

Notes gathered from Day 2 of HotMobile 2014
by Kartik Muralidharan (SMU), Terence Chen (University of New South Wales), and Shuo
Deng(MIT)

Some of the Q&A notes gathered by Stefan Saroiu (MSR)

1. Mobile Systems and Applications

This session focussed on the innovative use of smartphone sensors and wearable devices to support a host of novel context-based services. The first paper, “Sinabro: Opportunistic and Unobtrusive Mobile Electrocardiogram Monitoring System”, presented by Seungwoo Kang from Korea Advanced Institute of Science and Technology, talks about a mobile electrocardiogram (ECG) monitoring system that unobtrusively monitors the user’s heart rate during daily smartphone use. The prototype does this by attaching multiple sensing electrodes along the smartphone to capture bioelectric signals when the user touches the phones during regular interactions. Interesting questions were raised about the validity and limitations of the approach, but overall, the idea was well-received.

Q&A Session:

Prabal Dutta (U. Michigan): Do you have any evidence taking the heart rate when using a phone is representative of a heart rate?

A: People use their phones in the morning and that is the resting heart rate.

Prabal: What is the relationship between the resting heart rate and using the phone heart rate.

Matt Welsh (Google): How do you know the quality of ECG signal is adequate given the way the signal is collected?

A: We have focused so far on gathering the data and we would like to study this further in future work. We believe that detecting peaks is possible even if the signal is noisy.

MW: Does it require direct contact with skin? I always keep my phone in my pocket.

A: It does need contact with your skin.

Gaetano Boriello (UW): It needs two contacts with the skin on the opposite sides of the heart.

JeongGil Ko (ETRI): Depending on how many points you are touching, the quality of the signal varies. You should try to exploit that potential.

The following talk titled “Barometric Phone Sensors More Hype Than Hope!” by Kartik Muralidharan from Singapore Management University explores the use of the barometer sensor, available on newer smartphones, in determining the exact floor a user is on in an indoor environment such as a shopping mall. He shared their experience in using this sensor and showed that despite the various challenges, the sensor could still be used to detect the exact

number of floors changed by a user as well as the mode of transportation used to change floors.

Q&A:

Ramon Caceres (AT&T): One of the reasons for pressure changing over the day is weather. Have you considered measuring the quality of the barometers against ground truth?

A: We considered this, but everyone deploying this must now calibrate their phones. Also, we saw how the pressure reported in Singapore through the official channels is different from

Eyal de Lara (U. of Toronto): You could imagine installing a probe on each floor that would be reported to a central server. This would let you use the dependency or correction between each device and the probe.

A: Yes, we mention this possibility in the paper. However, this would need additional infrastructure cost.

Vishnu Navda (MSR): Given the dependency between the pressure and temperature, could you have use the temperate sensor to calibrate?

A: Temperature also varies when the user is moving. It is not clear this is usefl.

Gaetano Boriello (UW): We built multi-sensor boards in the past and the recommendation from that work was to think about relative change and not absolute values. For example, first figuring out when the user enters a building and then keeping track of the changes. Keeping track of changes you could do a pretty good job of what floor the user in. You can then enhance this with additional context information.

A: We will take a look at this other work.

Fabian Bustamante (Northwestern): Have you measured a baseline for each phone?

A: We have a reference point when the user is starting to move up and down, and using this reference we can keep track of the phone.

FB: How much variation do you get with the same stationary device over time?

A: We measured this in the paper.

In his paper, "QuiltView: a Crowd-Sourced Video Response System", Chen from Carnegie Mellon University, presents a system that leverages the ability of wearable devices such as Google Glass to provide videos as a response to context-based queries posted by users. A prototype of the system, which was also demoed during the workshop, showed how queries such as "Is the beach crowded?" could be better (and more faster) answered through brief video segments - using the Google Glass - rather than giving detailed verbal responses that could result in ambiguity. Chen indicated that a follow up of this work would answer key

questions regarding privacy issues of the current system as well as investigate the system optimizations that were discussed following his presentation.

Q&A Session:

Alec Wolman (MSR): Have you thought of the privacy issues of passerbys?

A: It is true that this is a concern. We haven't done much yet, but perhaps blurring/removing faces in the video helps.

Nate (Buffalo): What is the delay between capturing the video and answering the query?

A: Currently YouTube takes some time, anywhere between 10 seconds to minutes. However, perhaps we could use other systems in the future that remove this limitaiton.

Matt Welsh (Google): What is the contribution beyond an interesting application? Jelly is a similar example where people can post pictures of their questions and their friends can answer.

A: We believe that our system is possible only through a Glass-like device, whereas Jelly is different.

David Chu (Microsoft): Have you thought of more sophisticated optimization to solve flooding the users?

A: We have thought about connecting queries for the same context. There is a lot of potential here.

Archan Misra (SMU): Have you evaluate the impact of shifting the cognitive burden from answering the query to asking the query?

A: No, we have not done much. The focus of our project is to reduce the burden on the person answering the query.

The final presentation of the session, "LiveLabs: Building An In-Situ Real-Time Mobile Experimentation Testbed" was by Rajesh Krishna Balan from Singapore Management University. Rajesh talked about their plans and the lessons learned in putting together a mobile experimentation testbed that aims to support 30,000 users! He also highlighted some of the research challenges that the mobile community would need to tackle to make such a system practical. Several attendees questioned the finding of privacy being a secondary issue to most LiveLabs users as compared to energy concerns and whether this would hold across different demographics/cultures. Rajesh stated that he currently does not have an answer and it would be interesting to view the outcome of performing the same type of experiment with the same type of users in different environments. Discussions also included the need for different incentives, ranging from one time payments to building specialized mobile apps, in order to ensure continued user participation.

Q&A Session:

Roy Want (Google): Have you looked at Google analytics?

A: We have already connections with several companies and we're using their analytics platforms.

Roy: Why people don't use their app? Do they uninstall it? Is this a background app?

A: Multiple reasons, no single reason. Software updates can disrupt our services.

Q: Can you elaborate on the incentives provided?

A: Everything except data plans and phones, e.g., one-time payments, continuous payments, and so on. People come for different incentives, and we need to have a whole bag of incentives, but we haven't figured out which ones are most effective.

Nigel Davies (Lancaster): Do you have a sense how to incentivize people long-term beyond money.

A: Given our demographics, our users (students) want apps. We're building some of these opportunities to build apps for them.

Doug Terry (MSR): How do you know an experiment is successful or not.

A: Great question about experiment design. We assist you, and you can re-run your experiment given an if-else scenario.

At the end of the session (for panel), Matt Welsh from Google pointed out that in certain cases it is hard to scale up experiments, especially those involving real users. He started a discussion on whether as a community we are putting undue weight on conclusions drawn from such small scale experiments and if we should even attempt to answer a research question being fully aware that it will lack clarity owing to unavailable resources.

Panel Discussion:

Nigel Davies (Lancaster): What do you think would constitute a good study of privacy and implications to different demographics/cultures?

A (Rajesh): Very interesting question, and I don't know the answer. Perhaps doing the same type of experiment with the same type of users and run an experiment in different environments.

Roy Want (Google): Users have a hard time with contextual information. Aren't these systems overloading them?

Kartik (SMU): We are trying to not "bombard" the users with information and take steps on this.

Unknown: We use static preferences; users can use these preferences to reduce their burdens.

Archan Misra (SMU): Our experience in Asia is that retailers want to send everyone the coupons.

Matt Welsh (Google): It's hard to scale some of these experiments up, and hard to draw strong conclusions. This is a challenge for the community. Some of these questions cannot be answered given the small-sized of some of these experiments.

Satya(CMU): Science progresses by asking question. We might not be able to answer these questions in full generality.

Stefan Saroiu (MSR): All studies are relevant and you have to have a caveat

and quantify the answers.

Geoff Chalen (Buffalo): I agree with both points, we have to be careful when we try of extrapolate.

Fabian Bustamante (Northwestern): I think we need to change the expectations on how to write papers -- we should start by writing the biases of our studies upfront in our paper.

Stefan Saroiu (MSR): Quiltview is proposing a new model of query/answer information, where the user asks a query and the answer is a video. Are you going to continue to investigate alternative models of querying information?

Satya (CMU): Hard to tell what the right model. Nevertheless I think that given that human attention is limited even more limited than energy, the question is how can I borrow a little bit of attention to get out a lot of value?

2. Privacy

The two papers presented in this session exposes the security vulnerabilities of some mainstream services. The first paper, "Information Leakage through Mobile Analytics Services" was presented by Terence Chen from the University of New South Wales. He shows the ease with which an adversary, through device spoofing, can extract an individual's profile and mobile applications usage information on two major mobile analytics services, Google Mobile App Analytics and Flurry. He also demonstrates that it is possible to manipulate these analytics services to influence the ads served to users' devices.

The following talk, "Unveiling the Hidden Dangers of Public IP Addresses in 4G/LTE Cellular Data Networks" by Wai Kay Leong from the National University of Singapore, highlights attacks that are possible when using a public IP address provided by cellular services in Singapore. He interestingly shows how an attacker can completely exhaust a users monthly mobile data quota within a few minutes (in addition to other attacks) and talks about ways to mitigate the same.

The session ended with Matt Welsh from Google suggesting that it is the duty of the community to inform service provides when such vulnerabilities in their system are uncovered. He also pointed out that the tie between security and the need to root phones have to be broken in order to progress towards a more secure mobile platform.

3. Networking

Yongtae Park of Korea University opened this session with "Software Radio on Smartphones: Feasible? In his talk Yongtae demonstrates that the current smartphone hardware is capable of supporting the real-time execution of downloadable MAC/PHY software for low speed technologies such as ZigBee, and is close to providing real-time operation for higher-speed technologies like Wi-Fi. He highlighted the need to factor in the processor architecture like multiple cores and SIMD instruction sets in order for SDRs to be feasible. Overall, this paper generated excitement within the audience in the anticipation that SDRs will soon come into existence.

The next talk, "Prometheus: Toward Quality-of-Experience Estimation for Mobile Apps from Passive Network Measurements" was presented by Jeffrey Pang from AT&T Research

Labs. Prometheus uses machine learning techniques to passively estimate the Quality of Experience (QoE) metrics for evolving smartphone services such as VOIP and OTT video delivery. Compared to previous approaches, Prometheus does not require any control over app services or domain knowledge about how an app's network traffic relates to QoE.

The final talk of the session was by Shuo Deng from MIT's Computer Science and Artificial Intelligence Laboratory. In her presentation, "All Your Network Are Belong to Us: A Transport Framework for Mobile Network Selection", she address the issue of network selection under the umbrella of multiple radios (cellular radios, Wi-Fi, Bluetooth) and their corresponding modes available on current smartphones. She highlights that with increasing cellular speeds Wi-Fi and LTE throughputs are often comparable making the problem of network selection more challenging as compared to previous systems where Wi-Fi is always assumed to perform better. Their system Delphi, infers useful properties of different networks as well as shares what they learn with other nearby devices and finally selects a network using what it has observed locally and/or from its neighbours.

4. Energy and Performance

Wook Song from Seoul National University began the last session with his presentation, "Reducing Energy Consumption of Smartphones Using User-Perceived Response Time Analysis". In his work, Wook Song proposes a redefinition of the traditional 'system response time' metric to one that is more aligned towards the usage patterns of smartphones. Their new metric, 'display-centric response time' distinguishes a task execution in two parts, one affecting the visible portion of the display and the other not a affecting the visible portion of the display. Doing so allows them to apply intelligent system optimizations that can reduce the CPU energy consumption on smartphones significantly.

The following talk, "Cloud is not a silver bullet: A Case Study of Cloud-based Mobile Browsing" was by Ashiwan Sivakumar from Purdue University. His work aims at understanding the performance implications of mobile cloud browsing solutions by comparing the two extreme points – one that does not use the cloud at all, and another that primarily relies on the cloud. Given the ever increasing usage of the cloud to augment mobile activities the topic was considered timely and was well received.

The final presentation of the workshop was by Tarun Bansal from Ohio State University. In this talk, "RobinHood: Sharing the Happiness in a Wireless Jungle", Tarun proposes an interference nulling scheme that leverages the high density of the access points present in today's Enterprise Wireless LANs. With this scheme multiple mobile devices are able to transmit simultaneously to multiple access points (APs) within a single collision domain. Results from a trace-driven simulations showed that RobinHood provides a throughput improvement over TDMA and IEEE 802.11.