

Casting Doubts on the Viability of WiFi Offloading

Shu Liu, Aaron D. Striegel
Department of Computer Science and Engineering, University of Notre Dame
Notre Dame, IN, USA
sliu6@nd.edu, striegel@nd.edu

ABSTRACT

With the advent of the smartphone, mobile data usage has exploded which in turn has created tremendous pressure on cellular data networks. A promising candidate to reduce the impact of cellular data growth is WiFi offloading. However, recent data from our study of two hundred student smartphone users casts doubts on the reductions that can be gained from WiFi offloading. Despite the users operating in a dense university WiFi environment, cellular consumption still dominated overall data usage. We believe the root cause of lesser WiFi utilization can be traced to the WiFi being optimized for laptop WiFi reception rather than the more constrained smartphone WiFi reception. Our work examines the relationship of WiFi versus 3G usage through a variety of aspects including active phone usage, application types, and traffic volume over an eight week period from the Spring of 2012.

Categories and Subject Descriptors

C.4 [PERFORMANCE OF SYSTEMS]: Reliability, availability, and serviceability; C.5.3 [COMPUTER SYSTEM IMPLEMENTATION]: Microcomputers—*Portable devices*

General Terms

Measurement

Keywords

WiFi offloading, 3G, User behavior

1. INTRODUCTION

Wireless connectivity is undergoing a tremendous transformation as the bandwidth demands afforded by smartphones and other wireless devices have simply overwhelmed the data speeds and access paradigms offered by various third-generation (3G) technologies. Moreover, the predicted

wireless data tsunami encompassing 1000x growth over the next ten years for wireless shows all the early signs of being a conservative estimate [1]. Although technologies such as LTE herald the arrival of fourth-generation (4G) wireless technology, the new speeds often only temporarily satiate the need for additional bandwidth.

For many cellular providers, WiFi offloading, i.e. the users receiving data from 802.11-based hotspots, offers a significant appeal by reducing the strain on the already overloaded cellular infrastructure. Recent studies such as the one in [2] points to offloading offering gains approaching 65% of the total traffic volume. There are other works such as [3–6] that discuss the feasibility of WiFi offloading. However, our on-going large scale study of several hundred student smartphone users casts doubts on the viability of achieving such gains with WiFi offloading. In contrast to the prior work of [2], we have found that despite users operating in a dense university WiFi environment, the potential gains for the majority of users are quite muted. Rather than finding that WiFi usage dominates 3G usage, our study curiously finds that much of the user consumption of 3G dominates the consumption of WiFi.

At first glance, such a result would appear to counterintuitive. A dense WiFi deployment¹ would have ample WiFi access points placed throughout the buildings on the campus. The dormitory-oriented residential life of the campus meant that all study participants (all of whom were freshmen) would have ubiquitous WiFi coverage at night (dormitory) as well as during the day (classrooms/dorms). Coverage at the university is also frequently verified by employing roaming laptops by IT staff throughout campus. However, it is the usage of laptops to validate coverage as opposed to smartphones where the problem originates.

Consider the observation noted in Table 1 that compares the observed RSSI on a laptop versus the RSSI observed in the same time period and same location via a representative smartphone from the study. In the table, despite being in the same location, the smartphone observes a dramatically reduced signal strength (around 10dB) versus the laptop. The discrepancy between observed signal strength on the smartphone is not an isolated phenomenon to one particular smartphone but rather tends to be broadly indicative of many commercially available smartphones that are subject to marketing and development constraints. This is potentially significant for WiFi offloading as most dense WiFi

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.

CellNet'12, August 13, 2012, Helsinki, Finland.

Copyright 2012 ACM 978-1-4503-1475-6/12/08 ...\$15.00.

¹The University of Notre Dame was ranked as one of the top 20 wireless campuses in 2010.

environments (i.e. businesses) tend not to be designed at smartphones but rather tend to be designed at laptops.

Table 1: WiFi RSSI on Laptop and Smartphone

Access Point	RSSI on Laptop		RSSI on Smartphone	
	Avg.	Std Dev	Avg.	Std Dev
1	-75.27	2.76	-83.32	2.72
2	-75.22	2.79	-85.04	3.39
3	-75.40	2.87	-84.94	3.04
4	-71.99	2.06	-79.52	3.51
5	-72.05	2.06	-80.18	3.80
6	-73.88	5.13	-80.02	3.34

The net result of our preliminary findings over an eight week period of our smartphone usage data shows that the gaps in WiFi coverage for smartphones temper the perceived potential gains by WiFi offloading. To that end, our paper makes the following contributions:

- Our paper casts doubts on the large gains that may be realized in practice through WiFi offloading pointing to a more reasonable gain of 30%. While 30% is certainly beneficial, it represents a significant reduction versus gains predicted in prior work.
- We explore the relationship of WiFi data consumption and phone usage time deduced by screen session length. Moreover, our work includes accommodations for considering SMS/phone call/browser impacts with regards to phone usage.
- Our paper notes that users who are able to get reasonable WiFi reception tend to be users consuming significantly more traffic. Different from users with less WiFi consumption, their video/audio streaming traffic takes a large part of the total data usage.

2. DATA SOURCE

The dataset employed for the analysis in this paper draws from a large scale study NetSense launched in August of 2011. The premise of the study was to explore how always-on/pervasive networking impacted how younger individuals (students) made and kept friendship ties. A secondary aspect of the study was to explore the relationship between the social network and the underlying technical aspects of data consumption. In short, the study provided a smartphone (Nexus S) with unlimited data, unlimited texting, and unlimited mobile-to-mobile minutes plans for free in exchange for complete monitoring privileges (including length and parties involved but excluding content).

The monitoring portion of the phone was developed through a custom application/agent running on top of Android 2.3.4 (Gingerbread). The monitor itself gathers data with respect to all inbound/outbound communications (e-mail, SMS, phone), browser history, screen usage time (time on, time off), hourly traffic counts (cell, WiFi), and other assorted usage characteristics [7]. Data is temporarily spooled onto the SD card of the phone into a SQLite and then periodically relayed back to a central server which then parses/stores the information into a secure, central database.

In August of 2011, the phones were handed out one week prior to the fall semester to the incoming freshmen students




Figure 1: Average Downlink Traffic per Phone per Week

in the study and data collection has been carried since that time. The period in question analyzed in the paper represents eight weeks of continuous data from the end of January 2012 to the end of March of 2012 in the spring semester. The time frame is selected to ensure that the students have fully adapted to their phones and represent behaviors typical of their normal phone usage. The seventh week in the study represents spring break (March 10th - March 18th) which offers a likely period where students returned home for the week.

In total, we selected 131 users out of the study (62 female, 69 male) that have continuous data throughout the entire eight weeks and have more than five minutes per day of usage on the phone². The graphs and tables are based on the data collected from these 131 participants.

Figure 1 shows the average 3G and WiFi downlink traffic per phone per week across the study period. On average, the WiFi traffic takes approximately 30% of the total data consumption. The average WiFi traffic statistics excludes the data of devices with no WiFi usage. The underlying 802.1X WiFi infrastructure (primarily due to expired passwords) can prevent a user from using any WiFi traffic during that entire week. Once the password is updated (passwords are required to change every 90 days), the phone can correctly use WiFi once again. We note though that an inability to authenticate does not preclude monitoring signal strength (via beacons) and that users do not experience intermittent 802.1X issues (the phones either authenticate or do not contiguously).

As noted in the introduction, the average WiFi usage is curiously lower than the average 3G usage aside from the week of spring break (likely on a single AP at home) where WiFi barely eclipses 3G usage. In Figure 2(a), we select a small subset of phones that exhibit a high degree of 3G usage versus WiFi and plot their respective average traffic patterns for a given day during the semester. From 5 PM to 6 PM

²The primary reason for exclusion involved cases where a phone required repair for greater than a single day. In other cases, while the study requested that all students do not use a prior phone, the five minute threshold ensures that the smartphone was their primary phone.




Figure 2: Switch between 3G and WiFi

(1700-1800h), the phones exhibit a degree of asymmetry, favoring 3G over WiFi. Figure 2(b) explores this time period further noting the detected AP beacons but exceptionally low quality signal strengths over that time period. Hence, the phone naturally defaults to favor 3G over WiFi. Conversely, during the 10 PM to 11PM (2200-2300h) as noted in Figure 2(c) where the signal strength is better but not great, the ratio between 3G and WiFi changes to favor WiFi over 3G but not markedly so.

To verify that the users in this subset were indeed active (using the screen), Figure 3 plots the average screen session lengths over that same selected day. The screen session time captures via event when the screen is turned on (timestamp recorded) and when the screen is turned off (timestamp recorded). While the users were more active in terms of the number of sessions during the later time period (10PM to midnight), the users actually used more data during the earlier time period. We also note that the spike did not correlate with a check in by the agent as such data has been filtered from the results.




Figure 3: Screen On Duration in One Day

3. COMPARISON OF USER BEHAVIOR

We continue our analysis by exploring how user behavior changes if the user is able to get reasonable quantities of adequate WiFi smartphone coverage. Figure 4 shows the distribution of the percentage of WiFi downlink traffic as

a ratio versus total traffic for one week (2nd week). The average WiFi consumption ratio is around 30%. From the figure, we see that there are approximately 30% of participants whose traffic is offloaded to WiFi for the week by more than 50%. Conversely, we see that there are nearly 20% of the participants that week who used no WiFi traffic for the reasons mentioned earlier (incorrect password, etc.).




Figure 4: Distribution of Percent of WiFi Traffic

For further analysis in the paper, we subdivide the 131 participants into three groups based on their WiFi downlink consumption ratio versus their total traffic consumption: no WiFi, less than 50%, and more than 50%. We summarize how the participants break down into each group among the eight weeks in Figure 5. Similarly, the breakdown of males versus females is shown in Table 2. Although the numbers in the groups varied, the groups are reasonably well separated to analyze the participants in the various categories.

Figure 6 shows the average total downlink traffic per phone (3G+WiFi) of the different groups across the eight weeks. For the category of each particular participant, determination was computed on a weekly basis meaning that a user could move between categories. Interestingly enough, we begin to see a growing pattern of separation as the semester goes on and particularly so during the week of spring break. We posit that the week of spring break offered a much more consistent set of WiFi coverage versus what may occur on campus.

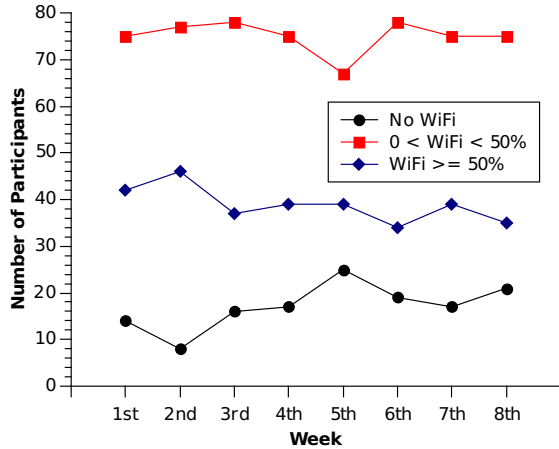


Figure 5: Numbers of Participant in Groups

Table 2: Number of Males and Females in Groups

Week	No WiFi		0<WiFi<50%		WiFi>=50%	
	M	F	M	F	M	F
1st	9	5	35	40	25	17
2nd	6	2	35	42	29	17
3rd	5	11	37	41	27	10
4th	7	10	36	39	26	13
5th	15	10	30	37	24	15
6th	9	10	39	39	21	13
7th	8	9	36	39	25	14
8th	10	11	36	39	23	12

3.1 Phone Usage

Related to usage, a secondary question arises with regards to the ratio of screen time to data consumption. The most straightforward way to indicate whether the person is using the phone is based on the phone screen status (on or off). When the phone screen is on, it is most likely because the person is actively using the phone (sending messages, checking e-mail, playing games, etc.). In order to explore the relationship between phone usage time and traffic usage, we collect the screen on session length on the phones by calculating the duration between screen on and screen off. On average, the screen session lengths vary from 10 seconds to more than 100 seconds. We categorize them into four ranges: (0, 30s), (30s, 60s), (60s, 90s) and more than 90s. Figure 8 lists the numbers of participants from three groups in different duration ranges. In the range of (60s, 90s) the number of people who use WiFi more than 50% is greater than the other two groups numbers. We calculate the weekly total screen on time and the corresponding daily average per phone as well. There are five ranges of daily average from less than half an hour per day to more than 2 hours per day. Figure 7 gives an example (data from the 2nd week): compared with purely 3G users (no WiFi), the number of participants who use WiFi more than 50% in each range is much greater. The fact that the phone is more useful (faster access) may encourage such additional usage.

We further calculate the screen session duration across different time slots in a single day to analyze user behavior. As shown in Figure 9, we divide one day into four time slots:

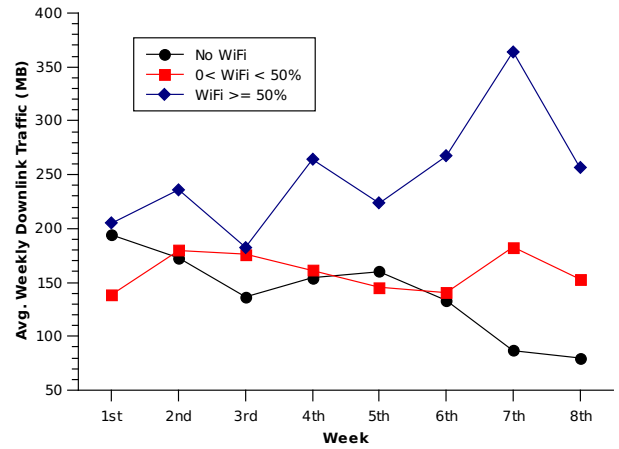


Figure 6: Weekly Average Total Downlink Traffic per Phone




Figure 7: Avg. Daily Screen Usage of Groups

morning (7am-12pm), afternoon (12pm-5pm), night (5pm-10pm) and midnight (10pm-7am). As would be expected with a student population, usage in the morning is quite low relative to usage in the evening.

Finally, we explore the relationship of WiFi to 3G usage with respect to various other services on the phone including the number of text messages, total phone usage, average phone call length, etc. Table 3 shows the numbers across the entirety of the study with users categorized again on a weekly basis. While WiFi dominant users tended to use their screen for longer average periods of time, they tended to less frequently use text messages relative to the non-WiFi-dominant users (455 average text messages sent/received per week versus 522 average text messages sent/received per week). Text messages do not count against the data count for either 3G or WiFi usage.

3.2 App Usage

The network bandwidth and throughput has a profound influence on application usage preference. When the network provides better service, users tend to use more intensive services. We look into the data of installed application

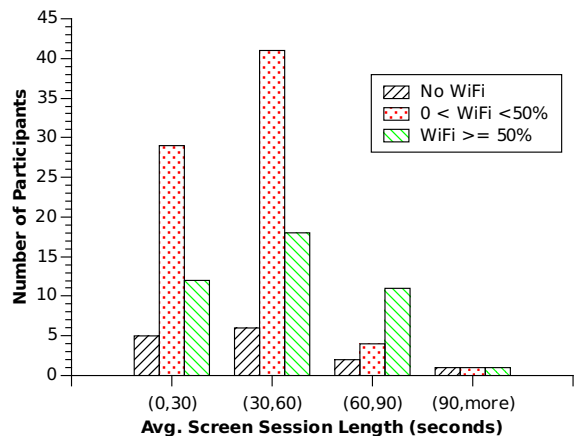


Figure 8: Avg. Screen session length of Groups

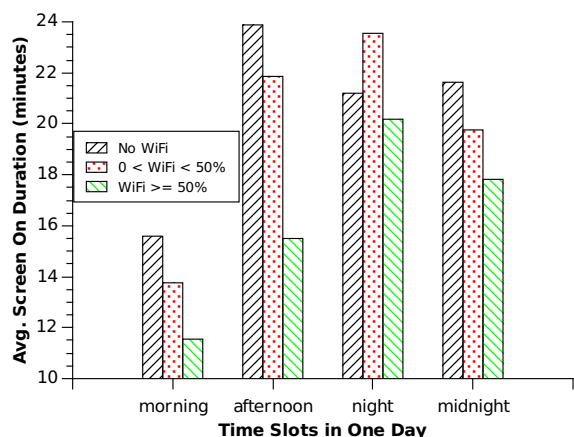


Figure 9: Avg. Screen on Duration in Different Time Slots

and application traffic which are logged by the service as part of the agent (hourly application data usage).

In Table 4 we summarize the top 5 applications appeared in the each group based on their total downlink or uplink traffic. For the phones using WiFi more than 50%, streaming apps such as Netflix and Pandora always are the top applications. For the participants using 3G only, their main activities on the phone are web browsing and email checking. Figure 10 demonstrates the weekly downlink traffic of the top 10 apps in different categories which include video & audio (Netflix, Youtube, Pandora and Google Music), social (Facebook and Twitter), tools (Browser, Gmail, Google Maps, Dictionary) and games (Zynga and Glu applications). Similarly, we present the results of uplink traffic in Figure 11. In both graphs, the traffic of video and audio applications increases dramatically when the participants use more WiFi. An interesting question emerges if the streaming consumption is user specific or rather becomes enabled by WiFi speeds implying that WiFi with continuous LTE would show the same usage patterns.

Table 3: Weekly SMS/Phone Call Comparison

Groups	No WiFi	0 < WiFi < 50%	WiFi >= 50%
Avg. Screen On Duration (seconds)	39.33	43.93	48.76
Avg. SMS All	436	522	455
Avg. SMS Sent	209	249	225
Avg. Number of Phone Calls	26	31	24
Avg. Phone Call Duration (hours)	1.03	1.60	1.27
Avg. Email All	89	95	90
Avg. Number of Browser Sessions	102	81	67

Table 4: Top Applications Categories

	No WiFi	0 < WiFi < 50%	WiFi >= 50%
Top Downlink	Browser Facebook Zynga Words Amazon Twitter	Browser Facebook Pandora Zynga Words Twitter	Netflix Browser Pandora Facebook Dictionary
Top uplink	Google Maps PhoneMonitor Browser Facebook Gmail	Google Maps Pandora PhoneMonitor Browser Facebook	Pandora Google Maps Glu Games PhoneMonitor Browser

4. RELATED WORK

There have been several recent studies proposed to reducing the strain on the overloaded cellular infrastructure with WiFi offloading being one of the main candidate [2–6]. In [3], WiFi connectivity is used to reduce the pressure on 3G spectrum when possible for transferring data. In [2, 4–6], delayed WiFi offloading is introduced to migrate data traffic from cellular networks to WiFi access points. However, such methods are not practical for most access patterns (any interactive app including web browsing and most streaming) as the point of using the smartphone is for data access that moment. Conversely, the de facto solution is upgrading the network to the next generation networks. The configuration details of both WiMAX and LTE technologies are summarized in [8]. Similarly, [9] provided an overview of the 4G

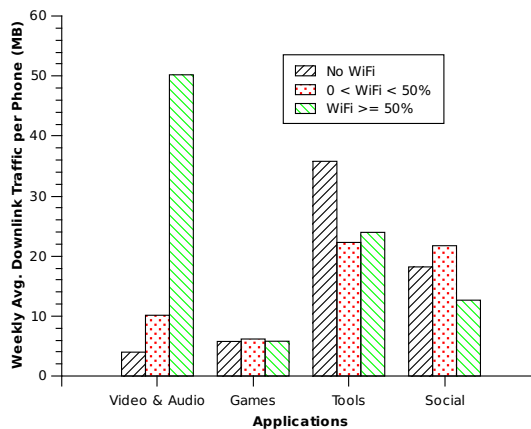


Figure 10: Weekly Application Downlink Traffic

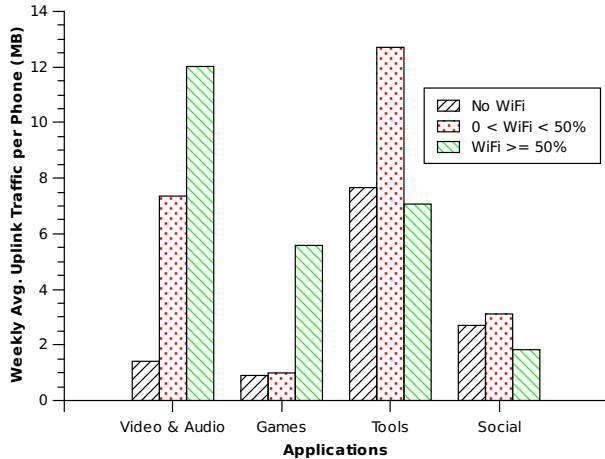


Figure 11: Weekly Application Uplink Traffic

evolution and categorized how such technologies can accompany a more user focused world of wireless.

Finally, there are several work related to the general analysis of user behavior on the mobile devices. Pentland, et al. [10] use Bluetooth signals as indicators for people nearby within the Bluetooth range (around 10m). Youngs, et al. [11] reported the communicative practices of 19 new mobile phone users and further extend the discussion to how mobile phones affects public perceptions. The call duration distribution of phone users in large mobile networks is analyzed in [12] and several distributions are proposed to model such specific user behavior. Dinda et al. [13] present some analysis of smartphone usage based on the data of 25 users in six months. Their findings include power consumption, network activity and etc. Different from their work, this paper focuses more on the change of user behavior between different networks instead of the general data collection.

5. CONCLUSION

In conclusion, we believe that additional scrutiny is needed with respect to the end benefits that can arise from WiFi offloading. Although individual handsets (ex. iPhone) may possess better WiFi characteristics, the heterogeneity of available handsets implies that it is exceptionally likely there will significant variations in the ability of phones to take advantage of WiFi offloading. We believe estimates approaching one third *or less* are much more realistic in terms of the actual impact of WiFi offloading. While certainly any offloading is desperately appreciated due to capacity shortages, WiFi offloading may be further clouded by bottlenecks in the next hop following the WiFi link as well. Furthermore, we believe significant attention should be paid to WiFi coverage to examine said coverage through the lens of a typical smartphone rather than a typical laptop or tablet device. Our future work includes exploring the viability of relaying or multi-pathing as a viable intermediate approach.

Acknowledgement

This work was funded in part by the National Science Foundation through grant IIS-0968529. We would also like to thank our collaborators, Prof. Christian Poellbauer, Prof. David Hachen, Prof. Omar Lizardo, Dirk Van-Bruggen, Chris Miller, and Lei Meng. Special thanks to Sprint who provided the ability to offer the phones and service plans for the experiments.

6. REFERENCES

- [1] Cisco, “Cisco VNI mobile data traffic forecast 2011-2016,” February 2012.
- [2] K. Lee, I. Rhee, J. Lee, S. Chong, and Y. Yi, “Mobile data offloading: how much can WiFi deliver?” in *Co-Next’10*. ACM, 2010, pp. 26–38.
- [3] A. Balasubramanian, R. Mahajan, and A. Venkataramani, “Augmenting mobile 3G using WiFi,” in *MobiSys’10*. ACM, 2010, pp. 209–222.
- [4] S. Dimatteo, P. Hui, B. Han, and V. Li, “Cellular traffic offloading through WiFi networks,” in *MASS’11*. IEEE, 2011.
- [5] B. Han, P. Hui, and A. Srinivasan, “Mobile data offloading in metropolitan area networks,” *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 14, no. 4, pp. 28–30, 2011.
- [6] S. Wiethölter, M. Emmelmann, R. Andersson, and A. Wolisz, “Performance evaluation of selection schemes for offloading traffic to IEEE 802.11 hotspots,” in *ICC’12*. IEEE, 2012.
- [7] S. Liu and A. Striegel, “Accurate extraction of face-to-face proximity using smartphones and bluetooth,” in *WiMAN’11*. IEEE, 2011, pp. 1–5.
- [8] A. Khan, M. Qadeer, J. Ansari, and S. Waheed, “4G as a next generation wireless network,” in *ICFCC’09*. IEEE, 2009, pp. 334–338.
- [9] M. Arshad, A. Farooq, and A. Shah, “Evolution and development towards 4th generation (4G) mobile communication systems,” *Journal of American Science*, vol. 6, no. 12, pp. 63–68, 2010.
- [10] A. P. Nathan Eagle and D. Lazer, “Inferring social network structure using mobile phone data,” *PNAS*, vol. 106, no. 36, pp. 15 274–15 278, 2009.
- [11] L. Palen, M. Salzman, and E. Youngs, “Going wireless: Behavior & practice of new mobile phone users,” in *CSCW’00*. ACM, 2000, pp. 201–210.
- [12] P. Vaz de Melo, L. Akoglu, C. Faloutsos, and A. Loureiro, “Surprising patterns for the call duration distribution of mobile phone users,” *KDD’10*, pp. 354–369, 2010.
- [13] A. Shye, B. Scholbrock, G. Memik, and P. Dinda, “Characterizing and modeling user activity on smartphones: summary,” *ACM SIGMETRICS Performance Evaluation Review*, vol. 38, no. 1, pp. 375–376, 2010.