

A Case for World-wide Network Measurement using Smartphones and Open Marketplaces

Dongsu Han and Srinivasan Seshan

February 2011
CMU-CS-11-106

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

Obtaining representative measurements with a large number of vantage points has been one of the key challenges in network measurements. Smartphones and open marketplaces for applications provide a unique opportunity for obtaining a large number of vantage points and enable measuring diverse properties of the Internet and mobile networks. This paper presents the methodology, practice and unique challenges involved in leveraging smartphones and open marketplace for large scale measurements. We examine issues with using this approach related to application adoption patterns and the diversity of users across geographic regions. Finally, we highlight examples of new measurement studies this approach enables, such as typical localization accuracy of mobile users, effectiveness of IP localization, and characterizing WiFi network environment.

This research was funded by the National Science Foundation under grant number CNS-0435382. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Keywords: cellular networks, 3G measurement, WiFi measurement, open marketplace, android

1 Introduction

One of the greatest challenges in network measurement is showing that the results are representative [8]. The most common approach is to collect measurements from many vantage points. However, obtaining access to various measurement locations has been difficult and recruiting users to perform measurements has typically been time-consuming. In this paper, we examine the use of smartphones and open marketplaces as one way to address this challenge and highlight some of the unique issues in using this approach. Note that some recent efforts [15, 21] have begun to use open marketplaces for smartphone applications to measure 3G networks and characterize the application and device usage of smartphones. We show that this approach can be used to measure more general, non-smartphone or cellular related network properties, such as the connectivity properties and communication environments of mobile users.

To explore the challenges of using smartphones and open marketplaces, we developed an application, called *WiFiMapper*, for the Android platform and made it available through Android Market. The application is designed to provide benefit to users by showing information about the network they access and at the same time report network-measurement results to a centralized server. Over 7,000 users from all over the world installed and used this application in just 20 days. We show that the approach enables new measurement studies by providing means to recruit a potentially large user-base, and explore the possibility of using the application user-base as a massive number of vantage points for continuous Internet measurements. We believe that applications of this type can provide a framework that continuously updates network information from users across the world.

This paper makes the following contributions. We present the methodology, practice and unique challenges involved in using open marketplaces to recruit many users and build an automated measurement framework. We demonstrate that the approach enables large-scale, world-wide measurements and that the approach is effective in collecting interesting measurements that were traditionally hard to collect. We further characterize the method's effectiveness by showing the application adoption pattern and analyzing the geographic diversity of the user-base. Finally, we demonstrate the capability of the approach by showing a diverse set of measurement results that this approach enables. Specifically, we verify the typical localization accuracy of Android phones, explore the possibility of bootstrapping the IP geo-location database, examine how IP addresses are geographically clustered, and characterize the WiFi network environment.

In Section 2, we present the challenges and our practice. Section 3 examines the application adoption pattern and the geographic diversity of user-base. Section 4 presents measurement results. We present related work in Section 5 and conclude in Section 6

2 Challenges

The approach of using smartphones and open marketplaces involves in a number of challenges that are different from traditional measurements. For example, the approach completely relies on users to download and use the measurement application and researchers must focus on reaching out to a large user base. We use our experience with *WiFiMapper* to illustrate these challenges as well

as some possible solutions. We breakdown the challenges in to challenges in development and deployment.

Challenges in development: Smartphone users are often battery-life conscious because the battery life of the smartphone is much shorter [7, 1]. WiFi and GPS localization are especially known to consume high power [12, 18, 9]. Smartphone platforms often report battery consumption of applications, so users can take actions such as deleting applications that use significant power. Therefore, designing an application that uses little power while still performing measurement and providing utility to users is important. In WifiMapper, we minimize the automatic background measurement, and give control over to the users for lengthy measurements to be less obtrusive as possible in terms of power consumption.

Another challenge is that there are many hardwares and software versions of the Android platform. To date, 24 different Android phones are available from at least 7 different vendors [2] and there are six versions of the Android software with three major versions¹ almost equally dividing the total share of 99% [6]. Because not all phones support all major versions, developers have to support all versions to reach out to a large user base. We used the most recent version of the API for development to read physical layer information such as signal strength of the various wireless networks, but used Java's reflection method [3] to remain backwards compatible.

Despite the abstraction that the Android platform and its API provide, hardware dependent issues still remain. For example, link setup time such as establishing WiFi connection and mobile network connection depends on the hardware, and time to get localization fix may be also different. This implies that things like timeout values and retry limits may have to consider the worst and the best case. Similarly, the battery capacities also differ and the same operations consume different amounts of energy on different hardware. Therefore, predicting battery lifetime is difficult and the applications may need to tune the frequency of different measurements on different hardware. We also found a version dependent bug in the Android API. The API call for disconnecting WiFi worked well in asynchronous fashion on G1 with Android 1.6, but would block and stall for tens of seconds on Droid with Android 2.1 platform. Therefore, optimizing for one device may not work for others, and testing on different devices is critical.

Challenges in deployment: Although the goal is to set up a framework for large-scale world-wide measurement, the methodology gives the experimenter no control over the scale and speed of deployment. This approach relies on users installing the application out of their own interest and for users to initiate more heavy-weight measurements (for reasons described above). Similarly, the experimenter has little control over when users may delete the applications, which has significant impact on the ability to perform longitudinal studies. However, while we do not have direct control over user behavior, we can indirectly influence them through the application design and the interface that the open marketplace provides.

When searching for new applications, users typically rely on the reputation and recommendation system of the market. The Android Market interfaces categorizes an application based on input from developers, and displays applications by rank. The ranking algorithm takes account for the

¹The three versions are v1.5, v1.6, and v2.1.

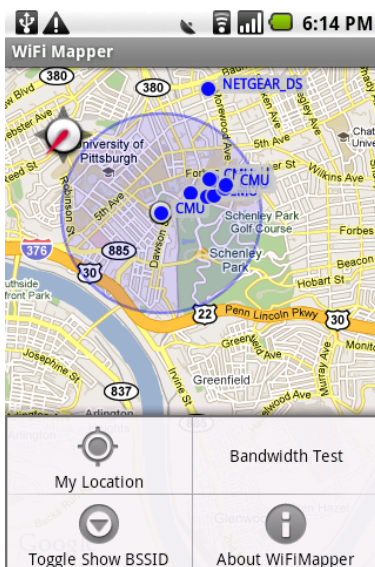


Figure 1: Screenshot of the application.

Event	Information collected
After WiFi Internet connection	BSSID, RSSI, channel, location, WiFi scan result, interface name, cellular network info
After throughput test	throughput, BSSID, RSSI, channel, location, interface name

Table 1: Information collected from the server

user rating, active installs, total installs (downloads), and updates of applications. Hence, user satisfaction and publicity are the key to the success of deployment. To incentivize users, WiFiMapper provides detailed information about WiFi connectivity and locations of hot spots they have previously accessed overlaid on a Google Map as noted in Section 2.1.

To gain publicity, we set up a website, posted an advertisement on Facebook’s Android community, and sent out emails to friends and family after the initial release. We were able to get more than 1,000 users within 4 days of the initial release. In Section 3, we present the adoption trend of the application. Currently, our application is within the top 15% ranked by the number of downloads [4]. However, the ratio of active installs to total installs ² has continuously decreased since the initial release suggesting that it is harder to keep the existing users than to recruit new users.

2.1 Android Application

Our Android measurement application, called *WiFiMapper*, is designed to provide benefit to users by showing information about the network they access. At the same time, The application collects information about the user’s WiFi access points (APs) and Internet connectivity.

²Number of active installs is number of total installs minus the number of uninstalls.

The application listens for WiFi events generated by the system in the background. When a WiFi connection is established, it attempts to connect to a centralized server to check Internet connectivity. If Internet connectivity is confirmed, the application reports its location, WiFi and cellular network information to the server. The application records information such as BSSID and RSSI of WiFi near-by networks, base station ID, RSSI, and type of the cellular network along with the location of the device. The information is kept in a local database and also sent to a centralized server. The centralized server records all the information sent from the application along with the public IP address that the phone was using through WiFi. However, we do not collect any easily identifiable user information for privacy reasons. Table 1 summarizes the information collected.

The application’s frontend displays the location of WiFi APs, their signal strength, coverage, channel, SSID and security settings on top of a Google Map. Figure 1 shows the screen-shot of the application. Users can also perform download throughput measurements on WiFi and cellular networks by clicking the “bandwidth test” menu option. The test performs a 9 second TCP transfer from the central server.

3 Methodology Evaluation

Ideally, we want a measurement application that is rapidly deployable, retains the user-base throughout the measurement period, and provides coverage of diverse geographic regions around the world. In this section, we answer the following questions: 1) How fast is the deployment and do users keep the application installed for a long period of time? 2) What is the geographical locations of users, and how diverse are they? These results are based on data collected from 7,652 users who downloaded our application during a period of 20 days.

3.1 Application Adoption

The effectiveness of the approach heavily relies on how readily users adopt the application. Figure 2 shows the number of installs and the number of active installs reported by the Android Market. The number of active installs is number of installs except the number of uninstalls.

The application gained immediate attention within the first 24 hours of release, recruiting more than 500 users. The initial exposure was from Facebook’s Android group, as our advertisement posting stayed at the top of the wall for a few hours following the release. In the Android Market, the application was visible in the “Just-In” tab of the “communication” category for the first 7 days, slowly moving towards the top of the list. At around day 7, it hit the top of the list and number of total downloads marked well over 2,000 the next day. We currently have over 7,000 downloads and 4,000 active installs and have an average rating of 4 stars out of 5.

Our results show that the combination of an effort to show users useful information and advertisements to special interest groups is enough to get a critical mass of users for measurements. However, the percentage of active installs did not keep up, but continuously fell from 80% to 58%, suggesting that keeping existing users is harder than recruiting new users. This further implies that longitudinal studies may be more difficult to conduct using this approach, and the application must

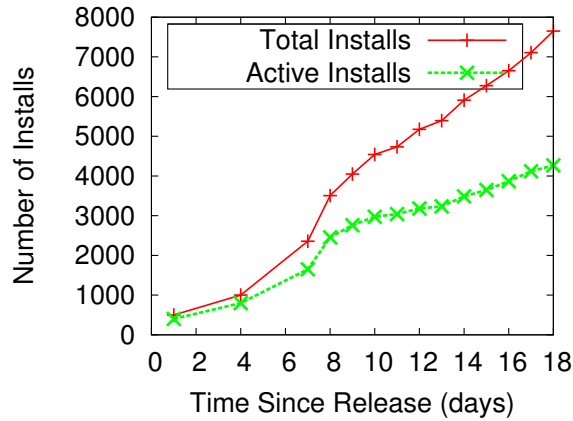


Figure 2: Growth in the number of users over time.



Figure 3: Reported locations of Access Points (over 13,000 WiFi APs).

be designed to provide greater continuous user benefit to retain existing users. We revisit this issue and examine a case of longitudinal studies in Section 4.1.2.

How do users react to updates?: We updated three new versions of the application after the initial release to add more functions and to fix bugs.

The initial update happened on day 6 after the release. Since the our application does not track users, we cannot determine the percentage of users who updated the application. However, 82% of the all measurements were reported by the new version three days after the update. However, during the three day period, the number of active installs almost doubled. As a result, the increase in use of the new version can be largely attributed to the increase in new users, who downloaded only the updated version, rather than existing users that updated their application.

United States	47.7%	China	1.3 %
United Kingdom	6.8%	Netherlands	1.3%
Germany	3.2%	Switzerland	1.2%
Sweden	2.9%	Argentina	1.2%
South Korea	2.8%	Denmark	1.2%
France	2.2%	Russia	1.1%
Canada	2.1%	Japan	1.1%
Spain	2.0%	Poland	1.0%
Italy	1.7%	Australia	0.9%
Norway	1.3%	Brazil	0.9%

Table 2: Distribution of measurements across countries

3.2 Geographic Diversity

One of the most important benefits of this methodology is that it enables measurements from all over the world. Here, we look at the diversity of geographic locations and distribution of measurements among different geographic regions. For visualization, we map the locations of APs that the application reports in Figure 3. Over 13,000 unique APs were reported from which users connected to the Internet. We see that the measurement covers area from all around the world, densely covering parts of North America, Western Europe and parts of East Asia. A total of 112 different countries and dependent territories were represented. Table 2 shows the top 20 countries ranked by the percentage of AP measurements. U.S. reported by far the most with all 50 states represented.

4 Data Evaluation

In this section, we present a few examples of general measurements that are enabled by smartphone deployments. In particular, we show results on how effective different localization schemes are in the real world and how users connect and deploy WiFi networks.

4.1 Localization

Smartphones are often equipped with GPS which provide highly accurate localization results. However, GPS is often not available indoors because of poor signal reception, and can be power hungry [12]. Also, most traditional end hosts, such as desktops and laptops, do not have GPS. Therefore, it is important to know how well alternative localization techniques perform, and how we can improve them. In this section, we examine the accuracy of wireless network localization and IP-based localization, and discuss how to improve them through our measurement infrastructure.

4.1.1 Wireless Network Localization

The application reports the location where they access APs. The Android platform supports both Assisted GPS (A-GPS) and wireless network localization. Wireless network localization uses the handset's signal strength measurements of cell towers and WiFi access points to determine location. Assisted GPS (A-GPS) uses wireless network localization to improve the startup performance and availability of GPS. In this section, we look at the typical accuracy of both GPS and network localization.

Out of more than 272,000 location reports, 67% were generated by wireless network-based localization and 33% by GPS. Each report also carries the accuracy of the localization fix. Note that many APs have multiple reports and 55% of APs only had reports with wireless network-based location.

Figure 4 shows the accuracy distribution of GPS-based and network-based localization reported by the Android system, and the actual accuracy of the network-based technique. The actual accuracy was estimated from the GPS result. We took the GPS result as the ground truth when it was taken within 30 seconds from the network-based location at the same AP and had an accuracy of at least 20m. The motivation is that we believe that GPS accuracy is far better understood than wireless localization accuracy. Median GPS and network localization accuracy is 48m and 150m respectively. The inaccuracy of GPS location is because A-GPS falls back to cell tower localization when the GPS reception is poor. Accuracy reported by GPS had spikes at (2, 4, 16, 32,...) and (3, 6, 12,...). In network-based localization, over 30% of accuracy was at 75m, and nearly 10% was at 1km. We think that these discrete steps in accuracy are artifacts of implementation, and 75m and 1km respectively represent default WiFi-based and cell tower based location accuracy. Finally, the reported accuracy of network-based location roughly matches the estimated actual accuracy of network location in terms of its distribution.

To see how often the actual location is within the accuracy from the reported location, we plotted ratio of the actual accuracy to the reported accuracy in Figure 5. When the actual location is within the reported accuracy from the reported location, the ratio is less than or equal to one. About 64% of the localization fix had was within reported accuracy and about 84% fell within two times the accuracy.

4.1.2 IP-based Localization

Another frequently used localization method on the Internet is IP localization [16, 5]. The most common approaches for IP localization infer the geographic location of Internet hosts from its IP address using delay measurement to landmark hosts, network topology information or commercial databases. Approaches using commercial database reports city level location of IP addresses but claim only 95% accuracy at the country level. While there are measurement studies showing the accuracy of a particular IP localization scheme, there is no work exploring the underlying limits of IP localization. For example, little is known about how frequently the IP address assignment changes and how ISPs assign addresses. We answer these questions by looking at how stable the individual IP assignments are and the area that a typical /24 network covers.

We observed 15,744 unique public IPs (12,892 unique /24 subnets) from 13,126 WiFi APs.

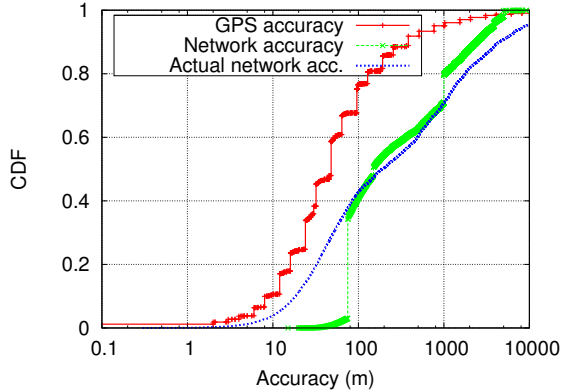


Figure 4: Reported accuracy of the localization results.

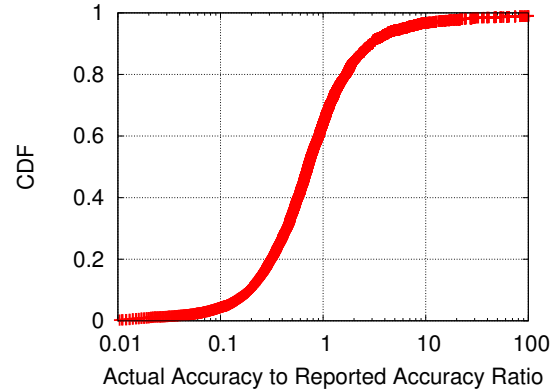


Figure 5: Estimated actual to reported accuracy.

N (days)	1	3	5	7	9	11
APs	13126	3660	19926	1116	697	404
Percentage	100	28	19	8.5	5.3	3.1

Table 3: Number of APs observed for N or more days

Table 3 shows the percentage and the number of APs that we observed for N or more days. Because of the user churn and the duration of study, only 1116 APs were observed for 7 or more days. To see the stability of the IP assignment of those APs, we plotted the number of unique IP and subnet addresses observed from the same AP in Figure 6. Almost 70% of the IP and 85% of /24 subnet assignment remained stable over 7 or more days.

To see the area that a /24 subnet covers, we calculated the bounding area of /24 subnets that were observed from multiple APs. Figure 7 shows the radius of the bounding area of 431 such subnets. 80% of /24 networks were clustered within a radius of 10km. Most notable outliers were DSL networks, a hotel hotspot provider, and a satellite Internet provider that spanned the continental US.

Discussion: We believe that smartphones can provide a bootstrapping mechanism for IP localization to improve its accuracy and provide accuracy information. Although our measurement does not have enough data to be conclusive, it suggests that IP assignment is often stable over multiple days, and most IP subnets are clustered within a small area with exceptions of some providers. We suggest that the stability of IP assignment can be exploited to provide finer grained location, and the geographic locality of subnets can be used to estimate the accuracy of localization reports.

4.2 WiFi Connectivity

The WiFiMapper application reports WiFi scan results when users connect to an AP. In this section, we characterize the environment in which users connect to WiFi networks from these mea-

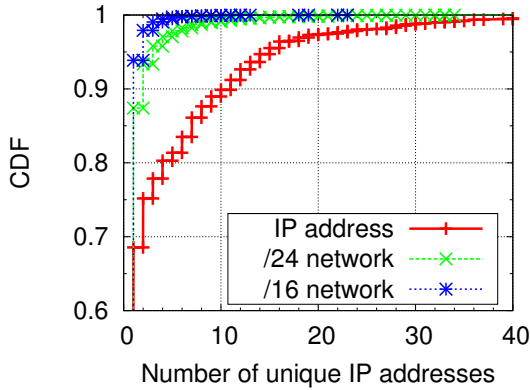


Figure 6: Host and Network IP address stability.

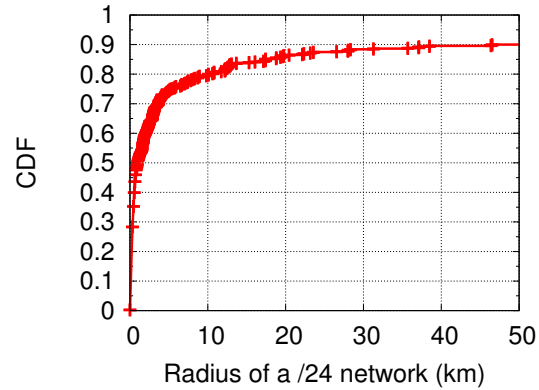


Figure 7: Radius of the enclosing area of /24 networks.

measurements. Figure 8 shows the number of APs observed when users access the Internet. We classify the environment into three equally probable regimes by the AP density. Lower 30% of times, users are in low density where they see one or two APs. In mid and upper 30%, users see 3 to 6 APs and more than 6 APs respectively.

To infer the performance impact due to interference in different regimes, we calculate the signal-to-noise-and-interference ratio (SINR) of the AP by subtracting the strongest co-channel interferer’s signal strength from the AP’s signal strength. Figure 9 shows the distribution of signal strength and the SINR of the AP in use categorized by the AP density. Note that if clients associated to the nearest AP and both APs and clients were randomly located in space, clients in dense environments would have higher signal strength from the APs. This is not the case in our measurements which shows that signal strength distribution of the AP is similar regardless of the AP density. This suggests that clients are about the same distance from their APs regardless of the density of AP deployment. The high density environment has slightly higher co-channel interference. This is explained by the fact that the SNR to the clients chosen AP is about the same while the SNR to the nearest interfering AP is likely higher in dense environments. Median SINR of high density environment was 11dB and low and mid were both 15dB. By mapping the SINR to achievable bit-rate, the measurements suggest that maximum achievable throughput may decrease significantly from 33Mbps at no interference to 7Mbps even at the co-channel interference level of low density environments [17].

Figure 10 shows the signal strength of the strongest alternative AP regardless of its channel. We can see a clear distinction and wider gap between different density than we saw in the SINR curves in Figure 9. The median gap between each density is 8 to 10 dB. This shows that while there is much stronger alternative APs in dense environment, they are often not in the same channel and, thus, they do not necessarily interfere in dense environments. This also suggests that channel selection is done more carefully in high-density environments than in other settings.

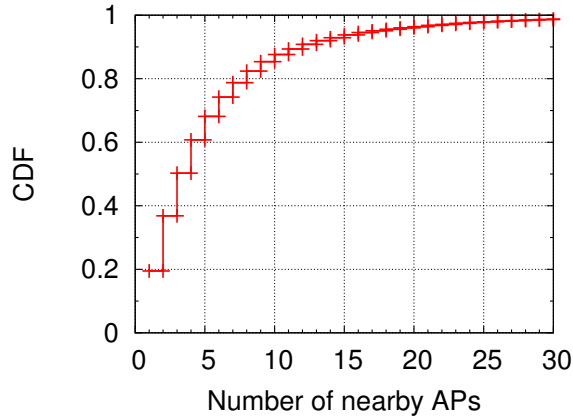


Figure 8: Nearby APs observed by WiFi scan




Figure 9: Signal strength of the APs in use

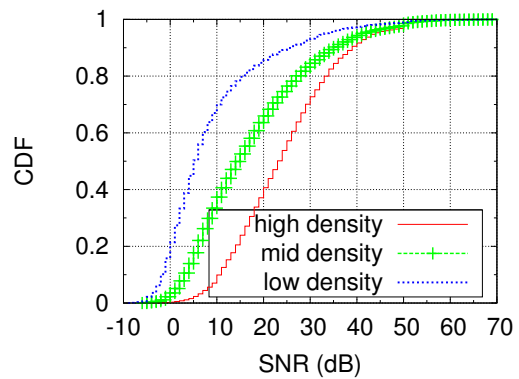


Figure 10: Signal strength of strongest alternative AP

5 Related Work

Recently, Huang et al. [15] developed applications for Android, iPhone and Windows Mobile, and performed the first wide-scale cross-platform study on smartphones. They focused on the 3G network performance of different providers in the U.S. and application performance, such as page loading time, Javascript execution time, and video streaming behaviors, on different platforms. Our study focuses on the effectiveness of the method of using smartphones and open marketplace and highlight new measurement studies the method enables.

Many user behavioral studies have also been conducted on smartphones to better understand how user and interact with the phones by measuring application, network and battery use [10, 13, 19, 20]. These studies typically perform an in-depth measurement or survey from a pre-selected sample of users by giving them instrumented phones that log user activities. Shye et al. used similar approach to ours by making a logger application through the Android Market to characterize the power consumption [21]. Our study shows that the method enables large scale measurements that goes beyond the scope of mobile networks and phone usage.

Finally, Apollo [14], NETI@home [22], and ONO [11] are systems that either rely on user participation or users using variants of existing popular applications to collect measurements from traditional end hosts.

6 Conclusion and Future Work

In this paper, we examined using smartphones and open marketplaces to create a large-scale measurement infrastructure. We demonstrated that the method is effective in recruiting a large user-base across the world. We showed examples of measurement studies the method enables and made observations on localization accuracy, WiFi network environments, stability of IP assignment and their geographic locality.

We believe that this approach can provide richer and broader set of information about user behavior and network environment that have strong systems implications. We plan to continue on expanding the infrastructure by adding more measurement and by providing richer information to the users.

References

- [1] T-mobile g1 battery life solution; turn off everything!, 2008. www.zdnet.com/blog/cell-phones/t-mobile-g1-battery-life-solution-turn-off-everything/210.
- [2] Android phones, 2010. www.androphones.com/.
- [3] Backward compatibility for applications, 2010. <http://developer.android.com/resources/articles/backward-compatibility.html>.
- [4] Distribution of app downloads in the android market, 2010. <http://www.androlib.com/appstatsdownloads.aspx>.
- [5] IP2location, 2010. www.ip2location.com.
- [6] Platform versions, 2010. <http://developer.android.com/resources/dashboard/platform-versions.html>.
- [7] Smartphone battery life: Cell phones with the best battery life, 2010. www.geek.com/smartphone-buyers-guide/battery/.
- [8] A. Akella, S. Seshan, and A. Shaik. Toward representative Internet measurements. In *Proc. New York Metro Area Networking Workshop (NYMAN)*, September 2003.
- [9] Niranjan Balasubramanian, Aruna Balasubramanian, and Arun Venkataramani. Energy Consumption in Mobile Phones: A Measurement Study and Implications for Network Applications. In *Proc. ACM IMC*, November 2009.
- [10] Nilanjan Banerjee, Ahmad Rahmati, Mark D. Corner, Sami Rollins, and Lin Zhong. Users and batteries: interactions and adaptive energy management in mobile systems. In *Proc. UbiComp*, 2007.

- [11] David R. Choffnes and Fabián E. Bustamante. Taming the torrent: a practical approach to reducing cross-ISP traffic in peer-to-peer systems. In *Proc. SIGCOMM*, 2008.
- [12] Ionut Constandache, Romit Roy Choudhury, and Injong Rhee. Towards mobile phone localization without war-driving. In *Proc. IEEE INFOCOM*, 2010.
- [13] Hossein Falaki, Ratul Mahajan, Srikanth Kandula, Dimitrios Lymberopoulos, Ramesh Govindan, and Deborah Estrin. Diversity in smartphone usage. In *Proc. ACM MobiSys*, San Francisco, CA, June 2010.
- [14] Marios Iliofotou, Georgos Siganos, Xiaoyuan Yang, and Pablo Rodriguez. Comparing BitTorrent clients in the wild: The case of download speed. In *Proc. USENIX IPTPS*, San Jose, CA, April 2010.
- [15] Junxian Huang, Qiang Xu, Birjodh Tiwana, Z. Morley Mao, Ming Zhang, and Paramvir Bahl. Anatomizing application performance differences on smartphones. In *Proc. ACM MobiSys*, San Francisco, CA, June 2010.
- [16] Ethan Katz-Bassett, John P. John, Arvind Krishnamurthy, David Wetherall, Thomas Anderson, and Yatin Chawathe. Towards IP geolocation using delay and topology measurements. In *Proc. IMC*, 2006.
- [17] D. Naudts, S. Bouckaert, J. Bergs, A. Schouttset, C. Blondia, I. Moerman, and P. Demeester. A wireless mesh monitoring and planning tool for emergency services. In *Proc. E2EMON*, May 2007.
- [18] Trevor Pering, Yuvraj Agarwal, Rajesh Gupta, and Roy Want. CoolSpots: reducing the power consumption of wireless mobile devices with multiple radio interfaces. In *Proc. MobiSys*, 2006.
- [19] Ahmad Rahmati, Angela Qian, and Lin Zhong. Understanding human-battery interaction on mobile phones. In *Proc. MobileHCI*, 2007.
- [20] Ahmad Rahmati and Lin Zhong. A longitudinal study of non-voice mobile phone usage by teens from an underserved urban community. Technical Report 0515-09, Rice University, May 2009.
- [21] Alex Shye, Benjamin Scholbrock, and Gokhan Memik. Into the wild: studying real user activity patterns to guide power optimizations for mobile architectures. In *Proc MICRO 42: IEEE/ACM International Symposium on Microarchitecture*, 2009.
- [22] Charles Robert Simpson, Jr. and George F. Riley. NETI@home: A distributed approach to collecting end-to-end network performance measurements. In *Proc. PAM2004 - A workshop on Passive and Active Measurements*, April 2004.