



## HotMobile 2013

*Aarathi Prasad and Michael Butkiewicz*

The Workshop on Mobile Computing Systems and Applications (HotMobile) has always been a smaller, more intimate forum, leading to more conversational discussions compared to larger conferences. This year's attendees engaged in group discussions on a plethora of topics during the conference's workshop sessions and breaks, all while enjoying delicious Southern food and walking along the beach of Jekyll Island, Georgia. This year saw the broadest participation yet, with approximately 90 attendees, due in part to the organizers' ability to negotiate low fees and, for the first time, offer student travel grants. (As grant recipients, we thank the sponsors for their continued support!) Approximately 12 women attendees also gathered for a Sisters lunch on the second day of the workshop (<http://anitaborg.org/initiatives/systers>).

This year's HotMobile highlighted how mobile computing has helped improve lives—for example, by providing better facilities to people in areas with resource constraints or by helping traffic-clogged cities with carpooling. Yet according to Thad Starner, the keynote speaker, mobile computing still has a long way to go in reaching its full potential and improving users' lives.

### THAD STARNER'S KEYNOTE

Starner presented his work on wearable computing and shared his successes in building prototypes that have predicted future trends. He highlighted the need for on-the-go computing, letting users complete tasks on mobile devices

without disrupting activities or conversations. He discussed Google Glass, a heads-up display for augmented reality that superimposes its display on what users would normally see. He also discussed Twiddler, a handheld device that can be used to control such displays. He reported that the device has a learning curve three times faster than a regular keyboard and allows typing of up to 130 words per minute. Combining Google Glass and Twiddler provides on-the-go computing, and Starner hopes they'll pave the way for seamless computer interactions.

Starner then presented two recent projects. The first project, Mobile Music Touch, teaches users through muscle memory to play music, even as they perform another task requiring their complete attention. Mobile Music Touch has helped patients who lost sensation in their fingers regain that ability. Starner also worked with researchers of wild dolphins to build a device that emits tones understood by dolphins.

Starner observed that people use wearable computing primarily for productivity tools, as opposed to what they said they might use it for—making emergency calls.

### HOTMOBILE SESSIONS

This year's program comprised the following seven sessions.

#### Cellular Billing

Himanshu Raj presented the idea of *split billing* between mobile operators and content providers, letting content providers pay for the traffic generated

by mobile customers using their services. This content subsidization is given to mobile operators as incentive to continue improving their network infrastructure. To provide trustworthy proof of a device's mobile traffic, SIMlets are created using trusted hardware. One audience member voiced concern over system scalability across different device types. Raj agreed that scalability is an issue, but the overhead of building the network-interface-card driver into the trusted computer environment was similar to adding the driver to the operating system.

Younghwan Go explained how TCP's retransmission mechanism could inflate data consumption for the average cellular user. Basing Go and his colleagues' findings on the policies of US and South Korean ISPs, Go went on to explain several attacks on the cell carrier and customer due to differing policies for handling retransmissions. An audience member asked if the user could simply account for retransmission through an app and judge data usage this way. Go responded by saying clients can't see packet loss if the loss happened in the Internet or backhaul network.

During the session discussion, an audience member asked for ways to prevent ISPs from overcharging. This led to an interesting discussion about whether economic issues would be solved in time by natural market forces. Some felt that users had little say over data plan prices, while others felt that user abandonment would cause prices to equalize to an acceptable level. When asked about the direction in which

## CONFERENCES

research should proceed, both Raj and Go agreed that better security as well as transparency in terms of ISP data billing were necessary to give users the power to decide how best to use their devices.

### Power Management

Zhan Ma presented a study on the relative power consumption of Active-Matrix Organic LED (AMOLED) displays on smartphones against that of multimedia apps. Ma and his colleagues investigated how energy-usage requirements of these displays differ based on the displayed content (color). They investigated this by performing a power-consumption benefit analysis of several automatic dynamic color-tuning techniques.

Mostafa Uddin explained how to leverage the low power consumption of the audio interfaces to reduce the wake-up events of a Wi-Fi interface in a power-saving mode. Using these audio beacons—emitted above the frequency of human perception—is envisioned for use in close, peer-to-peer situations, such as Wi-Fi Direct. One interesting question raised regarding combining audio and Wi-Fi technologies was how to deal with the different signals' various absorption and speed characteristics.

Marcelo Martins introduced the notion of *application modes* to address the option paralysis typical of other power-management systems. The system works by separating concerns between the application, OS, and user, exposing the user only to the high-level trade-offs between battery lifetime and functionality.

During the session discussion, an audience member asked if we've reached the point where the easy fixes have run out, meaning further optimization will require user input. The presenters were divided in their responses. Uddin thinks there are still many small areas left to investigate before the switch to a whole power-management system is required. Martins feels that, although passive solutions might have potential, we need

more aggressive approaches. Ma stated clearly that all power management must be automatic; manual approaches will never be approved in business situations for usability reasons.

One audience member asked, "Even if the hardware is as good as it can be in terms of power management, how do you handle being governed by the app developers, who might be running extraneous, power-consuming code on the phone?" Martins felt that the current app market's review-based systems could address the problem by letting users complain about energy-hogging apps. One audience member suggested showing an alert to the user, similar to the current "this app is using your location information," when an application starts performing power-hungry operations.

**Xuan Bao explained how computing systems can benefit from embedding human psychological factors deeper into their design, rather than just their interfaces.**

### Understanding Humans

Xuan Bao explained how computing systems can benefit from embedding human psychological factors deeper into their design, rather than just their interfaces. The areas of psychological integration shown were diverse, ranging from overnight data prefetching for mobile devices to psychology-aware compression. One audience member asked how far you could model human beings, given that what works for one person might not work for another. Bao responded that the system should only try to make high-accuracy predictions when enough data has been collected, and if user discomfort is detected, some kind of back-off mechanism is needed. Another audience member pointed out

that the required context sensing would consume additional energy.

He Wang proposed human recognition based on clothing identification rather than facial recognition. For each user, a fingerprint, based on image data, is created, starting first as a self-photograph, then becoming enriched as the system uploads additional information gathered from other system users who identify the person from different perspectives. An audience member asked if the constant changing of clothes would create a database-scalability issue. Wang replied that you could use location information to prune the database information.

Shahriar Nirjon explained the difficulty of extracting acoustic features in real time when using a high sampling rate, to support voice-driven apps on resource-constrained devices. By exploiting human speech's frequency domain sparseness, Nirjon and his colleagues approximated a widely used feature of speech, the Mel-Frequency Cepstral Coefficients. Nirjon added during an audience discussion that this same method could be applied to speaker identification as well.

The conversation during the discussion section focused on the privacy concerns raised by two of the papers. Bao stressed that in his approach, all of the topic analysis is done on the smart device; only the derived topics are used as queries over the network. Wang discussed how the major privacy concern of facial data is eliminated in his system by instead relying on clothing. In addition, data is only shared with people who can physically see you, further limiting opportunities of malicious data usage. Spurred by the public image collection of the InSight system, the discussion led to the question of how a user could be informed or opt out of such public-data-collection systems. Suggestions included a standardized sign stating that data collection was occurring in the area; a symbol, such as a QR code, which (when worn) would visibly tell systems not to track the user; or smartphone-generated

beacons, which would tell systems within transmission range of the user's desire not to be tracked.

### Sensors and Data

The second day of the conference started with a session on mobile interaction with sensors.

Yu Xiao presented a virtual machine (VM)-based deployment technique to reduce the difficulty of deploying large-scale crowd-sensing applications. This let phones interact with an application VM, a clone of the application server, on the cloudlet. Xiao pointed out that the VM ensured privacy by storing data in the trusted cloudlet instead of the untrusted cloud. However, an audience member pointed out that VMs were heavyweight and could affect performance and efficiency. The question about who owns the data and controls the VMs also came up.

Waylon Brunette presented the Open Data Kit 2.0, designed based on feedback from users of a previous ODK version. While replying to a question, Brunette pointed out that ODK 2.0 wasn't completely backward compatible with 1.0.

Supriyo Chakraborty presented a privacy mechanism to limit information disclosure to reduce inferences when the number of users isn't large enough to apply known anonymization techniques. This mechanism lets users specify sharing policies, based on their privacy and usability preferences. This led to a discussion about practical privacy, where the audience members pointed out that any method that lets users specify sharing policies is impractical, because users might not know what inferences can be made from their data.

### Panel on Mobile and Developing World

Following the session on sensors was a panel on mobile systems and the developing world, moderated by Gaetano Borriello. The panelists—Elizabeth Belding, Lakshmi Subramanian, and Bill Thies—spoke about their experiences working on mobile device

projects in the developing world. Belding explained that phone-based solutions should interoperate with feature phones, cater to illiterate users, use minimal power and limited bandwidth, and work even without cellular connectivity. She also highlighted the importance of working with people outside the field of technology, to accomplish tasks that aren't related to computer science and obtain continuous feedback from users. Subramanian talked about four of his recent projects, while Thies talked about how innovative ideas that would benefit the developed world were discovered while trying to find solutions for the challenges the developing world faces.

During the discussion, Matt Welsh asked whether the challenges of the developing world were simply

### Tan Zhang explored using TV-whitespace-based communication systems to provide connectivity to vehicles.

related to bandwidth and would hence diminish in the coming years. The panel agreed that the needs of the people will persist and so will the solutions, and that part of the goal of researchers in the Information and Communication Technologies and Development field is to accelerate the increase in bandwidth availability in these regions.

Furthermore, researchers should develop solutions that are meaningful to people in these regions, rather than to people in developed contexts. Rajesh Balan wondered how to get in touch with nongovernmental organizations, and the panel suggested using social and professional networks. Suman Banerjee pointed out one of the major challenges was sustaining these projects beyond three or four years; solutions suggested were to

make the project open source, find a partner in the developing world, or build a startup with the goal of sustaining the project.

### Mobile Cloud Interactions

Lars Tiede presented the concept of a display cloud model, where users could select local and remote programmable displays and add them to a user-specific cloud display using their mobile phone. This model was more flexible than the Ubuntu phone, he said, because it lets the user connect to any display.

Jeffrey Bickford presented a method that allowed seamless synchronization of both execution and data across multiple mobile devices. A VM image was maintained per device, and developers specified policies for checkpointing. Bickford explained that the takeaway from the feasibility studies was that changes in execution state and data weren't as drastic as expected.

Michael Butkiewicz presented WebSieve, which reduced the load time of webpages on mobile devices by loading objects with higher priority and lazy loading the objects with lower priority. The challenges, he said, were to infer object utility, understand the object dependencies, and predict the new load times. The audience pointed out that the object utility value could be used for other Web optimization techniques as well.

The discussion then moved on to different applications for display clouds.

### Vehicular Networking and Transportation

Tan Zhang explored using TV-whitespace-based communication systems to provide connectivity to vehicles. Using both a forward- and rear-mounted antenna, the system identified the best channel parameter conditions in advance of the rear-antenna transmission.

Wenjie Sha proposed a social vehicle-navigation system that uses driver-provided information in the form of 20-second voice tweets to calculate

## CONFERENCES

personalized driving routes. Sha pointed out that the system could complement Google traffic. One audience member had concerns about malicious users using the system to verbally abuse other drivers in the vicinity. Sha talked about using filters and reputation systems to address this issue.

Blerim Cici presented how, using Call Description Records containing phone-call location information from a large Spanish telecom, his group derived an upper bound on the potential effectiveness of ride sharing in Madrid. Audience members were concerned about the practicality of ride sharing.


### DEMO SESSION

In addition to the talks, there were 13 posters and six demos presented during the joint poster and demo session on the evening of the first day of the workshop.

During the “Bringing In-Situ Social Awareness to Mobile Systems:

Conversational Turn Monitoring and its Applications” demo session, the authors presented the idea of harnessing the omnipresence of smartphones to launch a new class of in-situ, socially aware apps. All displayed apps were visually geared toward an approachable, game and cartoon-like aesthetic. This demo won the best poster/demo award.

Sanorita Dey and Nirupam Roy presented a technique, SensorPrint, to fingerprint mobile phones using hardware sensors, such as accelerometers. Sensors exhibit diverse behaviors that vary with phones and models. An initial experiment confirmed the hypothesis, and a preliminary evaluation showed that a smartphone could be identified with an accuracy of 96 percent. An audience member pointed out that a malicious app or webpage could access the accelerometer and use this technique adversely to identify the phone. This poster won the second prize for the best poster/demo.

**W**e would like to congratulate the organizing committee, chaired by Sharad Agarwal, Alexander Varshavsky, Nilanjan Banerjee, Souvik Sen, Andreas Bulling, and Christos Efstratiou, for making this year’s Hot-Mobile such an enjoyable and successful event. 

**Aarathi Prasad** is a PhD student at Dartmouth. Contact her at [aarathi@cs.dartmouth.edu](mailto:aarathi@cs.dartmouth.edu)



**Michael Butkiewicz** is a PhD student at the University of California, Riverside. Contact him at [butkiewm@cs.ucr.edu](mailto:butkiewm@cs.ucr.edu).



## CONFERENCES *in the Palm of Your Hand*

**IEEE Computer Society's Conference Publishing Services (CPS)** is now offering conference program mobile apps! Let your attendees have their conference schedule, conference information, and paper listings in the palm of their hands.



The conference program mobile app works for **Android** devices, **iPhone**, **iPad**, and the **Kindle Fire**.

For more information please contact [cps@computer.org](mailto:cps@computer.org)

