Personalized Energy Consumption Modeling on Smartphones

Yifei Jiang[1], Abhishek Jaiantilal[1], Xin Pan[1], Mohammad A.A.H. Al-Mutawa[1], Shivakant Mishra[1], and Larry Shi[2]

[1] University of Colorado at Boulder, CO 80309, U.S.A.
{yifei.jiang,abhishek.jaiantilal,shivakaht.mishra}@colorado.edu
[2] University of Houston, TX 77204, U.S.A.
larryshi@cs.uh.edu

Abstract. Energy has emerged as a key limitation in smartphone usage. As a result, optimizing power consumption has become a key design issue in building services and applications for smartphones. Understanding user behavior and its impact on energy consumption of smartphones is a key step for addressing this problem. This paper provides an in-depth study of user behavior and energy consumption of smartphones by analyzing smartphone data collected from twenty smartphone users over a period of three months. In particular, correlations between power consumption and factors such as time of day, user's location, remaining battery power, recent phone usage history, and phone's idle and active states have been studied. The results show varied levels of correlations between a user's phone usage and these factors, and can be used to model and predict smartphone power consumption.

Keywords: mobile computing, power usage, mobile power usage, user study.

1 Introduction

Smartphones are becoming an amalgam of multiple services, including online, social, video, music, gaming, payments, location-based services, and augmented reality. Because of such multifarious functionalities and applications, energy consumption of smartphones is rapidly increasing and becoming highly variant. For example, a typical smartphone features around five hours talking time on 3G and 150 hours standby time. However, these figures drop sharply when users play games, surf the web, or use GPS navigation. Because of the high and variant energy usage, battery lifetime of smartphones is getting shorter and much more unpredictable. As a result, better energy management techniques for smartphones are needed.

An understanding of energy consumption of smartphones as well as user behavior that impacts this energy consumption is a key requirement for designing effective energy management techniques. For example, if we can predict that a user will be running some power-hungry application on their smartphone over

D. Uhler, K. Mehta, and J.L. Wong (Eds.): MobiCase 2012, LNICST 110, pp. 343–354, 2013.

the next two hours and they will be at a place where there is no source for charging the phone, the system can then recommend the user to charge their phone immediately. Another example: power-aware scheduling mechanisms can be designed to pause non-urgent applications and services if we can predict that the user will be running some power-hungry application soon. Finally, knowledge of a smartphone's power consumption behavior can aid in detecting abnormal conditions (like a malware is running on the phone).

The main goal of this paper is to study the relationship between user behavior, their environment and energy consumption. Our study is motivated by several important observations. First, user location plays an important roles in phone usage. For example, users may surf the web or read emails when staying at an airport, but they may play a game when they are at home. Second, the amount of battery power remaining in the phone impacts user behavior. For example, a user is less likely to play (power-hungry) games if the remaining battery lifetime is quite low. Finally, different locations have different communication environments due to varying GSM signals and availability of WiFi access points.

This study of user behavior and its impact on energy consumption of smartphones is done by analyzing user behavior and power consumption data collected from twenty smartphone users over a period of three months. In particular, we study correlation between power consumption and factors such as time of day, user's location, remaining battery power, recent phone usage history, and phone's idle and active states. The results show varied levels of correlation between a user's phone usage and these factors, and this information can be used to model power consumption of a user's smartphone.

The rest of the paper is organized as follows. Section 2 briefly discusses the important related work. Section 3 describes our data collection methodology. Section 4 presents the data we have collected. Four factors that impact energy consumption of smartphones have been identified: time of day, user location, remaining battery life, and user's recent phone usage history. An analysis reveals that energy consumption pattern varies based on when a phone is idle and when it is active. Section 5 provides this analysis. Finally, Section 6 provides a discussion and concludes the paper.

2 Related Work

Earlier research has shown that battery consumption and charging vary in context [14,4,2,11,12]. Trestian *et al.* [14] examined the correlation between location and application usage from 280,000 clients of a 3G mobile network and found that user interactions with mobile phone vary according to users' location. Similarly, Banerjee *et al.* [2,11,12] studied battery charging behaviors on mobile systems and noted that most charging behaviors were driven by context, including location and time of day. Falaki *et al.* [4] conducted a study of user interactions with mobile phones and examined the corresponding energy drain rates. They found that considerable variation among user activities resulted in diverse energy consumption patterns, which indicates that learning and adapting

Table 1. Data collection on smartphone

Data	Sample rate
WiFi fingerprint	Every 10 minutes
Battery current	Every 5 seconds
Remaining battery	Call back
GPS	Call back
CPU load	Every 1.5 minutes
Screen status	Call back
Phone calls	Call back
Time stamp	For each data log

to personalized user behavior is likely to be more effective for improved mobile energy management.

In a different context, extensive research has also been done towards understanding power management at the hardware level. Simunic *et al.* [13] provided system-level dynamic power management algorithms to schedule idle components into lower power states automatically. Carroll *et al.* [3] conducted exhaustive measurements and analysis on different components of mobile devices.

There are in general two categories of location sensing algorithms – geometry and fingerprint. Geometry based algorithms [1,6] cluster geo-coordinates that belong to the same meaningful location into the same cluster. The location is recognized by checking whether the device's current geo-location falls into the geometric shape of any known location. Assuming that the radio fingerprints of locations are rather stable and unique, fingerprint based approach [5,9,15,7,8] can tell whether two locations are close or far apart according to the similarity between two fingerprints. SensLoc [8] uses accelerometers to detect movement and turns off GPS/WiFi sampling when users are stationary in order to save energy, thereby making the location sensing more energy efficient. Besides location sensing, iLoc [10] provides a predictive model to forecast future location-state transitions.

3 Data Collection

We developed a data collection application for two smartphone platforms: Nokia N900 smartphone[1] running Maemo Linux and smartphones running Android[2]. The application is developed for monitoring user's power usage and their context information such as location information, battery status and time. Furthermore, the usage information of each hardware component, e.g., GPS, screen, and CPU, are also collected by the application. The data and sample rate we used are listed in table 1. The energy overhead is less than 5 mw.

[1] http://ml.cs.colorado.edu/~abhi/mobien/

[2] https://market.android.com/details?id=edu.colorado.mobien

To determine location, we used information from both the cellular sites and the WiFi access points [8]. The cellular information (GSM Cell ID and CDMA Base station ID) provide us with a coarse location, while the WiFi scan information provide us with a finer location. Note that we do not identify the actual geographical location of the smartphone (or the user), as we are mostly interested in differentiating among the different locations that the user visited. We did not use GPS to determine location because it consumes more power and does not work inside buildings.

Over a period of three months, we collected data from twenty users who installed our application. Ten of these users are known to us and they installed our application on our request, while others installed our application by picking up from the Android Marketplace. There were six Nokia N900 user and fourteen Android users. Among the users known to us, some are graduate students and others work in industries. Of the twenty users, data collected from six users (three Android and three Nokia users) was most complete, i.e. it spanned the entire three-month period. We have chosen the data collected from these six users for our detailed analysis reported in this paper. The data collected from the remaining fourteen users spanned from two weeks to less than three months. All observations that we report in this paper do hold for these data as well. As we are primarily interested in usage patterns of users with reference to time and location, we have excluded portions of the data collected when the phone was charging and when a user was traveling between locations. The charging and discharging behavior was caught by a system API of the OS platforms. We were able to distinguish whether the user was stationary or traveling between locations by observing the change in the Wifi access points.

4 Phone Usage Analysis

The overall energy consumption of a smartphone includes the energy consumed during the phone's *active* as well as *idle* states. In active state, smartphones provide one or multiple types of information to user, including screen output, audio ouput and vibration ouput; and in idle state, smartphones do not provide any information to user, but the device may run backstage process such as GSM communication and customized data collection application. In the active state, a user is running some application and using the phone. Energy consumed during the active state is on average ten times higher than the energy consumed during the idle state. Thus active state is mostly responsible for the overall energy consumption. After a detailed analysis of the power consumption data collected, we have identified four factors that impact smartphone usage. These factors are: time of day, user location, remaining battery level, and recent phone usage history. We first analyze these factors, and then in the next subsection we explore how these factors affect energy consumption in active and idle states of the phone.

We first look at the percentage of the time that phone is in use across 20 users in our dataset. As showed in Figure 1, the y-axis shows the fraction of phone

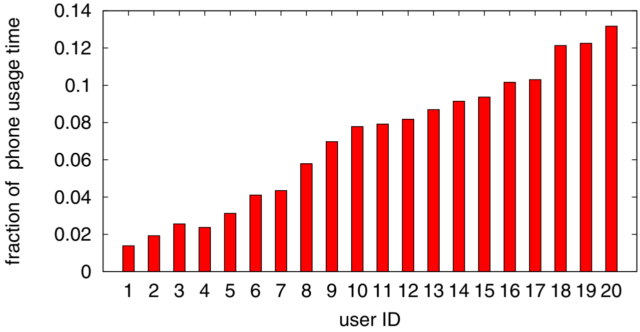

Fig. 1. Correlation between Time and Phone usage

usage time and x-aixs is the different user IDs sorted by phone usage percentage. The plot shows that 1) most of the users did use their phone actively, and on average they used their for 90 minutes per day. 2) the phone usage profile are very different across different users.

4.1 Time of Day

Logically, phone usage is related to users' daily schedule and routine activities and thus likely to be highly correlated with time of day. For example, a user may always sleep around 11 PM and get up at around 7 AM. Therefore, his/her phone is likely to be, everyday, in an idle state from 11 PM to 7 AM. On the other hand, the phone usage is typically high during some routine activities, e.g. during lunch time, users may often use their phones to check email.

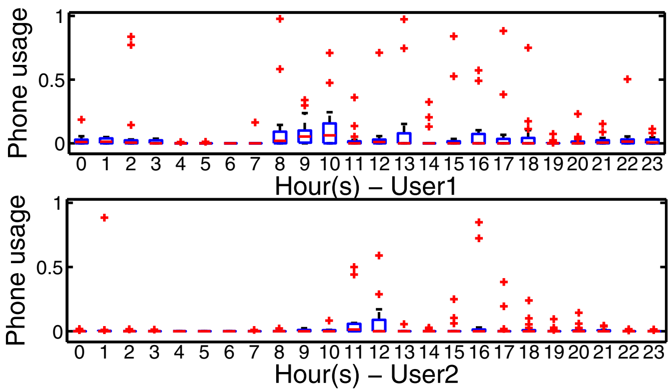

Fig. 2. Correlation between Time and Phone usage

Figure 2 shows the box-plot of phone usage over the 24 hours of a day for two users in our three month experiment. The phone usage is calculated as a percentage of the phone's active time during an hour. Red plus signs in this plot show the outliers. For example, for user 1 in hour 2, the percentage of time that the phone is in active state is below 3% all the time during the three-month period except for the three instances that are shown as outliers. An important observation we made from this figure is that the usage pattern for a user is generally quite uniform (smaller-sized box with low number of outliers). Although, this figure presents the data collected from only two users, we have also observed similar behaviors from all the twenty users based on the data we collected. This implies that a user's phone usage does not vary significantly everyday (over different hours). Thus, it is possible to model a user's smartphone power consumption based on the time of day and also make fairly accurate prediction about how much power will be consumed in the future. Of course, the presence of some, albeit small number of outliers as well as large-sized boxes for certain hours indicates that there is a possibility of error in this prediction especially when the prediction is made for those hours.

Also, Figure 2 shows that the phone usage of different users follow different patterns over time. For example, user 1 has high phone usage between 8 AM to 10 AM, while user 2 has high phone usage between 11 AM to 12 PM.

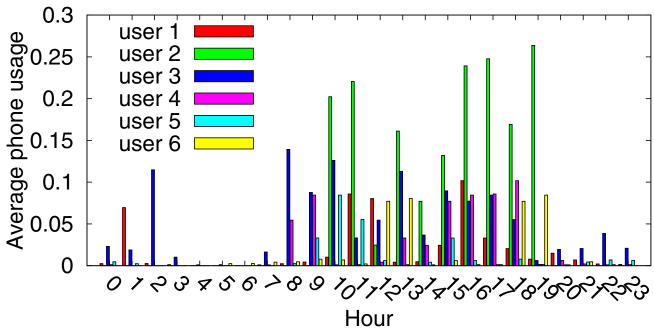

Fig. 3. Correlation between Time and Phone usage

Figure 3 shows this difference more clearly. Here we depict the average phone usage of six users over 24 hours. This figure indicates that different users have different usage pattern over time. Therefore, a personalized analysis is required.

4.2 Location

Location is another important factor that affects a user's phone usage because locations are typically indicative of the user's activities and associated phone usage. Figure 4 shows the box-plot of phone usage at different locations. This figure includes the data collected from six users. Note that we consider the same physical location visited by two different users as two separate locations in this figure because different users may have different behaviors at the same location. The phone usage is computed for each visit to the location.

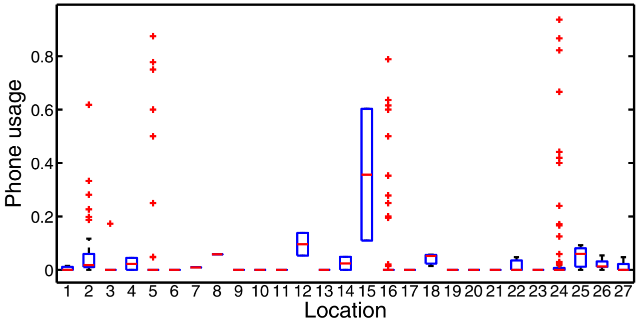

Fig. 4. Correlation between Location and Phone usage

We identify three types of location-based patterns when users visit different locations. First location-based usage pattern: phone usage is almost constant (smaller-sized box) for each visit. The phone usage at locations 4, 6-14, 17-23, etc. follow this pattern. This type of pattern happens at the locations where users have similar behaviors for each visit, or users never use their phones at those locations. Second location-based pattern: phone usage is similar for most visits but occasionally deviates from the normal profile (small-sized box and some outliers). The phone usage at locations 2 and 16 follows this pattern. The outliers here are typically caused by emergent or irregular events, such as an infrequent call. Third location-based pattern: phone usage varies a lot over different visits to the same location. The phone usage at location 15 follows this pattern. Thus, a power consumption model based on a user's locations that considers either the first or the second location-based pattern can make fairly accurate predictions about how much power will be consumed based on where a user is located.

4.3 Time of Day and Location

While both time of day and location correlate quite well with power consumption, we note that there is a possibility of a higher correlation between power consumption and (time of day, location). The intuition behind this is that a user may use his/her phone for different purposes at the same location at different times. For example, a student may check emails during daytime in the university, while play games at night. Figure 5 shows the box-plot of phone usage

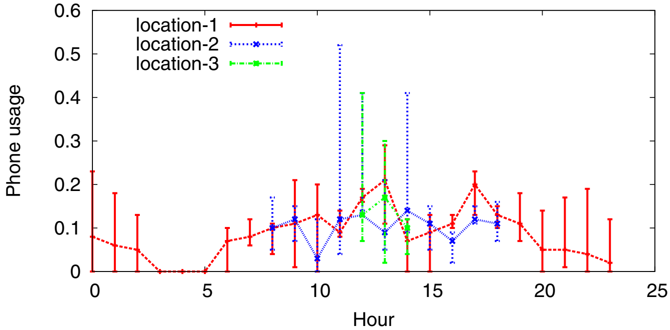

Fig. 5. Correlation between phone usage and (location, time of day)

for one user[3] at different locations and time. The figure shows three locations that the user frequently visited during the study period. Location 1 is most likely the user's residential place. Meanwhile, the user spent significant amount of time during day time at location 2 and location 3. Comparing location 2 with location 3, location 3 has a much narrow visit window. Figure 5 does not show locations with stay time less than one hour. According to figure 5, the user visited certain locations during predictable hours of a time (e.g., location 3). Figure 5 suggests that phone usage varied during the time when the user was at a particular location. Phone usage was more evenly distributed across time at residential place than other locations. During busy hours of a day (11am to 2pm), there was generally more phone usage and larger variance in usage pattern.

4.4 Battery State

We now consider battery state collected from the current battery level displayed on smartphone UI. In our analysis, we use eight discrete battery states (1-8) to examine the correlation between battery state and phone usage. Battery state 1 indicates that the battery is close to empty, and battery state 8 indicates that the battery is fully charged.

Figure 6 shows the correlation between the battery state and phone usage for six users. As shown in the figure, the average phone usage decreases with a decrease in the remaining battery level. This behavior is uniform across all users; it clearly indicates that a user's behavior in terms of the running applications is affected by the current state of the phone battery. Phone usage is low when the remaining battery level is low because 1) users want to keep the phone alive till the next charging; and 2) the actual remaining battery time is hard to estimate so that users tend to rely on the battery state provided by phone UI. This indicates that remaining battery time can be used to predict user behaviors in terms of if and how much they will use their phones.

[3] Data from the other users show similar trend. However, due to page limit, we cannot incorporate each user's location-time figure.

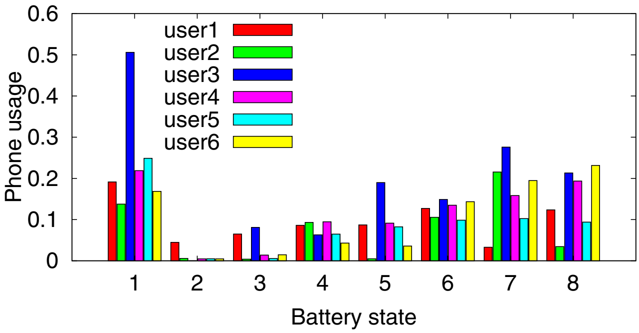

Fig. 6. Battery state vs Phone usage

The main exception here is when the battery state is 1, i.e. when the remaining battery time is extremely low. In that state, phone usage is actually higher than all other states for all six users. We traced this abnormality to the system periodically providing the user an alarm for low-battery, e.g. turning on screen, making a sound, vibrating, etc., which made the system active.

4.5 Recent Phone Usage History

Recent phone usage history may be used to predict a user's current phone usage. This hypothesis is based on our observation that a user's phone usage pattern typically doesn't change abruptly. So, users who have been using their phones quite actively in the recent past will continue to use their phones actively in the near future and vice versa for users who do not use their phones actively.

In order to explore this hypothesis, we analyzed the data collected from consecutive days for each user. If a user has a 30 day history period, we consider 29 pairs of data points for days, i.e (1,2), (2,3), ..., (29,30). For each user and a given average phone usage value on the previous day, we find the average phone usage value for the next day; e.g. for value x=0.1, we find all the tuples for which previous day value was 0.1 and then find the average of the next day's phone usage which we denote by y. We plot the (x,y) value correspondingly on the x/y-axes in Figure 7.

As shown in Figure 7, for most users, if the phone usage on previous day was higher, the phone usage on the current day will be, on average, higher. An interesting observation shown in the figure is the following: if a user has very low phone usage (e.g. < 0.02) on one day, usually, the phone usage would go up significantly on the next day.

5 Energy Consumption Analysis

5.1 Idle State

Energy consumption in idle state is dependent on two factors: programs running in the background and the ambient signal quality. Background programs directly

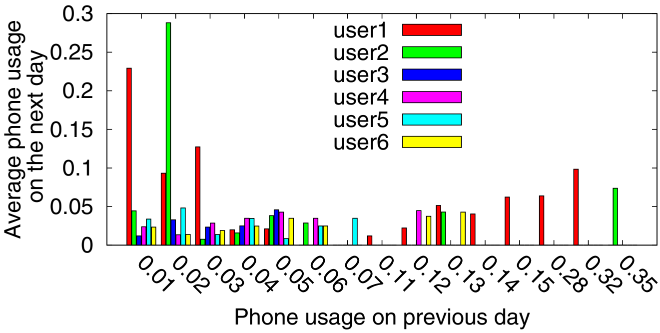

Fig. 7. Recent (Phone usage) history vs Phone usage

determine the usage/activeness of hardware components, e.g. CPU, networking and sensors; whereas, ambient signals affect power consumption of radio components, e.g. poor GSM/CDMA signal strength causes higher energy consumption. In particular, the second factor is related with the smartphone's location. In Figure 8, we plot the correlation between the location and the energy consumption in idle state. As shown in the figure, for most locations, the variability in power consumption is relatively low (small-sized boxes with low number of outliers). This indicates a very strong correlation between the power consumption and the location when the phone is in idle state.

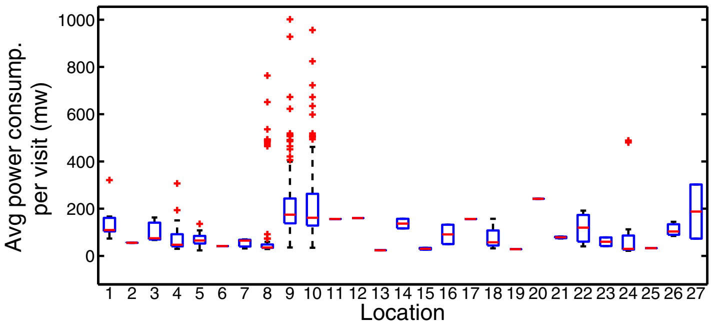

Fig. 8. Location vs Power consumption in Idle state

5.2 Active State

Energy consumption of a smartphone in active state is much higher than the average energy consumption in idle state. In this state, energy consumption depends on the running applications/software and the utilized sensors. The running applications are varied and user-specific. Since location is an important context that affects the running applications on the phone, we explore the correlation between energy consumption in active state and the location.

Figure 9 shows a plot of this correlation. Based on this plot, we can clearly divide user locations into two categories. The first category consists of locations where power consumption is relatively constant (small-sized boxes and low number of outliers). Examples include locations 2-6, 8, 10, and so on; for these locations, the running applications are strongly correlated with the location. The second category consists of locations where power consumption is highly variant (large-sized boxes and/or high number of outliers). Examples include locations 1, 7, 11-13, and so on.

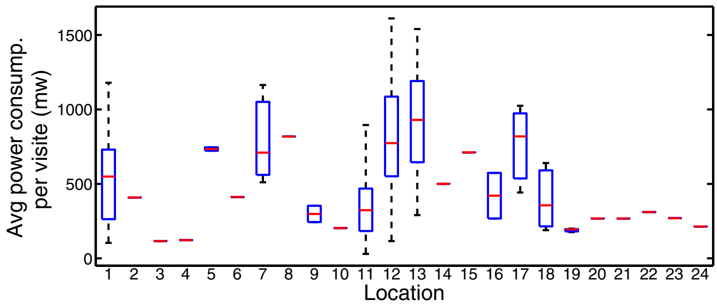

Fig. 9. Location vs Power consumption in Active state

Our analysis of power consumption in active and idle states indicates that a user-centric power consumption model, which incorporates these two states along with time of day, location, remaining battery level, and recent phone usage history, will help to provide more accurate predictions about power consumption.

6 Conclusions

In this paper, we studied and analyzed the interrelationships between smartphone usage behavior, environment, and energy consumption. We have identified four factors — time of day, user location, battery state, and recent phone usage history — that can be used to predict phone usage and energy consumption. We also explored the impact of location on energy consumption under both idle and active phone states. Our study results provide useful insights for modeling smartphone energy consumption, particularly taking into account the context of location.

While we have focused on individual factors and analyzed their impact on energy consumption, it would be interesting to investigate if there exists any stronger correlation between energy consumption and a combination of two or more factors. For example, our results show that there is a high correlation between energy consumption and time of day. Similarly, the results also indicate that there is a high correlation between energy consumption and user's location. We may find an even stronger correlation between energy consumption and <time of day, location>. As part of continuing efforts, we are developing quantitative energy consumption models based on the data collected from users. Then we are going to integrate these models into power management services present in smartphones. The end result will be an intelligent and energy efficient mobile platform based on energy consumption patterns of individual smartphone user.

References

1. Ashbrook, D., Starner, T.: Using GPS to learn significant locations and predict movement across multiple users. Personal Ubiquitous Computing 7(5), 275–286 (2003)
2. Banerjee, N., Rahmati, A., Corner, M.D., Rollins, S., Zhong, L.: Users and batteries: interactions and adaptive energy management in mobile systems. In: UbiComp 2007 (2007)
3. Carroll, A., Heiser, G.: An analysis of power consumption in a smartphone. In: USENIX ATC (2010)
4. Falaki, H., Mahajan, R., Kandula, S., Lymberopoulos, D., Govindan, R., Estrin, D.: Diversity in smartphone usage. In: MobiSys (2010)
5. Hightower, J., Consolvo, S., LaMarca, A., Smith, I.E., Hughes, J.: Learning and recognizing the places we go. In: ACM Ubicomp, pp. 159–176 (2005)
6. Kang, J.H., Welbourne, W., Stewart, B., Borriello, G.: Extracting places from traces of locations. ACM Mobile Computing Communication Review 9(3) (2005)
7. Kim, D.H., Hightower, J., Govindan, R., Estrin, D.: Discovering semantically meaningful places from pervasive rf-beacons. In: ACM Ubicomp (2009)
8. Kim, D.H., Kim, Y., Estrin, D., Srivastava, M.B.: SensLoc: Sensing everyday places and paths using less energy. In: ACM SenSys (2010)
9. Laasonen, K., Raento, M., Toivonen, H.: Adaptive On-Device Location Recognition. In: Ferscha, A., Mattern, F. (eds.) PERVASIVE 2004. LNCS, vol. 3001, pp. 287–304. Springer, Heidelberg (2004)
10. Ma, Y., Hankins, R., Racz, D.: iloc: a framework for incremental location-state acquisition and prediction based on mobile sensors. In: CIKM (2009)
11. Rahmati, A., Qian, A., Zhong, L.: Understanding human-battery interaction on mobile phones. In: MobileHCI (2007)
12. Rahmati, A., Zhong, L.: Fast track article: Human battery interaction on mobile phones. Pervasive Mob. Comput. 5 (October 2009)
13. Simunic, T., Benini, L., Glynn, P., De Micheli, G.: Dynamic power management for portable systems. In: MobiCom (2000)
14. Trestian, I., Ranjan, S., Kuzmanovic, A., Nucci, A.: Measuring serendipity: connecting people, locations and interests in a mobile 3g network. In: IMC 2009 (2009)
15. Yang, G.: Discovering Significant Places from Mobile Phones – A Mass Market Solution. In: Fuller, R., Koutsoukos, X.D. (eds.) MELT 2009. LNCS, vol. 5801, pp. 34–49. Springer, Heidelberg (2009)