

A Global Measurement Study of Context-Based Propagation and User Mobility

Rod Meikle and Joseph Camp
Department of Electrical Engineering, Southern Methodist University (SMU), Dallas, TX
{rmeikle, camp}@smu.edu

ABSTRACT

Mobile phones are becoming a powerful platform for global-scale measurements due to their ever-increasing programmability and prevalence. Moreover, advanced sensing capabilities have allowed mobile phones to become aware of the user's context, potentially leading to performance improvement. Such context awareness could be exploited to optimize a wireless network connection since wireless channels are known to depend on the surrounding environment. The viability of context-aware wireless performance improvement would heavily depend on whether differences in context had meaningful performance distinction and whether the training overhead per context encountered would overwhelm potential gains. In this paper, we perform a large-scale measurement study of regional performance based on a users context and characterize user mobility around the world. To do so, we deployed WiEye, an Android-based wireless sniffer which has collected over 50 million measurements from over 30 thousand unique users. We categorize measurements according to land use and political divisions to investigate whether distinct levels of performance exist as indicated by wireless path loss. We then examine user mobility patterns via subtractive fuzzy clustering to determine how many different contexts a user typically encounters. Our results show promise for context awareness since distinct levels of performance are observed per land use class with only one or two contexts being typical per user.

Categories and Subject Descriptors:

C.2.1 [Computer-Communication Networks]: Network Architecture and Design-Wireless Communication

General Terms:

Measurement, Performance, Experimentation, Design

Keywords:

Wireless Propagation, User Mobility, Mobile Applications

1. INTRODUCTION

In a matter of days, mobile phone applications can be downloaded in any part of the world on a device that stays with the user at all times. Thus, there is a tremendous opportunity for gaining measurements of nearly any type of event on a global scale. Correspondingly, there has been a rise in mobile sensing in a wide vari-

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.

HotPlanet'12, June 25, 2012, Low Wood Bay, Lake District, UK.
Copyright 2012 ACM 978-1-4503-1318-6/12/06 ...\$10.00.



Figure 1: Over 50M measurements collected from WiEye, our published Android-based application (30k+ active users).

ety of applications from physical activity monitoring [1] via mobile phones [2], including studies on human mobility [3] and context-aware applications based on user behavioral patterns and locations [4]. In fact, context awareness has been leveraged on mobile devices to increase network performance [5,6] and aid social applications [7] via measurements on a large scale. Considering these works, a great potential exists for both technical and societal advances via context awareness as well as an opportunity for new perspectives on in-situ global wireless performance and usage.

Motivated by the potential gains of context-aware protocols and applications, in this paper, we perform a large-scale measurement study to broadly characterize wireless cellular network performance and user mobility patterns in key cities around the world. To do so, we use data collected from our deployed application (WiEye) which was downloaded by over 30,000 unique devices over 120,000 times and has logged over 50 million measurements. A spatial distribution of the WiEye data points is shown in Figure 1. During the course of our data collection and analysis, we have determined that trends in signal propagation exist among various similarly defined regions in different parts of the world based on land use. In particular, we use path loss to characterize the propagation characteristics of a given region [8]. A similar procedure could be used for on-the-fly estimation of regional performance for context-aware applications such as in-situ training to optimize performance [9] or enabling optimized mobile handoff procedures based on directionality, speed, and current distance and propagation characteristics from cell towers [10]. We also examine user mobility within these regions, using a fuzzy logic technique of subtractive clustering [11] to determine groupings of user locations. We show how multiple users display similar mobility patterns and display trends among locations and signal strength across similar regions. Such a user mobility characterization is useful to understand a typical smartphone user and could be used to trigger context-aware training in commonly traversed areas [9].

Our three main contributions are as follows. First, we introduce the measurement platform and methodology for capturing a

large number of data points from unique client devices around the world. As might be expected, most of the data points come from the world's most populated cities. Thus, later sections focus on context-based propagation characteristics and user mobility patterns from these cities. Smartphones have become highly prevalent in the developed world not only as a communications device but also a measurement study platform [12], especially Android-based smartphones [13]. Thus, we choose to capitalize on the Android-based platform and popularity of free, widely-available applications, offering users a functional WiFi scanner in exchange for their choice to contribute data to our database (covered by an IRB to protect our participant's data). Data sourced from installations of WiEye are submitted to a central database server to ensure that only fresh, valid data will be gathered from clients.

Second, considering the data we have collected, we study regional propagation characteristics and mobility patterns of users. In particular, since each measurement contains a GPS coordinate, we examine certain wireless characteristics on a region-centric level using the land use classification provided by the open, community-moderated database from OpenStreetMaps [14]. This overlay of basic land use allows us to create generalizations of: (i) regions which users are most commonly located, and (ii) propagation characteristics in the regions themselves. Despite not knowing the exact transmitter and receiver characteristics for the cell tower and user device, we are still able to examine wireless signal propagation characteristics for these regions using a perceived line-of-sight reference point and distance [8] (discussed in detail in Section 3). As a result, we examine characteristics such as path loss and shadowing in order to draw conclusions about their effects on the system.

Third, we examine user mobility patterns from our measurements. As explored in previous measurements of mobile phone users [12], such large-scale studies produce data that can be filtered by users to examine typical user mobility patterns. Along the same vein, we analyze the mobility patterns of our top measurement contributors (giving us the highest confidence of the mobility patterns). Similar work has been done in the area of user location [15], but they have not examined the wireless characteristics of the mobility patterns. This classification of users is achieved by characterizing where a smartphone user typically resides throughout the day/week and what signal strength is realized for the various user types. The GPS data collected also allows the examination of user speed and the possible effects it has on both signal strength and user mobility. We examine the top users in various geographical designations within a radius of urban centers and specific land uses. Surprisingly, we find that, to a large degree, user mobility patterns can be classified as unimodal or bimodal. Since most context-based services require an initial overhead, such an awareness of typical user locations could trigger training to be performed [9] in the typical environment in which they operate to allow large gains in performance for the given scenario.

2. ANDROID-BASED MEASUREMENT PLATFORM

Accurately measuring wireless characteristics of in-situ client devices is often a challenging, labor-intensive task. For example, in [16] we performed a concerted field trial effort over multiple months to obtain 138 measurements to plan the backhaul of a mesh network. By comparison, mobile phones offer multiple orders of magnitude greater numbers of measurements in the same amount of time due to the high level of penetration possible through mobile application markets. The Android Market, for instance, allows users with diverse hardware, wireless carriers, and user locations to gain access to the same smartphone software to extend the functionality of their devices. Thus, we released an application that offers users a beneficial service of sniffing WiFi and cellular net-

works to find optimal channel conditions in return for a chance that they might participate in our measurement campaign. Due to the limitations imposed by the Android platform, we were only able to utilize gathered information from users who were serviced by GSM-based cellular networks, as CDMA-based networks such as those run by Verizon and Sprint in the United States do not allow the reporting of unique network identification information as gathered by our usage of the Android API. However, the number of usable data points we collected, or those that include valid GSM cellular network identification (CID and LAC values), accounted for over 22 million measurements, greatly exceeding our previous in-field measurement efforts.

2.1 WiEye Android Application

The application we created for the data collection is currently available for download and usage via the Google Android Market under the name WiEye. The user interface is displayed in Figure 2 and shows the user experience of running the application. Users are able to view all Wi-Fi access points that are within range of their cellular device in both graphical and tabular form. All data collection is done in the background, either continuously while the user is running the application or periodically if the user has opted in to background data collection to assist our study. Our data collection has been approved by the Southern Methodist University Institutional Review Board, a human subjects research committee, ensuring that all ethical precautions have been taken in collecting data from the users of our application.

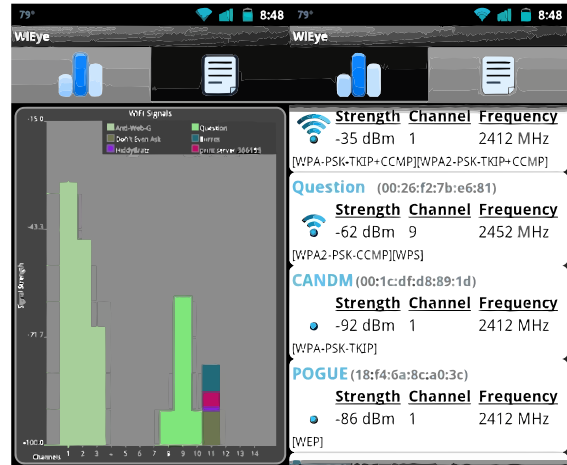


Figure 2: WiEye (Android-based) application interface, displaying the current activity level of different wireless channels.

2.2 Data Collection Methodology

Data from participating Android phones is reported by the application at intervals determined randomly by the Android device so as to submit data to the central database between 4 to 6 times per day (once every 4-6 hours) to help mitigate instantaneous server load to the database as well as to monitor a given region periodically throughout the day. We have made changes to the reporting frequency multiple times, as finding the ideal frequency to report data to the server is challenging. In determining the frequency of data reporting, we need to consider not only the need for timely data to maintain relevance to the current location of the user, but also that the user experience of having the application installed on their device does not cause the user to uninstall the application by using too much battery or data. The periodic nature of data reporting also serves to allow all regions to be defined via collecting a

large number of data points over time, as opposed to collecting a large set of potentially skew data in a single submission.

The data is sent to us from the Android application via a PHP script running on an Apache web host. The script accepts a HTTP request to store data from clients with the correct key value as well as a complete set of parameters that are expected by the database. This data is then stored into a MySQL database along with other measurements from the particular version of the application which submitted the measurements. This ensures that the database will be populated correctly and that our data is grouped relative to the information that is contained therein, as each version of the application reports slightly different data.

The reported data by the application has changed as we have progressed in our understanding of the Android platform's capability as a measurement device; we are currently storing data from the GPS as to the user's physical location in the form of a latitude and longitude point, the accuracy of this location, and the speed at which the device is traveling. We also capture data from the cellular radio in the form of the connected tower and visible towers to the device at the time of reporting, a unique identifier for the device, the cellular carrier on which the device is operating, the signal strength of the connected tower and other nearby visible towers, as well data from the phone itself in the form of the system software version that is running and the hardware identifier of the device. The data collected in the initial release of our application utilized only the triangulation functionality available to cellular devices with no GPS hardware, which on Android is aided by Google's servers in determining location. The accuracy is reported for each location measurement, and with the addition of GPS hardware functionality, our average accuracy estimate for any given location dropped from a 1219 meter radius to a 469 meter radius.

Our database currently consists of over 50 million reported cell tower instances with client locations where nearly half of the measurements have GPS data for client location. Devices lacking GPS hardware capabilities use a combination of cell tower triangulation and a less accurate proprietary assisted positioning system provided by Google to users of Android called "My Location" [17]. The application is available for public use via Google's Android Market and currently has over 130 thousand unique users worldwide.

3. CONTEXT-BASED PERFORMANCE

In this section, we examine characteristics of our measured wireless signal data across multiple, arbitrarily defined designations provided by OpenStreetMap, which is an open, community moderated database of land use classifications and political boundaries [14]. Use of a database such as this allows us to focus on the location of data points received from users with respect to users surroundings. Also, we investigate trends presented by our data in this context: user density in a region, expected signal strength for a region, and cell tower density within regions.

3.1 Characterizing Propagation Per Land Use

The region designations offered by our chosen region database include multiple specific region types, but for this work we will be classifying regions into the following subgroupings: (i) highly populated regions with high building heights and density (referred to as "commercial"), (ii) highly populated regions with low building heights and density (referred to as "residential"), (iii) relatively low population with low building heights and density (referred to as "park"), and (iv) other regions for which insufficient data is available to characterize. After making these classifications, our region groupings can be visualized as follows.

As indicated in Figure 3, the majority of our data falls into residential, commercial, and park classifications. The classifications for our land use come from the OpenStreetMap database, and are

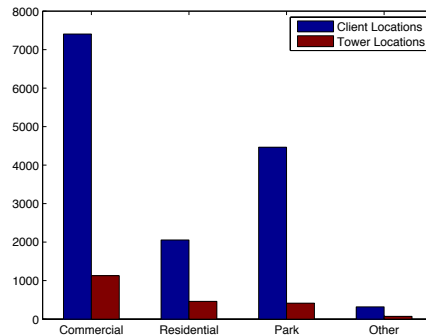


Figure 3: Total data from typical land use classes in London.

grouped at this high level to obtain sufficient data points from which we can draw conclusions. The commercial classification includes the following designations from the database: building, city, office, postcode, shop, and highway (with sub-designations of primary and motorway). Our residential classification includes the following designations: house, suburb, town, village, highway (with sub-designations of residential and suburb). Our park classification includes: amenity, county, historic, tourism. The amenity designation includes a number of auxiliary buildings and features located in many differing areas, however the majority of them share similar signal propagation characteristics, and without grouping them we would have insufficient data to draw accurate conclusions.

Regional Path Loss Calculation. In characterizing our regions, we wish to examine wireless signal propagation within each region, considering a unique reference power and distance for each cell tower. Even though the transmitter characteristics of the cellular towers such as power and specific antenna configuration are unknown to us as cellular providers do not publish information of that nature, we are able to make a determination as to the path loss seen by the clients of each tower using the following procedure. We consider the path loss exponent α that best describes the relationship of each signal strength measurement $P_{dBm}(d)$ at a given distance d according to the following equation found in [8].

$$P_{dBm}(d) = P_{dBm}(d_0) - 10\alpha \log_{10} \left(\frac{d}{d_0} \right) + \epsilon \quad (1)$$

We utilize a reference power p_0 from a perceived line-of-sight path at a reference point with a distance d_0 for all regions. The point is found by considering the data point with the largest euclidean distance in the positive direction from the minimum-mean-square curve fit of all measurements from a given cell tower. Each cell tower has a unique reference distance and power to account for the unique transmitter characteristics of the cell tower. Most importantly, this technique allows us to abstract away these unique transmitter characteristics which are often unavailable for the thousands of cell towers in our database. We then calculate the path loss α for each tower using the data for all client measurements in a given region surrounding the tower with consideration of the shadowing component ϵ and similar to our technique in [16]. We treat towers with clients in two or more regions as separate towers and calculate the path loss for each set of clients, assuming both sets of clients will lead to a systemically different path loss and to be able to include the results in each set of region characterizations.

Figure 4 shows path loss distributions per region within the Greater London area. As indicated by Figure 4(a), the commercial region displays an average path loss value near 4.5, which is expected for an area heavily obstructed by buildings. Figure 4(b) describes towers in the residential areas of Greater London, and has an observed average path loss value around 4. This is slightly unexpected; previous studies have measured the path loss for residential areas with




Figure 4: Gaussian distribution of path loss exponent within each regional classification in London.

a low average building height closer to 2 or 3, as there are fewer obstructions between the base station and the client radio. Our observed higher average may exist due to inconsistencies in land usage classification at such a broad level, causing us to consider non-exclusive residential areas in our measurements, or due to regulations regarding placement of cellular base stations in residential regions of the city. However, the residential region still exhibits a lower path loss and thus better performance than the commercial region, which is expected. Figure 4(c) shows the path loss distribution for the parks and amenities in Greater London. The average path loss in this region is close to 4 as well, but a number of varying land use designations in this region exist, leading to a larger degree of inaccuracy than other regions. The park region also does not follow the normal shaping of the other two regions, most likely due to multiple region distributions inclusion that we are unable to separate due to inaccuracies in our land usage database. Clearly, the path loss exponents shown in these figures depend on location, however similar types of regions as defined by these figures can show differing data across multiple cities due to the impact of characteristics specific to the individual locations. These path loss exponents also range from 2 to 5 in line-of-sight to dense urban areas, and thus when trying to predict mobile handoff, there can be significant skew when not using data from the precise environment the handoff would be occurring within.

Finding Cell Tower Locations. Client data throughout the regions, however, only gives limited insights, as the more important metric to define the propagation of regions is data concerning the towers located within the regions. The data we have accumulated for towers comes from Google’s AGPS database, which is publicly available for query for any mobile device visiting the Google Maps website that is capable of reporting a cell tower identifier, or CID, as well as a location area code, or LAC to Google. Our database of cell tower locations is based on responses from this service based on submitting CID and LAC values for all reported towers falling within a realistic distance (9km) from the client claiming to be in range of any given tower. Using these tower locations along with the client GPS data, we construct an approximation of the locations of the large number of cellular towers visible to the clients.

Analyzing these towers with respect to the clients reporting SNR values for them allows us to construct an approximate path loss and shadowing for the towers in each region, as seen in Figure 4. The path loss values shown here indicate that propagation is stronger in regions that have larger populations and lower average building height and density. This occurs consistently across the majority of regions when analyzed individually, and corresponds as well to these regions having a higher tower density as seen in Figure 3. These path loss measurements overall help us to generalize the signal propagation characteristics for each region. These measurements, while based on signal strength alone in this case, extrapolate to make generic assumptions concerning cellular link performance, as indicated in previous studies [18], signal strength is a relatively accurate indicator in most networks as to performance.

3.2 Regional Propagation in Key World Cities

After considering regions on the city-level scale, we proceed to examine multiple cities across the globe in a similar fashion, as our client and tower location methodology as well as our land use classification database both are either dynamically or community updated, and as such contain data for regions in different cities around the world. We examine the average path loss exponents for our region types in five large cities across multiple continents.

	London	Moscow	L. A.	Beijing	Sydney
Commercial	3.983	3.233	3.887	3.507	3.732
Residential	3.587	3.536	2.962	3.141	1.298
Park	3.890	3.854	4.983	3.731	3.414

Table 1: Path loss exponents for regions in key world cities.

As seen in Table 1, the path loss for region types worldwide varies greatly, but displays notable characteristics within each city between the regions. For most cities, the path loss is commonly best (lowest) in residential areas. The path loss in park areas is variable, as our designations include multiple types of building within the park region in addition to traditional free space areas, making the park designation directly comparable in a given city. Also, cities with a larger density of buildings in their central areas tend to exhibit a higher path loss exponent in their commercial region, due to the worsened signal propagation seen by users.

4. USER MOBILITY PATTERNS

User mobility and user context, or the generalized location of a user, can also be examined using our collected data. User location is a very important component of wireless performance, and as discussed earlier, has been examined before in detail [3], but not on consumer hardware at the scale examined in this paper. As we examined and categorized regions, we can also use our data to examine and categorize users both on a specific level, for those with a sufficiently large amount of reported data, and on a general level, using our specific users as a guide to the behavior of other users of our application and smartphones in general. This classification of user mobility patterns gives a basis for examining training algorithms based on which regions and for how long a device can expect to exist within a region.

4.1 Considering Users with Most Measurements

To accomplish this categorization, we examine the same Greater London region as explored previously in Section IV, but now approach collected data from the aspect of individual users and their mobility over time. We examine the top ten contributing users in an attempt to classify mobility patterns on a per user level.

Figure 5 shows all reported data points from the top ten contributors from the area. At first, this graphical representation appears to be contain a relatively small amount of data upon initial inspection, but this map contains a number of small groupings of points not



Figure 5: Measurements from top ten most frequently reporting users in London.

visible at this level, in total 9709 data points in very small groupings. Upon realizing this, user location distribution modes can be seen to follow specific clustering distributions.

We classify a user’s position distribution using the subtractive clustering method [11], a process using fuzzy logic methods to determine central points of location clusters from a user’s reported measurements. These points are then used as a central location, and surrounding data points are counted to observe the percentage of points within a cluster to all reported data from the specific user. Data for the top 5 London users can be seen in Table 3. Commonly, user locations appear to be similar to distributions with one or two modes, unimodal or bimodal respectively, as users tend to make the majority of measurements from clustered locations.

	Primary Location	Secondary Location
User 1	0.7173	0.2259
User 2	0.8530	0.1335
User 3	0.9818	No Cluster Found
User 4	0.9891	No Cluster Found
User 5	0.8665	No Cluster Found
User 6	0.9335	No Cluster Found
User 7	0.6210	0.3676
User 8	0.9979	No Cluster Found
User 9	0.5816	0.2667
User 10	0.9732	No Cluster Found

Table 2: Top ten users in London and their percentage of data reported from each user’s most common locations.

Considering these ten users as archetypes for most users of the application, and by extension smartphone users in general, we can classify users into certain common mobility patterns: unimodal as shown in Figure 7, or users having a single location in which they report all of their data points, bimodal as shown in Figure 6, or users reporting a majority of data points from two distinct locations and in some cases a transit path between locations, and transient, or users who have too few data points to accurately categorize (as many users uninstall the application after submitting only a few data points). These classifications allow the application of context aware methodologies, as they provide a common usage pattern that can be expanded upon in future works to reduce unnecessary training or even specify parameters for training algorithms.

In examining the data for Greater London users in Table 2, we determine that the majority of users are bimodal, spending time mainly in two diverse locations. Also notable is that all of the users displaying unimodal characteristics, while fewer in number than the bimodal users, are located within residential areas, and of the

users exhibiting bimodal characteristics, all have one mode in a residential area and the other mode in a commercial area. These users appear to post the majority of their data during weekday daytime hours in commercial areas and at all other times post the majority of data in residential areas, reinforcing the traditional work day schedule. Examination of this data by time period yields similar results to what is displayed in Table 2 for weekday data, whereas during weekends, most bimodal users tend to adopt the unimodal pattern. Examining data by device type within London yields very similar results between devices, as the majority of reporting users tend towards a very limited number of devices. A device-centric analysis may be a topic for future work when more data can be collected from new devices that enter the consumer market.

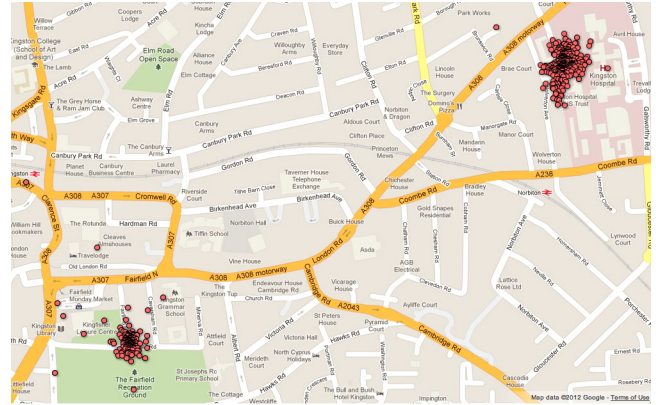


Figure 6: Bimodal mobility pattern in London (User 7).

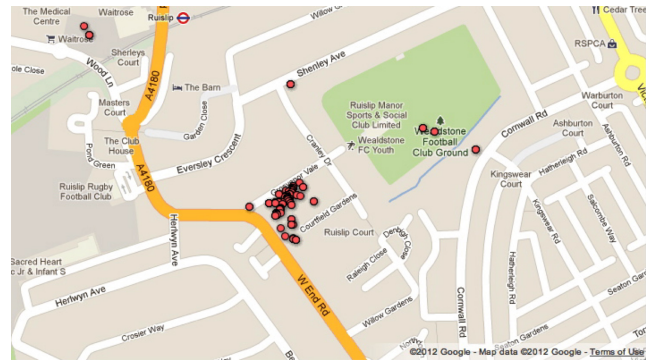


Figure 7: Unimodal mobility pattern in London (User 10).

4.2 User Mobility in Key World Cities

As we have now established a perspective on common user mobility patterns from one large metropolitan area, we can examine other cities around the world as most users exhibit similar mobility patterns within cities. Table 3 assists us in attaining this perspective, describing the top five commonly reporting users for five cities across multiple continents.

As shown in Table 3, the majority of users belong to the bimodal classification, with two cluster centers each. A visualization of a bimodal user’s data set and thus mobility pattern from the Greater London area can be seen in Figure 6. Users in this designation are common worldwide, as indicated by Table 3, bimodal characteristics are the most common for any user regardless of city.

In comparison, some users adhere to a unimodal mobility pattern, commonly not reporting data outside of a residential area. This pattern can be visualized in Figure 7, as exhibited by a user in the Greater London area. This designation is the second most common location distribution for a user to follow, which is intuitive in

	London	Moscow	L.A.	Beijing	Sydney
User 1					
Primary	0.7173	0.7515	0.9875	0.8630	0.9463
Secondary	0.2259	0.1949	-	0.1671	0.9463
User 2					
Primary	0.8530	0.8186	0.2671	0.6344	0.6387
Secondary	0.1335	0.1718	0.1845	0.3584	0.2156
User 3					
Primary	0.9818	0.8530	0.6114	0.8310	0.9559
Secondary	-	-	0.0616	-	-
User 4					
Primary	0.9891	0.6208	0.9998	0.7245	0.6533
Secondary	-	-	-	0.2755	0.1267
User 5					
Primary	0.8665	0.6708	0.9895	0.9693	0.8169
Secondary	-	0.1904	-	-	0.1408

Table 3: Percentage of time spent in users' most common locations (primary and secondary) for the top-five, total-measurement users in key world cities.

that while the majority of the world population using smartphones works at a separate location from where they live, the second most common group either works from home or is otherwise occupied in the same place they live, and has only a single cluster of points. As expected, none of the top ten contributors are classified as transient. This classification is reserved for the users who only submit a small number of data points via the application or uninstall the application after a short time of running it.

5. RELATED WORK

Work has been done in the area of wireless propagation studies across various region types [8,16,19]. Many studies have also been done on mobile devices in daily use as shown in multiple studies [7,12,20] as well as on cellular networks in controlled environments [18]. In these works, wireless performance and channel quality depend on the environment of transmitter and receiver. Moreover, as indicated in [8], wireless propagation can be characterized by path loss and shadowing in different environment types. The exact relationship between the path loss exponent and shadowing have been studied for specific environments and transmitter/receiver characteristics, such as our previous work [16]. In contrast to prior studies that have depended on exact cities, types of environments, and transmitter/receiver characteristics, we show that using an overlay of publicly available land use maps and heterogeneous mobile phone and cell tower characteristics, similar propagation characteristics can be found per region without the excessive measurement time and infrastructure previously necessary [16]. In fact, our proposed technique can be used on-the-fly by users in specific contexts according to the needs of their application.

Other smartphone based work has been done in the field of mobile measurements, including both studies on the identification and categorization of vehicle traffic using data from mobile devices [21] as well as a BlackBerry-based study to examine user behavior and the challenges of collecting user data from smartphone users [12]. Previous studies that have been done to examine properties of cellular networks [18] measure data in a more controlled environment. In contrast, by providing an application offering relatively unique functionality to the user that is readily available and presented, without cost, publicly to all users of the Android Market, we leverage a large user base to gather measurements via programmable smartphones in the field with typical use. Moreover, our methodology allows a highly-scalable study of regional propagation and user mobility within a framework for context-aware applications.

6. CONCLUSION

In this work, we performed a large-scale measurement study of regional performance based on a user's context and characterize user mobility in key cities throughout the world. To do so, we deployed WiEye, an Android-based wireless sniffer which collected over 50 million measurements from over 30 thousand unique users. We first categorized measurements based on land use and found that distinct levels of performance exist as indicated by wireless channel path loss. We then examined the user mobility patterns via subtractive fuzzy clustering and determined that unimodal and bimodal user mobility patterns typically occur. In future work, we will study the role of urban versus suburban on our regional classifications and leverage our measurement framework for context-aware applications such as mobile handoff and rate adaptation.

7. ACKNOWLEDGMENTS

The authors would like to thank Amir Ghadiry and Brian Tanous for their work in the initial design and release of WiEye as part of their undergraduate senior design project. This research was supported by NSF grants CNS-0958436 and CNS-1150215.

8. REFERENCES

- [1] C. Otto, A. Milenkovic, C. Sanders, and E. Jovanov, "System architecture of a wireless body area sensor network for ubiquitous health monitoring," in *Journal of Mobile Multimedia*, Huntsville, AL, January 2006.
- [2] N. D. Lane, E. Miluzzo, H. Lu, D. Peebles, T. Choudhury, and A. T. Campbell, "A survey of mobile phone sensing," *IEEE Communications Magazine*, September 2010.
- [3] P. Hui, R. Mortier, T. Henderson, and J. Crowcroft, "Planet-scale human mobility measurement," in *ACM HotPlanet*, San Francisco, CA, June 2010.
- [4] "Sensay: A context-aware mobile phone," in *Proceedings of the 7th IEEE International Symposium on Wearable Computers*, White Plains, NY, October 2003.
- [5] A. Hassawa, N. Nasser, and H. Hassanein, "Tramcar: A context-aware cross-layer architecture for next generation heterogeneous wireless networks," in *Proceedings of the IEEE ICC'06*, Istanbul, Turkey, June 2006.
- [6] L. Ravindranath, C. Newport, H. Balakrishnan, and S. Madden, "Improving wireless network performance using sensor hints," in *Proc. of USENIX NSDI*, San Jose, California, Apr. 2010.
- [7] T. Xu, Y. Chen, J. Zhao, and X. Fu, "Cuckoo: Towards decentralized, socio-aware online microblogging services and data measurements," in *ACM HotPlanet*, San Francisco, CA, June 2010.
- [8] T. Rappaport, *Wireless Communications, Principles & Practice*, ser. Emerging Technologies Series, T. Rappaport, Ed. Upper Saddle River, New Jersey: Prentice Hall, 1996.
- [9] J. He, H. Liu, P. Cui, J. Landon, O. Altintas, R. Vuyuru, D. Rajan, and J. Camp, "Design and experimentation of context-aware link-level adaptation," in *Proc. of IEEE INFOCOM*, Apr. 2012.
- [10] A. Giannoulis, M. Fiore, and E. W. Knightly, "Supporting vehicular mobility in urban multi-hop wireless networks," in *ACM MobiSys*, 2008.
- [11] S. Chiu, "Fuzzy model identification based on cluster estimation," *Journal of Intelligent and Fuzzy Systems*, vol. 2, no. 3, September 1994.
- [12] E. Oliver, "The challenges in large-scale smartphone user studies," in *ACM HotPlanet*, San Francisco, CA, June 2010.
- [13] ComScore, "2012 Mobile Future in Focus," <http://www.comscore.com/2012/MobileFutureinFocus>, February 2012.
- [14] "OpenStreetMap Foundation," www.openstreetmap.org, March 2012.
- [15] N. Patwari, I. Alfred Hero, M. Perkins, N. Correal, and R. O'Dea, "Relative location estimation in wireless sensor networks," *IEEE Transactions on Signal Processing*, vol. 51, no. 8, pp. 2137–2148, August 2003.
- [16] J. Camp, J. Robinson, C. Steger, and E. Knightly, "Measurement driven deployment of a two-tier urban mesh access network," in *ACM MobiSys*, Uppsala, Sweden, June 2006.
- [17] Google, "My Location Service," <http://www.google.com/mobile/gmm/mylocation/index.html>, March 2012.
- [18] X. Liu, A. Sridharan, S. Machiraju, M. Seshadri, and H. Zang, "Experiences in a 3g network: Interplay between the wireless channel and applications," in *ACM MobiCom '08*, San Francisco, CA, September 2008.
- [19] G. Stuber, *Principles of Mobile Communication*, 4th ed. Boston: Kluwer, 2000.
- [20] N. Balasubramanian, A. Balasubramanian, and A. Venkaramani, "Energy consumption in mobile phones: A measurement study and implications for network applications," in *ACM IMC'09*, Chicago, IL, November 2009.
- [21] J. Yoon, B. Noble, and M. Liu, "Surface street traffic estimation," in *ACM MobiSys*, 2007.