealis’ accuracy in AP direction estimates using both per-location measurements and user navigation experiments. We also examine the energy consumption of Borealis on both Droid and G1 phones.

6.1 Accuracy of Borealis Direction Estimation

We start from examining the accuracy of Borealis’ directional analysis. Figure 7(a) shows the cumulative distribution of the angular error in AP direction estimation. The angular error is the absolute difference in angular degree, between the estimated AP direction and its true direction. The results show that Borealis is fairly accurate in the Simple LOS environment – in 80% of locations, Borealis produces no more than 30° angular error. For the Complex LOS and NLOS environments where multipath propagation becomes dominating, Borealis can still maintain an error of no more than 50° and 65° for 80% of locations, respectively.

Sources of Large Errors. The above results show that occasionally, Borealis does produce large errors in direction estimation, particular for the Complex LOS and NLOS environments. To understand the cause of such errors, we studied the signal strength profile of locations with error higher than 60°. We observed that in most cases, the signal profile displays multiple dips, creating multiple peaks in the Diff function used in the directional analysis. In this




Figure 7: The performance of Borealis directional analysis. (a) The CDF of the angular error for the three propagation environments. Borealis is fairly accurate for most of the test locations. (b) When Borealis produces larger errors, the signal profile often displays multiple dips, which creates multiple peaks in the Diff function. In this case, Borealis chooses the highest peak to estimate AP direction. (c) We observe a general trend where the confidence value of a direction estimation scales inversely with the angular error.

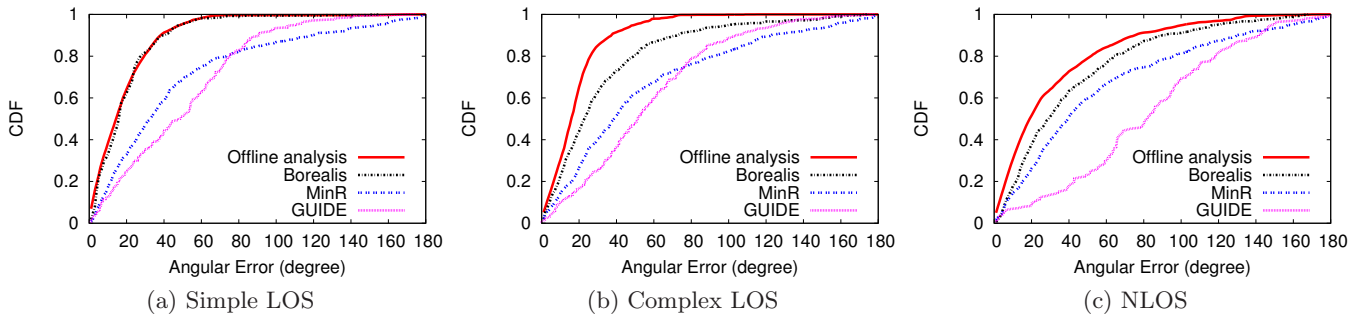


Figure 8: Comparing Borealis to Offline Analysis, MinR, and GUIDE in the three environments. Borealis significantly outperforms MinR and GUIDE, and is within a small distance from its offline version (Offline Analysis).

case, Borealis’ direction analysis uses the highest peak to estimate AP direction (based on eq. (5)). Figure 7(b) plots the probability of having multiple peaks in Diff, as a function of the angular error. Clearly, the larger the angular error, the more likely that it is caused by multiple dips in the measured signal profile. We also examined the spatial measurement locations with large estimation errors ($> 80^\circ$), and did not find any strong correlation between the two.

Confidence Prediction. Different from conventional solutions, Borealis also measures the confidence of its direction estimation. Figure 7(c) plots the relationship between the angular error of the direction estimation and the confidence value, using measurement data from all three environments. We observe a general trend that the angular error scales inversely with the confidence value.

An interesting question is whether the pattern observed in Figure 7(b), *i.e.* multiple dips in the signal profile or multiple peaks in Diff, can be used to refine the confidence estimation. The answer is no. This is because having large errors always maps to having multiple peaks in Diff, but not vice versa. From Figure 7(b) we see that the pattern appears even for cases with small angular errors.

6.2 Comparison to MinR & Offline Analysis

To examine the optimality of Borealis within the proposed directional analysis, we now compare Borealis’ sector-based

estimation to MinR, a simple estimation method, and Offline Analysis, the offline version of Borealis that uses training data to optimize direction decision. The comparison to MinR allows us to understand the gain of sector-based analysis, while the comparison to Offline Analysis allows us to understand the distance between Borealis and the “upper bound” performance of our directional analysis. The comparison is shown in Figure 8 in terms of the statistical distribution of the angular error.

We make two key observations. First, there is a large performance gap between MinR and Borealis. For example, for 80% of locations, the bound on angular error of MinR is 60° for Simple LOS, 120° for Complex LOS, and 135° for NLOS. This is roughly 2 times the Borealis’ estimation error. As we have discussed earlier, such large error is because inherent signal variations create random ripples in signal strength profile. Thus a single point based direction analysis is highly sensitive to such random variations, leading to large errors. These errors are further exacerbated by multipath propagation in complex environments.

Second, the gap between Borealis and its offline-trained version is rather small, and even negligible in Simple LOS. We note that the offline version has the advantage of optimizing the direction estimation mechanism based on training data, so that it is able to recognize certain patterns in complex environments and produce a more accurate decision. On the other hand, the cost of such small improvement

is the large measurement overhead and the fact of being an offline learning solution requires actual knowledge of the AP direction. Overall, because the gap between the two methods is small, we conclude that Borealis is a practical and effective solution for determining AP direction in real time.

6.3 Comparison to GUIDE

We now turn our attention to GUIDE [10], the most relevant work in AP direction prediction using signal measurements. Similar to Borealis, GUIDE operates online, uses only a single receiver and requires a small set of measurements (at three locations). Different from Borealis, GUIDE applies a triangle-gradient based solution to estimate the AP direction.

We implemented GUIDE on our smartphones and experimented it in the three environments. Figure 8 compares the performance of GUIDE to that of Borealis. We see that for all three environments, Borealis significantly outperforms GUIDE. The key reason behind such large performance gap is that GUIDE (and other gradient based solutions) assumes that received signal strength degrades with the distance between the transmitter and receiver. In practice, this does not always hold, even in the simple LOS environment. Our own experiments show that the above assumption breaks in roughly 30% of the measurement cases.

For fairness, we did not compare Borealis to other AP localization methods, such as [28] and [11]. This is because these designs either require sophisticated directional antenna [28] or large measurements [11].

6.4 Locating Indoor APs

We also examined the scenario where the AP is placed indoors and an outdoor user collects signal measurements to estimate AP direction. Specifically, we consider the complex LOS setting in Figure 3 but place the AP inside the office trailer so that the wall of the trailer blocks all paths from the AP to the measured locations. The AP is located in the center of the trailer, away from windows and doors. We repeat the experiments at the same measurement locations used in our previous experiments.

In Figure 9, we compare the performance of both Borealis and GUIDE to prior results when the AP is placed outdoors. We see that multipath propagation in the indoor AP scenario does degrade the accuracy of direction estimation. However, the degradation is nearly negligible for Borealis. More importantly, Borealis still significantly outperforms GUIDE when locating indoor APs. These results show that Borealis is effective, and more accurate for locating indoor APs using outdoor signal measurements than GUIDE.

6.5 Borealis Navigation Efficiency

We also evaluate Borealis' end-to-end performance in terms of user navigation. We randomly selected 40 starting points in the three environments. They are within 60-140 meters from the AP and have RSS values around -90dBm. Further locations are not considered because no AP is detected. This range is similar to those used by prior work on outdoor AP location [10, 11].

We compare two Borealis navigation designs: i) navigation with periodic directional analysis (*e.g.* the tester rotates every 20 meters), and ii) confidence-guided adaptive navigation which determines the next location of directional anal-

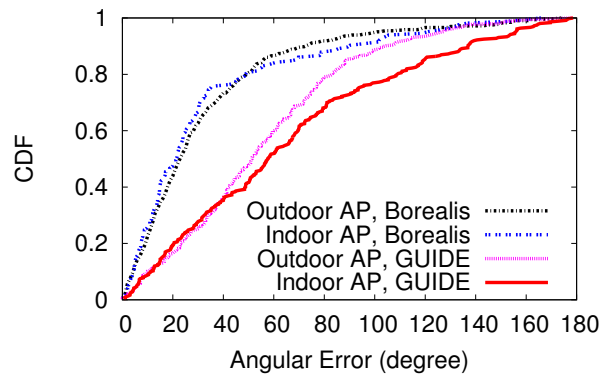


Figure 9: The performance of directional analysis when outdoor users locate an indoor AP using Borealis or GUIDE. The accuracy is comparable to the cases where the same AP is placed outdoors.

ysis using the confidence of the current estimation. Specifically, if the confidence is above 0.8, the user rotates again after walking 30 meters. If the confidence is less than 0.3, she rotates after 10 meters.

We use two performance metrics: *navigation overhead* and *measurement frequency*. The navigation overhead is the normalized extra distance traveled to locate the AP: $\frac{\text{navigation distance}}{\text{shortest distance}} - 1$. We compute the navigation distance using GPS records of the navigation, and the shortest distance from GPS coordinates of the AP and the starting point. Because the shortest path might be blocked by obstacles, this metric serves as the upper bound on the extra distance traveled to locate an AP. On the other hand, the measurement frequency defines average distance between two consecutive user rotations. It should be close to 20 meters for the periodic method, but larger for the adaptive design.

Table 1 lists both metrics averaged from the 40 experiments, using the periodic and adaptive methods. We see that in both Simple and Complex LOS environments, the adaptive design not only reduces the navigation overhead but also the measurement frequency. In particular, even in Complex LOS, it reduces the navigation overhead by half, and only invokes user rotation every 30+ meters. The gain is smaller in NLOS due to lower confidence in our directional analysis.

Figure 10 plots three sample navigation paths for the NLOS environment. We see that the navigation paths generally follow feasible walk paths, which demonstrates the efficiency of our proposed solution. Buildings and trailers do affect the accuracy of Borealis' direction estimates. But such errors are easily avoided by the user moving to a different location. An interesting observation is that multipath propagation is not always harmful. Often the strongest signal component comes from a path circumventing the obstacle, and the measured signal profile will indicate a strong dip near the direction of the open path. While it might not point to the exact AP direction, it certainly helps the user to identify a feasible path to the AP.

Borealis vs. GUIDE. We repeat the above experiments using GUIDE and its navigation procedure presented in [10]. For practical reasons, we stop GUIDE's navigation process when its navigation overhead is greater than 200% for Simple

		Navigation Overhead	Measurement Frequency
Simple LOS	Periodic	48%	22.59m
	Adaptive	15%	31.26m
Complex LOS	Periodic	74%	21.63m
	Adaptive	37%	30.62m
NLOS	Periodic	134%	20.45m
	Adaptive	107%	21.47m

Table 1: Performance of Borealis’ navigation in the three propagation environments. Adaptive navigation guided by the confidence value not only reduces measurement frequency, but also shortens navigation distance.

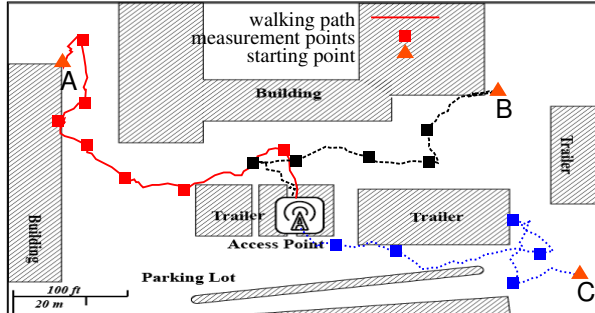


Figure 10: Sample navigation paths of Borealis in the NLOS environment. Points A, B, C mark the three starting points, and squares mark the locations of user rotation.

LOS and 500% for Complex LOS and NLOS. Note that under this constraint, Borealis, using both periodic and adaptive navigation, can always reach the AP within a distance of 2m. But GUIDE’s navigation rarely finds the AP. Only for 15% of all cases is GUIDE able to approach the AP within 11m. For all other cases, the user is still 35-184m away from the AP and does not have an accurate angular direction to the AP location. This result is consistent with the results offered by [10]. It is not surprising, since Borealis significantly outperforms GUIDE in direction analysis (Figure 8).

6.6 Energy Consumption

We now evaluate Borealis’ battery consumption using both Droid and G1 phones. Our analysis leverages the battery usage summary tool offered by Android. In addition to identifying the total battery usage of each Borealis’ directional analysis operation, we also study the distribution of energy costs across the different components involved.

We configure our energy experiments as follows. Because the Android battery report has a coarse granularity (per 1% battery usage for G1 and 5% for Droid), we use a brand new, fully charged battery for each phone and run Borealis repeatedly to drain the battery. For each experiment, we also call the Android PowerManager API to log the phone battery level every 10 minutes. We verified offline that this API has negligible impact on battery usage. Using the battery log and the Borealis trace, we compute the energy consumption of each Borealis operation. We use two Droid and two G1

	Droid	G1
% of battery consumed per Borealis operation	0.36%	0.78%
% of battery consumed ignoring Display and Standby	0.15%	0.29%

Distribution of Energy Usage across Components		
Display	54%	32%
Cell Standby	3%	31%
WiFi Radio	5%	12%
OS	21%	11%
Other Borealis Activity	17%	13%

Table 2: Energy consumption analysis of Borealis’ directional analysis on Droid and G1 phones.

phones in our experiments, and show the average results for each phone category.

Table 2 summarizes the results of our energy experiments, including a detailed breakdown to five major components (as reported by the Android battery usage summary). The total energy consumption of a single Borealis directional analysis takes 0.36% of the total battery for Droid phones and 0.78% for G1 phones. The majority of energy cost comes from Display (for both phones) and Cell Standby (for G1). If we remove these factors, the battery use of a Borealis direction analysis operation reduces to 0.15% for Droid and 0.29% for G1. The OS component of Borealis consumes a substantial bit of energy. This is not due to computation, but the fact that each Android app must run inside its own lightweight virtual machine.

Finally, we note that WiFi uses much less energy compared to other components, because Borealis’ signal measurements are passive and do not involve packet transmission. More importantly, Borealis only requires signal measurements during user rotation, and each navigation only requires a user rotation every 20-30m. Therefore, we conclude that normal usage of Borealis will not significantly impact the battery life of a smartphone.

7. CONCLUSION

In this paper, we described Borealis, a smartphone-based system for locating WiFi access points in real time. While our tests show Borealis to be effective on Android phones in different environments, earlier measurements suggest that the same techniques would be effective on other smartphone platforms as well.

More importantly, the underlying principle behind Borealis, using signal dips from blocking obstacles to locate wireless transmitters, is general and could be applied to locate other types of transmitters. For example, Figure 11 plots two sample signal strength profiles obtained from a rotating user holding a USRP2 GNU radio operating on 2.4GHz and 5GHz. It is clear that the same signal-blocking artifact exists for these frequencies as well. For transmitters on lower frequencies that penetrate deeper and exhibit more multipath propagation behavior, users can potentially “rotate” or move around larger obstacles such as trees, vehicles, or buildings. Development of these systems could address significant network management issues in the future, such

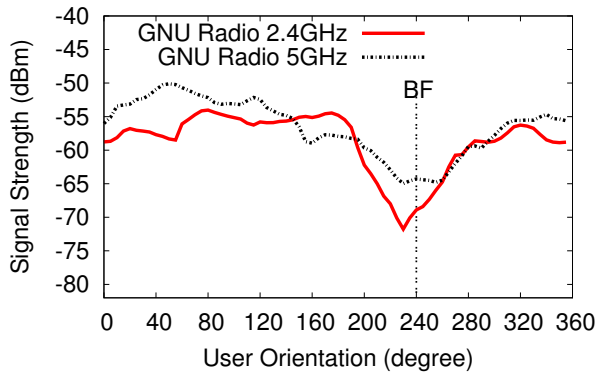


Figure 11: Sample signal strength profiles as a user rotates, measured using USRP2 GNU radios operating on 2.4GHz and 5GHz. We observe the same signal-blocking artifact. As before, BF marks the direction where the user has her back facing the AP.

as the enforcement of authorized transmissions in dynamic spectrum networks.

8. ACKNOWLEDGMENTS

We thank our shepherd Samir Das and the anonymous reviewers for their helpful suggestions. This work is supported in part by NSF Grants CNS-0916307, IIS-0847925, CNS-0832090, and CNS-0905667. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

9. REFERENCES

- [1] Google WiFi. <http://wifi.google.com/>.
- [2] WiFi data offload: Turning challenges into opportunities. <http://wifidataoffload.wordpress.com/2010/12/04>.
- [3] AT&T now has 460 mobile data offload points in NYC. <http://www.4gtrends.com/articles/37936/>, June, 2011.
- [4] BAHL, P., AND PADMANABHAN, V. N. RADAR: An in-building RF-based user location and tracking system. In *Proc. of INFOCOM* (2000).
- [5] BRACEWELL, R. *Pentagram Notation for Cross Correlation. The Fourier Transform and Its Applications*. New York: McGraw Hill, 1965.
- [6] BULUSU, N., HEIDEMANN, J., AND ESTRIN, D. GPS-less low-cost outdoor localization for very small devices. *IEEE Personal Communications* 7, 5 (2000), 28–34.
- [7] CHAN, Y., TSUI, W., SO, H., AND CHING, P. Time-of-arrival based localization under nlos conditions. *IEEE Transactions on Vehicular Technology* 55, 1 (2006), 17–24.
- [8] FINK, J., AND KUMAR, V. Online methods for radio signal mapping with mobile robots. In *Proc. of ICRA* (2010).
- [9] GHADDAR, M., TALBI, L., AND DENIDNI, T. Human body modelling for prediction of effect of people on indoor propagation channel. *Electronics Letters* 40, 25 (2004), 1592–1594.
- [10] GONZALEZ, M., GOMEZ, J., LOPEZ-GUERRERO, M., RANGEL, V., AND DE OCA, M. GUIDE-gradient: A guiding algorithm for mobile nodes in WLAN and Ad-hoc networks. *Wireless Personal Communications* 57 (April 2011), 629–653.
- [11] HAN, D., ANDERSEN, D., KAMINSKY, M., PAPAGIANNAKI, K., AND SESHAN, S. Access point localization using local signal strength gradient. In *Proc. of PAM* (2009).
- [12] HE, T., HUANG, C., BLUM, B., STANKOVIC, J., AND ABDELZAHER, T. Range-free localization scheme for large scale sensor networks. In *Proc. of MobiCom* (2003).
- [13] HO, K. C., AND XU, W. An accurate algebraic solution for moving source location using TDOA and FDOA measurements. *IEEE Transactions on Signal Processing* 52, 9 (2004), 2453–2463.
- [14] KIM, M., AND CHONG, N. RFID-based mobile robot guidance to a stationary target. *Mechatronics* 17, 4-5 (2007), 217–229.
- [15] KIM, M., AND CHONG, N. Direction sensing RFID reader for mobile robot navigation. *IEEE Transactions on Automation Science and Engineering* 6, 1 (2009), 44–54.
- [16] LAMARCA, A., CHAWATHE, Y., CONSOLVO, S., HIGHTOWER, J., SMITH, I., SCOTT, J., SOHN, T., HOWARD, J., HUGHES, J., POTTER, F., ET AL. Place lab: Device positioning using radio beacons in the wild. *Pervasive Computing* (2005), 116–133.
- [17] NASIPURI, A., AND LI, K. A directionality based location discovery scheme for wireless sensor networks. In *Proc. of WSN* (2002), ACM.
- [18] NICULESCU, D., AND NATH, B. VOR based stations for indoor 802.11 positioning. In *Proc. of MobiCom* (2004).
- [19] PACCHIANO, R. Track down rogue wireless access points. WiFi Planet Tutorial, March 2006.
- [20] PAGES-ZAMORA, A., VIDAL, J., AND BROOKS, D. Closed-form solution for positioning based on angle of arrival measurements. In *Proc. of PIMRC* (2002).
- [21] PEARSON, K. Mathematical contributions to the theory of evolution. iii. regression, heredity, and panmixia. *Philosophical Transactions of the Royal Society of London. Series. A* 187 (1896), 253–318.
- [22] POULSEN, K. Wardriver pleads guilty in lowe’s wifi hacks. SecurityFocus, June 2004. <http://www.securityfocus.com/news/8835>.
- [23] ROBERTS, B., AND PAHLAVAN, K. Site-specific RSS signature modeling for WiFi localization. In *Proc. of GLOBECOM* (2009).
- [24] RYCKAERT, J., DE DONCKER, P., MEYS, R., DE LE HOYE, A., AND DONNAY, S. Channel model for wireless communication around human body. *Electronics Letters* 40, 9 (2004), 543–544.
- [25] SANI, A., ZHONG, L., AND SABHARWAL, A. Directional antenna diversity for mobile devices: Characterizations and solutions. In *Proc. of MobiCom* (2010).
- [26] SAVVIDES, A., HAN, C., AND STRIVASTAVA, M. Dynamic fine-grained localization in ad-hoc networks of sensors. In *Proc. of MobiCom* (2001).
- [27] SHAH, S., SRIRANGARAJAN, S., AND TEWFIK, A. Implementation of a directional beacon-based position location algorithm in a signal processing framework. *IEEE Transactions on Wireless Communications* 9, 3 (2010), 1044–1053.
- [28] SUBRAMANIAN, A., DESHPANDE, P., GAOJIAO, J., AND DAS, S. Drive-by localization of roadside WiFi networks. In *Proc. of INFOCOM* (2008).
- [29] SUN, Y., XIAO, J., LI, X., AND CABRERA-MORA, F. Adaptive source localization by a mobile robot using signal power gradient in sensor networks. In *Proc. of GLOBECOM* (2008).
- [30] WELCH, T., MUSSELMAN, R., EMESSIENE, B., GIFT, P., CHOUDHURY, D., CASSADINE, D., AND YANO, S. The effects of the human body on UWB signal propagation in an indoor environment. *IEEE Journal on Selected Areas in Communications* 20, 9 (2002), 1778–1782.