Characterizing and Modeling User Activity on Smartphones: Summary

Alex Shye Benjamin Scholbrock Gokhan Memik Peter A. Dinda

Dept. of Electrical Engineering and Computer Science

Northwestern University, Evanston, IL, USA

{shye, scholbrock, memik, pdinda}@northwestern.edu

ABSTRACT

In this paper, we present a comprehensive analysis of real smartphone usage during a 6-month study of real user activity on the Android G1 smartphone. Our goal is to study the high-level characteristics of smartphone usage, and to understand the implications on optimizing smartphones, and their networks. Overall, we present 11 findings that cover general usage behavior, interaction with the battery, power consumption, network activity, frequently-run applications, and modeling usage states.

Categories and Subject Descriptors: C.3 [Special-Purpose and Application-Based Systems]: Real-Time and Embedded Systems; C.0 [Computer Systems Organization]: General-Modeling of Computer Architecture; H.1.2 [Models and Principles]: User/Machine Systems-Human Factors

General Terms: Measurement, Human Factors

Keywords: Human Factors, Embedded Systems

1. INTRODUCTION

We present a comprehensive analysis of real smartphone usage from a 6-month study, involving 25 users, and one specific smartphone, the Android G1. The main goal of our study is to observe the high-level workload characteristics of real smartphone users in the wild and understand the implications of these characteristics for optimizing smartphones. Specifically, we are interested in the following questions:

- What does typical user activity look like? How does it vary across users?
- What are the most energy consuming hardware components? What are the execution characteristics of the most power hungry components?
- What are the network connectivity characteristics on a smartphone?
- Can we detect high-level patterns in user activity?

For reference, a more detailed presentation of this work is available elsewhere [4].

2. METHODOLOGY

Target Architecture. Our target smartphone in this paper is the HTC Dream, marketed as the Android G1, a cellular phone platform built by HTC that supports the open source Google Android mobile device platform. A high-level diagram of the G1 is shown in Figure 1. The G1 has a 3.2inch HVGA 65K color capacitive touch screen, uses a Qualcomm MSM7201A chipset, and a 1150mAh lithium-ion battery. The Qualcomm MSM7201A chipset contains a 528MHz ARM 11 apps processor, a ARM 9 modem processor, QDSP4000 and QDSP5000 high-performance digital signal processors, 528MHz ARM 11 Jazelle Java hardware acceleration, quadband GPRS and EDGE network, integrated bluetooth, and wifi support.

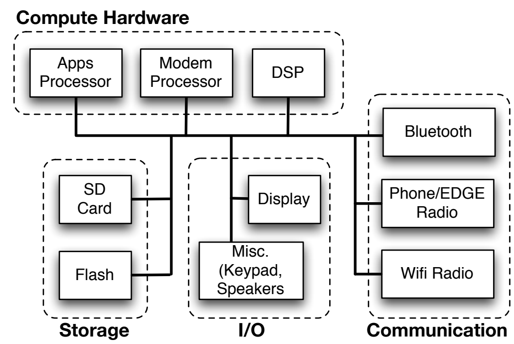

Figure 1: High-level diagram of our target mobile architecture, the Android G1 smartphone.

Logging User Activity. To study the real usage of the G1, we have developed a logger application that logs user activity events, as well as system-level performance measurements. The logger is developed as a normal Dalvik executable using the Java standard libraries available in the Android framework. It runs on the G1 without any special hardware or OS support.

Obtaining users. We posted online advertisements and physical flyers for anonymous volunteers on various university campuses, technical news web sites, and Android-related forums. Overall, we collect logs from 52 users from April–November 2009. For this paper, we use the logs from the 25 users with the longest total recorded time. The data from these 25 users represents approximately 1329 days ($\sim$ 3.6 years) of real user activity, with an average of approximately $\sim$ 53 days of logged user activity per user.

Power estimation. We use a regression-based power estimation model for the Android G1 that has been proposed and validated in prior literature [3].

3. OVERVIEW OF FINDINGS

We summarize our findings in Table 1. The table includes our 11 main observations, split into several categories: gen-

SIGMETRICS'10, June 14–18, 2010, New York, New York, USA.
ACM 978-1-4503-0038-4/10/06.

General Observations	Implications
1. Users recharge their phones on a daily basis, and use their phones until the battery is low in $\sim 20\%$ of the cases when it is unplugged for over 4 hours.	Battery management is a significant part of the smartphone user experience, including daily charging and frequent low battery indicators.
Improving Energy Efficiency	**Implications**
2. There is significant variation in the usage behavior of our users. For example, some users are heavy phone users, while others may be heavy wifi users.	It is critical to perform real user studies when studying smartphones. In addition, the large variations in usage activity may represent optimization opportunity [1, 2].
3. Active phone use consumes 53.7% of the total system power and 11% of the usage time. Of the hardware components, the screen and the CPU consume the most power, 19.5% and 7.3% of the total power, respectively.	Active phone use consumes the majority of the power, even though it accounts for a small fraction of the total use time. The screen and CPU require the most attention with respect to energy efficiency.
4. On average, the phone is in the idle state 89% of the time and accounts for 46.3% of the total system power.	Reducing the power consumption of the idle state should also be a high priority for improving battery life.
5. Most users do not switch between multiple brightness levels, nor do they install power management software.	Automatic brightness adjusting optimizations (as well os other power optimizations) should be included with smartphones.
6. The CPU utilization is typically either at 100%, or under 10%.	Dynamic CPU scale-down optimizations should be used for saving power.
Networking	**Implications**
7. EDGE network session durations follow a power-law distribution.	Session durations can be modeled with a General Pareto Distribution, with a a shape parameter of 2.7454 and a scale parameter of 2.4732.
8. Wifi network session durations appear to be the sum of several distributions and the wifi network traffic is highly dependent on daily usage modes.	Modeling wifi session durations warrants more investigation and shows promise of revealing trends in user behavior. Session durations can be modeled with a MMPP.
Application Usage	**Implications**
9. A significant portion of CPU utilization is attributable to OS-level processes.	From the perspective of mobile computing, OS developers must be aware of the broader impacts of frequently-run code (e.g., on power consumption due to the CPU).
Usage Patterns	**Implications**
10. From a large space of possible states, only a few significant states and transitions are required to meaningfully represent smartphone usage patterns.	Building a useful state-transition graph to model smartphone user behavior from a large dataset is tractable.
11. Automatic clustering of usage logs closely matches manually selected states of interest.	Meaningful states may be extracted automatically from input data to build a state-based model of user behavior.

Table 1: Summary of the 11 main findings in this paper.

eral, energy-related, networking, application, and usage patterns. Each observation includes a description of the implications on energy-efficient design and/or optimization for the Android G1 smartphone.

In general, we see that studying real user activity yields many interesting findings. Our user activity traces show that (1) battery life is a significant part of the user experience, (2) there exists large variation in usage patterns across our users, (3) the screen and CPU are good targets for power optimization, (4) clear differences in EDGE and wifi network traffic can be seen in usage behavior, (5) OS activity plays a large role in execution, and (7) user activity can be automatically clustered to produce Markov decision processes for modeling individual users (an example shown in Figure 2). The last of these findings is particularly significant as it points towards automatic generation of models for various aspects of smartphone usage (e.g., modeling EDGE network traffic from individual devices).

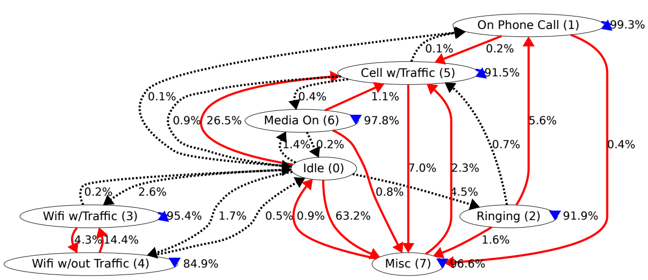

Figure 2: Example of a user activity model derived from clustering real user activity traces.

4. REFERENCES

[1] P. Dinda, G. Memik, R. Dick, B. Lin, A. Mallik, A. Gupta, and S. Rossoff. The user in experimental computer systems research. In *Proceedings of the Workshop on Experimental Computer Science*, June 2007.

[2] A. Shye, B. Ozisikyilmaz, A. Mallik, G. Memik, P. A. Dinda, R. P. Dick, and A. N. Choudhary. Learning and leveraging the relationship between architecture-level measurements and individual user satisfaction. In *Proceedings of the Intl. Symposium on Computer Architecture*, June 2008.

[3] A. Shye, B. Scholbrock, and G. Memik. Into the wild: Studying real user activity patterns to guide power optimizations for mobile architectures. In *MICRO*, December 2009.

[4] A. Shye, B. Scholbrock, G. Memik, and P. A. Dinda. Characterizing and modeling user activity on smartphones. Technical Report NWU-EECS-10-06, Northwestern University, Evanston, IL, March 2010.