Characterizing Mobile User Habits: The Case for Energy Budgeting

Ramanujan K. Sheshadri*, Ioannis Pefkianakis§, Henrik Lundgren†, Dimitrios Koutsonikolas*
Anna-Kaisa Pietilainen‡, Augustin Soule†, Jaideep Chandrashekar†
*University at Buffalo, SUNY, Buffalo, NY
{ramanuja, dimitrio}@buffalo.edu
†Technicolor Research Labs, Paris
{henrik.lundgren, augustin.soule, jaideep.chandrashekar}@technicolor.com
‡Inria Paris-Rocquencourt
anna-kaisa.pietilainen@inria.fr
§HP Labs, Palo Alto, CA
ioannis.pefkianakis@hp.com

***Abstract*—In this paper, we collect and analyze data from 85 smartphone users over a 9 month period. Different from existing work, we study device usage patterns in concert with network performance in space and time. Our results uncover predictable mobility patterns, where users are moving between hubs (i.e., home or workplace) and transit locations. In hubs, users are typically connected using Wi-Fi, while in transit locations cellular connectivity dominates with highly varying performance (from EDGE to HSPA+). Interestingly, there are set of apps over time running on user devices, independent of the location, network conditions, and device resources (e.g., battery level). These apps can aggressively use the network, which leads to significant device resource consumption (e.g., energy), as shown by our controlled experiments. We discuss how our findings can be used to budget mobile device available resources and improve user experience.**

I. INTRODUCTION

The use of mobile devices, such as smartphones and tablets, is steadily increasing and it is predicted that the traffic from wireless and mobile devices will account for 61% of the IP traffic by 2018 [1]. Despite the increasing popularity of mobile devices, mobile user habits, mobile device usage patterns, network performance, and their spatio-temporal relationships, which together determine user experience, still remain largely unexplored. In this paper, we seek to fill this gap by collecting and analyzing traces from 85 smartphone users (volunteers) over a 9 month period. To this end, we implement CellOScope [2], an Android application which periodically collects location and device usage information (running applications, traffic etc.), and instruments network measurements to determine network performance (throughput). This is different from existing studies which mainly collect and analyze either network performance data (e.g., [3]) or device usage patterns (e.g., [4], [5]) but not both.

We first investigate the mobility patterns of the users in our dataset. We classify a user's locations into *hubs* and *transit* locations, according to the frequency of visits and the duration that a user spends at a certain location. A hub is typically their home or workplace, which the user visits often and stays for a long time. Our results show that we can often predict with certain accuracy the times at which a user will be at a hub location. Our network performance analysis further shows that Wi-Fi is typically used at hub locations and provides on average 50.9% higher throughput than cellular.

We then look at the applications used at the two types of locations. Interestingly, our trace shows that smartphone users typically use almost the same set of applications at both hubs and transit locations. Twitter, Facebook, and Email are among the most popular applications used by the users in our dataset. These applications generate significant network activity, attributed to periodic syncs. These periodic syncs can run even in the background irrespective of whether the user is actively using the application or not. Such background network activity leads to significant energy consumption as shown by our controlled experiments. Half an hour of Twitter and Facebook sync operation over a cellular network can reduce a typical smartphone's standby time by 3 hrs and 1.5 hrs respectively. Based on our findings, we propose the design of a user profile-based resource budgeting system in mobile devices, which can defer delay-tolerant applications and/or operations (such as periodic syncs), until more network or device resources (e.g., a better network or a power plug) are available.

The rest of the paper is organized as follows. Section II discusses related work. Section III presents the implementation of CellOScope and an overview of our dataset. Sections IV and V investigate user mobility and network performance, respectively. Section VI studies user applications and their impact on energy consumption. Finally, Section VII discusses a user profile-based resource budgeting design in smartphones and concludes the paper.

II. RELATED WORK

Considerable work has been done in the past towards developing user-specific solutions to contain resource wastage in smartphones. Building user profiles by identifying patterns in user data is shown to be an effective option [6]. To this extent, our work is complementary to many other works [3], [5], [7], [8], [9], [10], [11], [12]. The novel contribution of our work is the study of device usage and network performance of a diverse set of smartphone users in concert with their behavior (mobility). We believe that our approach of combining network, device usage, and user behavior allows us to build a more credible and strong user profile framework that can be used for smartphone resource budgeting schemes.

Researchers in [5], [8] use only the phone usage patterns to build user profiles to estimate rate of battery drain. Such estimates, which do not consider network performance, will often be incorrect. As we show in Section VI, energy drain in smartphones for a given application varies with different networks and hence, involving network performance in user profiles is important.

On the other hand, studies that consider the available network patterns [3], [9], [10], [11] fail to take into account the user phone usage preferences. The potential drawback is that this may cause them to miss out on prioritizing the data transfers or at times even allow background applications to transfer wasteful data that the user rarely needs. This may further aggravate both energy consumption and user experience.

Finally, works which study both the application and the network performance of various smartphone users [4], [12], [13] analyze the problem from the network perspective and not from the user's perspective.

III. Data Collection

Collecting reliable data from smartphone users in real time is a major challenge due to a number of reasons. Users are reluctant to share their data due to privacy and security, as well as energy concerns. In the past, researchers have studied different ways of collecting user data [14]. Conducting surveys/asking participants about their behavior is one of the most popular practices. However, such surveyed data is often unreliable and incorrect [14]. Another popular method is by installing a custom application on user's smartphone. This method provides more flexibility in monitoring the user's phone activities reliably but it comes at a cost. Any custom application which monitors user activities has to run continuously in the background thus depleting the user's phone battery. Additionally any active measurements involving sending/receiving data over cellular networks can eat up a huge chunk of the users data budget. Thus, finding a right balance between data collection and user preferences is always crucial.

For our data collection, we used CellOScope – a smartphone data-collection system. CellOScope has two components: a) an Android based smartphone app which was installed on participants' smartphones manually and was also made available on google play for download and b) a data collection server. CellOScope[1] data collection involves three important components of the user's smartphone data:

- Location: The CellOScope application tracks the user's geographical locations using the GPS coordinates. In case the GPS coordinates are unavailable, it uses the location coordinates provided by the cellular operator.
- Network: CellOScope conducts periodic active/passive monitoring of the user's network activities. Passive monitoring involves tracking the user's network preferences (Wi-Fi/cellular) and recording the available cellular networks. During active monitoring, CellOScope measures the downlink HTTP throughput by sending 100KB of HTTP data from our servers to the user's phone over the network the phone is currently connected to. This helps us record the HTTP throughput performance at different times over different networks.
- Applications: The third component of data collection involves user's applications usage patterns. CellOScope tracks the active/running applications on the user's smartphone. Data counters record the data sent/received in bytes for each application since the last reading was taken for that application. This helps us analyze various patterns in user's smartphone usage and also identify rogue applications involved in unusually high data transfers.

[1]CellOScope users have agreed to all our terms and conditions and due legal process was followed in collecting sensitive information.

TABLE I. **Overview of the dataset.**

Collection dates	02-21-2013 - 11-30-2013
Number of users	85
Location of sampled users	Paris, Fr
Sampling frequency	1 hr or 15 min
Avg. #of samples per user	362
Types of networks sampled	Wi-Fi, HSDPA, HSPA, HSPA+, EDGE, GPRS

A. Challenges and Limitations

The most important challenge in data collection is to encourage the users in sharing their data. To incentivize the users, the application provides useful information about the user's network performance, as well as location and mobility patterns using Google Maps. In order to avoid excessive usage of the phone's battery, we stop sampling during night when the user is most likely to be at home. The sampling frequency of CellOScope was set to 1 hr by default and for users who agreed to higher sampling frequency it was set to 15 min.

Out of the total 85 users in our dataset, 15 users have continuous data for days ranging from 3 weeks to 7 months. For the rest of the users, we have discontinuous data samples ranging from 2 to 100 days and spread over 9 months. The main reason for this is that CellOScope stops sampling when the (monthly) cellular data usage of CellOScope crosses the user set threshold. Table I gives some additional specifics of the dataset we collected over 9 months.

IV. Understanding User's Mobility

In order to predict spatial device usage patterns, we first try to understand user mobility. To identify patterns in user mobility, we classify user locations as: a) hubs and b) transit locations. The intuition behind this classification is to identify locations where a user visits frequently and stays for a long time. For example, a user's home and workplace can be classified as hubs. When the user is at a hub, we assume that she can recharge her phone, and thus would not worry about wasting the phone's battery. On the other hand, transit locations are those which a user rarely visits for a brief duration and may or may not find charging opportunities. These transit locations can be places which the user visits

TABLE II. **Hubs and transit locations for different values of α.**

α	Avg. Hubs	Avg. Transit
0.20	1.07	25.2
0.15	1.32	25.0
0.10	1.61	24.7
0.05	2.22	24.1
0.01	5.76	20.6

during her commute between hubs. When the user is at a transit location, she may want to conserve her phone's resources like battery, and cellular data. To classify a location as a hub or a transit location, we consider two parameters: *Frequency* and *Duration*. They are calculated as follows:

$$Frequency = \frac{No.\ of\ days\ user\ visits\ a\ location}{Total\ no.\ of\ sampling\ days}$$

$$Duration = \frac{Hours\ user\ spends\ at\ a\ location\ each\ day}{Total\ sampling\ hours\ per\ day}$$

The *frequency* represents the user regularity in visiting a particular location. The *duration* represents the fraction of time spent by a user at a particular location each day. The parameter *duration* is calculated per day and *Avg. duration* gives the average time spent by the user at a location each day. We then calculate the *Location Popularity (LP)* for every location visited by a user as:

$$LP = Frequency\ *\ Avg.Duration$$

We categorize a location as a hub if $LP \geq \alpha$, and as a transit location if $LP < \alpha$ where α is an adjustable parameter. Table II shows the average number of hubs and transit locations for all users. For the remaining of the paper, we select $\alpha = 0.05$, since we consider that a typical user is most likely to have two hubs (home and workplace).

To investigate the possibility of predicting a user to be at a hub or a transit location at any given time of the day, we consider the 15 users for whom we have continuous data samples for over 3 weeks. The continuous data for these users help us track user mobility with greater certainty. For our analysis, we separately consider weekdays and weekends in order to investigate if the user mobility patterns change during these two time periods. Figures 1 and 2 show the probability heatmap during weekdays and weekends, respectively. The grey scale indicates the probability that the user is at a hub at a particular hour. A fully white block indicates a probability of 1 for a user to be at a hub and a fully black block indicates a probability of 0 (i.e., the user is at a transit location).

The heatmaps show diversity among columns, which suggests that individual users exhibit different behavior. This suggests that adaptation of services given a user's spatio-temporal profile can benefit from personalization. Very light colors (white/light-grey) and very dark colors (black/dark-grey) indicate that we can predict whether or not a user is at a hub with high certainty. We observe that several users have monotonous mobility habits, as there are clearly such patterns at certain times of the day. For example, user 11 is almost always at a hub in the mornings, while user 6 is almost always in transit during weekday morning hours. On the other hand, for some users the probability is somewhere around 0.5 for certain time periods. This makes it challenging to predict the user's location with greater probability. However, for most of the users we see that it is possible to predict their location at a given hour with greater certainty with either probability > 0.8 or probability < 0.3. By comparing the weekday and weekend heat maps, we can see that the pattern changes for many users. While user 11 reinforces the probability of staying at a hub all day long, we have less certainty about the location of user 6 as compared to weekdays. Overall, it seems slightly easier to predict user's presence at hubs during the weekends. This may suggest that user mobility is more predictable during weekends than during weekdays.

V. NETWORK PERFORMANCE

In this section we analyze a) network preference of users at hub and transit locations and b) performance of Wi-Fi and cellular networks.

Figure 3 shows the average usage time of a particular network when users are at hub and transit locations. We see that when users are at a hub, Wi-Fi is the most preferred network option. On average, users connect to Wi-Fi for over 75% of the duration they are at hubs and use cellular networks for the remaining duration. On the other hand, when users are at transit locations they tend to use the cellular network most of the time. The cellular connectivity at transit locations can vary from EDGE, GPRS to 3.5G (HSPA+)[2]. Interestingly, users connect to low speed cellular networks like EDGE, and UMTS for little more than 20% of the duration they spend at transit locations.

Next, we investigate the throughput performance of Wi-Fi and cellular networks. Figure 4 shows the results of our HTTP throughput measurements. We show the throughput recorded for downloading 1000KB over Wi-Fi and cellular network at hub and transit locations. Clearly, Wi-Fi provides better HTTP downlink throughput than the cellular network at both hubs and transit locations. In the median case, the Wi-Fi HTTP throughput at hubs and transit locations is 4.46Mbps and 2.7Mbps, respectively, and the cellular HTTP throughput is 2.21Mbps and 0.8Mbps, respectively. Wi-Fi at hubs offers much better throughputs when compared to the throughputs of cellular networks. Figure 5 further breaks down the HTTP throughput offered by each type of cellular network. The Wi-Fi HTTP throughput is included for comparison. In the median case, even the high speed cellular networks – HSPA and HSDPA – offer throughput 88% and 62% lower than the throughput offered by Wi-Fi.

VI. APPLICATIONS

In this section, we investigate the user's application usage patterns at hubs and transit locations. In order to understand what applications are popular at hubs and transit locations, we introduce a simple metric called *Application Popularity (AP)*. The AP is calculated as:

$$\frac{\#\ Samples\ of\ an\ app\ running\ at\ hub\ or\ transit}{Total\ no.\ of\ samples\ at\ hub\ or\ transit}$$

Table III lists some of the most popular applications at hubs and transit locations among all users along with their popularity. We notice that gaming and music applications have a relatively high popularity at transit locations as compared to hubs.

[2]At the time of data collection none of the operators in Paris offered LTE connectivity.

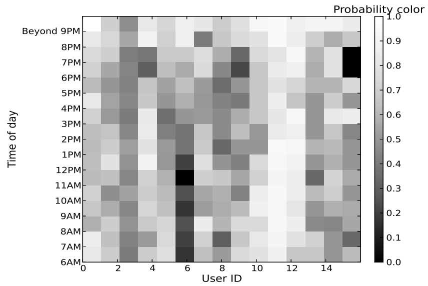

Fig. 1. Probability heatmap for users in hubs during weekdays.

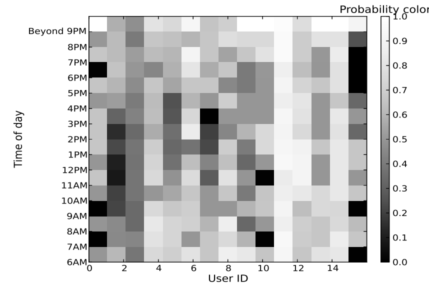

Fig. 2. Probability heatmap for users in hubs during weekends.

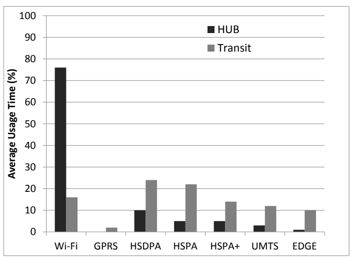

Fig. 3. Avg. network usage at hubs and transit locations.

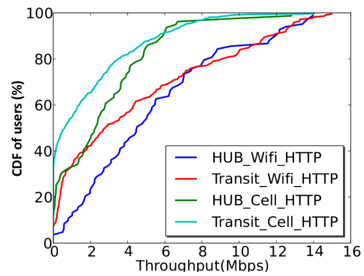

Fig. 4. HTTP throughput at hubs and transit locations.

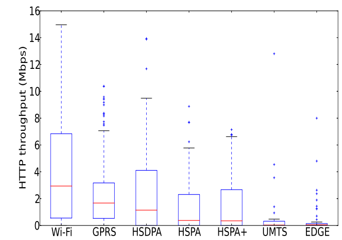

Fig. 5. HTTP throughput in Wi-Fi and cellular networks.

TABLE III. POPULAR APPS IN HUBS/TRANSIT.

App type	App name	App Popularity in Hubs	App Popularity in Transit
Network	Email	0.42	0.48
	Browser	0.54	0.56
Social Network	Twitter	0.71	0.76
	Facebook	0.42	0.45
Location	Google Maps	0.22	0.29
Games	Sudoku	0.40	0.77
	Candy Crush	0.10	0.33
System	Music Player	0.18	0.52

In contrast, other popular applications appear location agnostic exhibiting similar popularity at both hubs and transit locations. The reason for such a behavior is that these applications require constant synchronization with their servers to receive periodic updates and hence, they run in the background irrespective of whether the user is actually using the application or not. Similar observations were made in [15] for social networking applications like Facebook.

The sync applications, which periodically schedule data exchanges irrespective of the available network conditions, can have serious implications on the phone's resources. In order to quantify the impact of periodic data exchange by sync applications on the phone's resources, we first examine the bytes downloaded by the most popular sync applications across all users in our dataset. The graphs in Figures 6-9 show the data (per hour) exchanged by social networking, maps, email, and browser applications for every user. We see that the social networking applications are the most aggressive in exchanging data. For the median user, the social networking applications exchange 145KB of data on average per hour while maps, email, and browser applications exchange only 16KB, 16KB, and 32KB of data, respectively.

The next important question is to understand the impact of these periodic data transfers on the phone's battery life. To answer this question, we conducted a few controlled experiments on a case study user.

A. *Controlled Experiments*

In order to understand the impact of periodic data transfers which usually run in the background by different sync-enabled applications on the energy consumption of a smartphone, we selected a case study user from our dataset and tried to emulate the periodic data syncs under various network conditions that are similar to the network conditions which the case study user experienced at hubs and transit locations. The idea behind this experiment is to obtain an estimate on how different network conditions impact the energy drain in smartphones during periodic syncs.

For this experiment, we selected the 3 most popular applications (Twitter, Facebook, and Google Maps) which require periodic syncs and usually run in the background irrespective of whether the user is likely to use them or not. As we show in our experiments below, these applications are ignorant of the available network conditions and schedule data exchange periodically.

1) Experiment Methodology: We used the Samsung Galaxy S4 smartphone to conduct this experiment. We ran the 3 applications serially in the background with the screen off. We measured the energy using the Monsoon power monitor [16]. We set the application sync settings similar to what the case study user used during data collection. For Facebook, the periodic sync duration was set to 30 min and for Twitter it was set to 5 min. For Google Maps, users do not have an option to set the sync-duration. We allowed each application to run in the background for 30 min so as to ensure that all three applications had a chance to sync at least once during

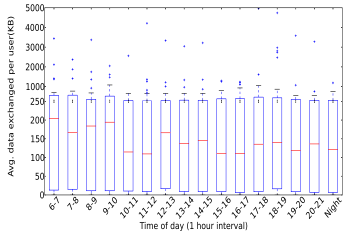

Fig. 6. Data exchanged by social networking apps.

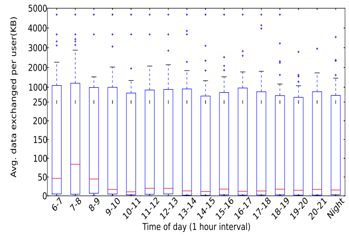

Fig. 7. Data exchanged by Google maps.

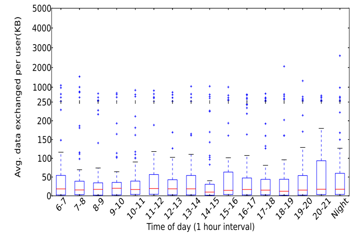

Fig. 8. Data exchanged by email apps.

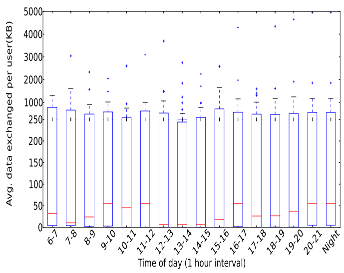

Fig. 9. Data exchanged by browser apps.

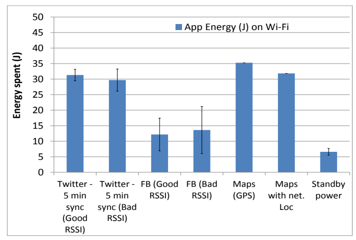

Fig. 10. Wi-Fi energy consumption for case study user.

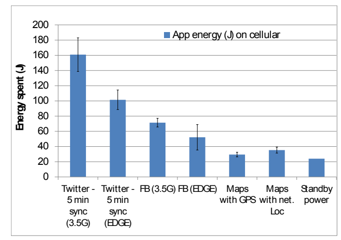

Fig. 11. Cellular energy consumption for case study user.

the experiment period. Figures 10 and 11 show the mean and standard deviation of Wi-Fi and cellular energy measured for the three applications over 5 trials. The Wi-Fi network setting emulates the condition when the case study user is at a hub and the cellular network setting emulates the conditions when case study user is at a transit location.

2) Analysis: We first analyze the energy consumption of the background syncs of the social networking applications – Twitter and Facebook with HSPA+ and EDGE cellular networks. This is the energy the case-study user is most likely to incur at transit locations. As shown in 11 when the user sets 5 min sync period, the mean energy consumed in 30 min by Twitter was 160J over HSPA+ network and 100J over EDGE network. The Facebook sync operation is not as aggressive as Twitter. Yet, the energy consumed by Facebook in 30 min is still high, 70J and 54J over HSPA+ and EDGE respectively. In contrast, the energy incurred for the same duration when the phone was in standby mode (connected to the cellular network, but with no active sync) is just 24J. This means that 30 min of Twitter sync operation over HSPA+ and EDGE networks would reduce the phone's standby time by 3 hrs and 2 hrs respectively, while a single Facebook sync operation would bring down the phone's standby time by 1.5 hrs and 1 hr respectively.

Next we analyze the energy consumption for Twitter and Facebook sync operations over Wi-Fi. To emulate extremes in Wi-Fi signal quality which the user is most likely to witness at hubs, we conduct the experiment with high (-45dBm) and low signal strengths (-85dBm). The total energy cost incurred for the Twitter sync was 82% and 72% lower than the energy cost over HSPA+ and EDGE networks, respectively. Similarly, the total energy cost for the Facebook sync over Wi-Fi was 78% and 70% lower than the energy cost over HSPA+ and EDGE networks, respectively. From the user's perspective, enabling sync operations during her stay at a hub over Wi-Fi is more energy efficient than running these syncs at transit locations over a cellular network. Surprisingly, we do not see significant difference in the energy consumption for the Twitter and Facebook sync operations under good and poor Wi-Fi conditions. This observation is in contrast to the claims made in [17] where authors show high energy consumption for Wi-Fi transfers during low signal strength. This is because the idle energy dominates the energy consumed during the actual sync and hence there is little difference in the total energy consumed during good and poor Wi-Fi conditions.

We now look at location applications like Google Maps. For location applications, users can opt between GPS based location and the default option of location provided by the network provider. Our experiment reveals that for either of the options Google Maps consumes roughly the same energy. This makes us believe that the energy incurred is actually to run the application. In fact, the energy consumed for running Google Maps for 30 min is higher than the energy consumed by running Facebook sync operation over Wi-Fi for the same duration.

We now briefly analyze the possible reasons for the difference in energy consumption for application sync-operations under HSPA+, EDGE, and Wi-Fi networks. The low energy consumption by Wi-Fi as compared to the cellular networks is due to the difference in tail energy. The graphs in figures 12, 13, and 14 show the duration of tail for a single Twitter sync operation over HSPA+, EDGE, and Wi-Fi. We observe a prolonged tail of over 10 sec for HSPA+ and over 3 sec for EDGE; in contrast, there is no tail for Wi-Fi. The longer the tail the higher is the energy consumption. This is in accordance with the previously observed results in [18].

The important take away point is that many popular applications exhibit significant background activity and exchange data periodically. This can lead to significant energy consumption which varies with the network type and is almost independent of the signal strength. Furthermore, background

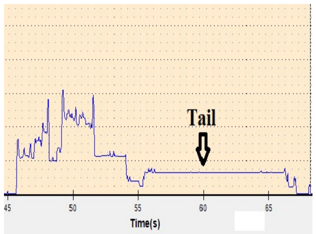

Fig. 12. Tail duration in HSPA+ network.

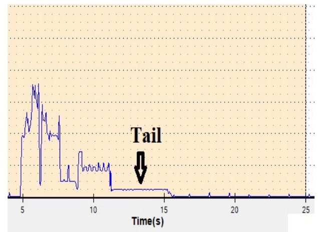

Fig. 13. Tail duration in EDGE network.

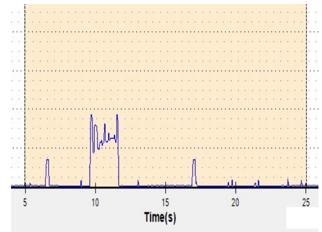

Fig. 14. Tail duration in Wi-Fi network.

syncs run irrespective of the available resources like battery and network type.

VII. Discussion and Conclusion

In this paper, we collected and analyzed user behavior (mobility), network performance, and device usage data from a set of 85 smartphone users. Our study uncovers significant resource (energy) consumption from background applications and sync operations, which run independently of the available mobile device resources and the network conditions. Our results motivate the design of user-profile based resource budgeting in mobile devices.

A profile-based resource budgeting design includes three key building blocks: *a)* A *mobility tracker* which can predict when a user will be located at hubs or transit locations. Our results show that this prediction is possible with a certain confidence. *b)* A *network and device monitor* which can profile the connected network types and the available device resources (e.g., battery) at hubs and transit locations. *c)* An *application monitor* which profiles the running apps and feeds a resource budget module which can defer execution of power hungry/delay-tolerant applications based on the available device resources. We have implemented the mobility tracker, the network and device monitor, and the application monitor in CellOScope (Section III). To construct the user profile in an energy efficient manner, it is important to collect user feedback at coarse time scales and avoid using any power hungry sensors. Apart from the above modules, the resource budgeting design should be able to predict the resource (energy) consumption by applications until the user reaches a hub.

A similar solution was proposed by [8] for resource budgeting in smartphones. The authors used user profiles to maximize the telephony talk time by postponing sync activities of different applications. They used past call records as input to make decisions. However, with advancements in smartphones such decision process may not be that simple, as there can be many other dependencies. For example, many VOIP/video applications which are popularly used as an alternative to telephony services depend on the available network. Additionally, many of these applications use inherent sensors that may cause further energy drain. In order to have an effective decision process, it is important to understand the available resources and their usage in greater depth. We leave the resource budgeting design and implementation as a future work. Overall, we consider our work to be an important step towards understanding the mobile user, and designing user profile-based resource budgeting on mobile devices, which can enhance user experience.

Acknowledgment

This work was supported in part by the European ICT FP7 User Centric Networking project (grant no. 611001) and NSF grant CNS-1422304.

References

[1] "Cisco visual networking index: Forecast and methodology," http://www.cisco.com.

[2] "Google playstore," https://play.google.com/store/apps/details?id=com.technicolor.celloscope&hl=en.

[3] A. Balasubramanian, R. Mahajan, and A. Venkataramani, "Augmenting mobile 3g using wifi," in *Proc. of ACM MobiSys*, 2010.

[4] H. Falaki, D. Lymberopoulos, R. Mahajan, S. Kandula, and D. Estrin, "A first look at traffic on smartphones," in *Proc. of ACM IMC*, 2010.

[5] D. Ferreira, A. K. Dey, and V. Kostakos, "Understanding human-smartphone concerns: A study of battery life," Technical report.

[6] N. Ravi, james Scott, L. han, and L. Iftode, "Context-aware battery management for mobile phones," in *Proc. of IEEE Percom*, 2008.

[7] N. Vallina-Rodriguez, P. Hui, J. Crowcroft, and A. Rice, "Exhausting battery statistics - understanding the energy demands on mobile handsets," in *Proc. of ACM MobiHeld*, 2010.

[8] E. Jung, F. Maker, T. L. Cheung, X. Liu, and V. Akella, "Markov decision process (mdp) framework for software power optimization using call profiles on mobile phones," Technical report, 2010.

[9] N. Santhapuri, J. Manweiler, R. R. Choudhury, and S. Nelakuditi, "Bytestogo: Offloading 3g via wifi prioritization," Technical report, 2010.

[10] J. Manweiler, N. Santhapuriy, R. R. Choudhury, and S. Nelakuditi, "Predicting length of stay at wifi hotspots," in *Proc. of IEEE INFOCOM*, 2013.

[11] V. A. Siris and D. Kalyvas, "Enhancing mobile data offloading with mobility prediction and prefetching," in *Proc. of ACM MobiArch*, 2012.

[12] Q. Xu, J. Erman, A. Gerber, Z. M. Mao, J. Pang, and S. Venkataraman, "Identifying diverse usage behaviors of smartphone apps," in *Proc. of ACM IMC*, 2011.

[13] J. Huang, F. Qian, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck, "A close examination of performance and power characteristics of 4g lte networks," in *Proc. of ACM Mobisys*, 2012.

[14] F. B. Abdesslem, I. Parris, and T. Henderson, "Reliable online social network data collection," Technical report.

[15] F. Qian, Z. Wang, Y. Gao, J. Huang, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck, "Periodic transfers in mobile applications: Network-wide origin and impact and and optimization," in *Proc. of WWW*, 2012.

[16] "Monsoon power monitor." http://www.msoon.com/LabEquipment/PowerMonitor/.

[17] N. Ding, D. Wagner, X. Chen, A. Pathak, Y. Hu, and A. Rice, "Characterizing and modeling the impact of wireless signal strength on smartphone battery drain," in *Proc. of ACM Sigmetrics*, 2013.

[18] N. Balasubramanian, A. Balasubramanian, and A. Venkataramani., "Energy consumption in mobile phones: A measurement study and implications for network applications," in *Proc. of ACM IMC*, 2009.