# HotMobile 2014

**Morgan Vigil and Paul Schmitt,** *University of California, Santa Barbara*

The 15th Workshop on Mobile Computing Systems and Applications (ACM HotMobile 2014) was held February 26th and 27th in Santa Barbara, California. The program included 22 papers, 19 posters, and seven demos, and had more than 90 attendees participating in the two days of presentations and lively discussion.

## KEYNOTE: THE SELF-DRIVING CAR

The workshop began with Andrew Chatham, principle software engineer at Google, delivering a keynote that focused on Google's self-driving car project. Chatham described the motivation behind the project, arguing that human error and slow reaction times cause the vast majority of accidents and traffic delays. It thus stands to reason that minimizing the role of humans should improve traffic safety and efficiency.

The talk focused on the sensors and technologies used in self-driving cars and the interplay between them and the real-world transportation infrastructure. According to Chatham, prior work used two general models for enabling such technology: smart cars and smart worlds. Both approaches are somewhat lacking, however, because real-time decision making in smart cars results in the inability to travel at reasonably high speeds, and altering the existing infrastructure is prohibitively expensive and difficult to change. The proposed solution from Google is a hybrid approach: real-time decisions by cars, augmented with an idea of what the car will encounter next in the form of digital maps.

Chatham described the challenge of building logical representations of the transportation infrastructure. In order to map the world, Google researchers first drive while collecting 3D data from the laser to obtain a pixelized view of the world. Using this data, as well as the real-time reading from the laser, the cars can reason about the environment.

The future challenges for self-driving cars include the degradation of sensors, driving in adverse conditions (such as snow), and reacting properly when others on the road don't act as expected. However, Chatham believes that software and hardware solutions will address these issues.

## CROWDSOURCING

The first session focused on mobile crowdsourcing applications and the optimization of user data collection. Noah Klugman (University of Michigan) presented "Grid Watch: Mapping Blackouts with Smart Phones," a crowd-sourcing system that aims to use smartphones to sense power outages rather than rely on expensive smart meters deployed by power companies. Grid Watch is a background process on Android or iOS phones that checks the accelerometer and the microphone when it detects that the phone has been removed from a power source. Leveraging the fact that phones are typically moved when unplugged by users and that the phone microphone can detect the "hum" of AC mains, the system determines whether the power loss is from normal use or a power outage.

The two papers that followed discussed approaches and lessons learned from mobile data crowdsourcing. In the former, the researchers created a crowdsourcing platform that falls between locally administered and app store-based methodologies. Evan Welbourne (Samsung Research America) showed that their methodology is feasible and provided several lessons learned from a user study. The latter paper explored different incentive mechanisms commonly used in crowdsourcing—micropayments and weighted lotteries. John Rula (Northwestern University) explained that user compliance and effort is markedly different depending on the type of incentive a user chooses. Using such knowledge, researchers looking for a particular type of participation can better choose which type of incentive to offer. Alternatively, those looking to cover the broadest range of users should explore multiple incentives.

## SMARTPHONE SENSORS

The second session included papers that focused on pushing the boundaries of what we can do with smartphone sensors. Romit Roy Choudhury (University of Illinois) presented "Injecting Life into Toys," a paper looking at what a smartphone placed inside a child's toy could potentially offer. The prototype system takes the sensor stream and attempts to learn and classify gestures from the child. A system such as this

could aid in passive recognition of early development in a natural environment.

Anandatirtha Nandugudi (University of Buffalo) presented "The Mote is Dead. Long Live the Discarded Smartphone!," where he proposed using discarded smartphones as sensing platforms rather than motes. The constant demand by consumers for higher-powered mobile devices leads to millions of discarded devices that are increasingly powerful and include a rich set of sensors. The paper includes a preliminary study and finds many cases where discarded smartphones can replace motes for sensing applications. Nandugudi put forth that there is significant opportunity for leveraging these devices to sense the world.

Mingming Fan (University of North Carolina, Charlotte) presented the last paper of the session, "HiFi: Hide and Find Digital Contents Associated with Physical Objects via Coded Light." The system uses coded light and a small sensor attached to a smartphone to localize the device and to store and retrieve information on arbitrary physical objects. Fan showed that coded light can lead to fast, accurate localization without introducing visual distraction.

## AUTHENTICATION

The next session was kicked off by Landon Cox (Duke University) presenting, "VeriUI: Attested Login for Mobile Devices." Cox put forth the concept of "attested login," where user credentials coming from a mobile device are augmented with a certificate that includes information about the device's hardware and software. There is a critical tension between usability and security, with passwords representing usable but less secure solutions, and two-factor authentication being more robust but less usable. VeriUI attempts to solve this by only allowing access to services from trusted environments. Cox argued for a hardware-isolated environment for credentials and user intent. A system prototype was tested and shown to work with existing third-party applications,

and the system was more responsive than thin-client approaches.

Napa Sae-Bae (Polytechnic Institute of NYU) next presented "Hand Authentication on Multi-Touch Tablets," where the researchers explored biometric authentication using hand geometrics on multitouch tablets. The system uses interactive "challenges" that users complete to gather hand geometry measurements. The system improves on prior work in that no memorized credentials are required and there is reduced vulnerability to replay attacks. Limited user testing was done and the system showed promising results. Sae-Bae demonstrated that the system functions on existing hardware. Future work will include lowering the number of challenges and a larger user test.

The last two papers presented in the session focused on implicit authentication. Tao Feng (University of Houston) discussed "TIPS: Context-Aware Implicit User Identification Using Touch Screen in Uncontrolled Environments," where user gestures on mobile devices can identify different users. The work first explored gestures gathered in an uncontrolled environment and found significant differences between users. Feng said that measuring gestures within individual applications can increase system accuracy. TIPS was implemented on Android, and a user study was conducted across multiple hardware platforms. The system achieved 90 percent accuracy during the study.

The final paper of the session, "Towards Application-Centric Implicit Authentication on Smartphones," was presented by Urs Hengartner (University of Waterloo). Hengartner also spoke of application-centric implicit authentication and challenged device-centric authentication. He argued that application developers know best when and how to include authentication in their software. While such a change will increase the burden on application developers, it could provide smarter,

more fine-grained security. Hengartner also announced that the burden could be eased somewhat by using a shared library, and that his research group has created the openly available library "Itus."

## MOBILE SENSING

This session focused on unobtrusive, mobile sensing. Seungwoo Kang (KAIST) began the session with a presentation of "Sinabro: Opportunistic and Unobtrusive Mobile Electrocardiogram Monitoring System." Sinabro attempts continuous monitoring of heart rate via electrodes formed around a mobile phone, rather than featured as part of a separate wearable device. The work observes that a significant number of opportunities exist throughout the day where users' heart rate can be monitored based on their hand positioning around their mobile phone. However, not all of these opportunities are equal—accurate measurement requires both hands to make skin contact with the device for 10 to 30 seconds. Future work seeks to normalize heart rate data based on user context.

Kartik Muralidharan (Singapore Management University) presented "Barometric Phone Sensors—More Hype Than Hope!" Muralidharan discussed the limitations of using barometric sensors on smartphones for determining location in indoor settings. Indoor environments provide a host of challenges for smartphone barometers, including climate control and diverse floor heights. While the work demonstrates that barometric sensors cannot accurately identify location on their own, they are more successful when coupled with other sensors. Likewise, the paper demonstrates a higher rate of accuracy identifying location based on relative barometric changes rather than absolute levels.

The third paper of the session, "QuiltView: A Crowd-Sourced Video Response System," presented by Zhuo Chen (Carnegie Mellon University), was also awarded "best demo" at the

conference. The QuiltView crowd-source system leverages wearable devices, such as Google Glass, to answer scenario-specific queries with video clips. The advantage of this approach is that it lets people respond to queries with information-rich and contextual responses with minimal interruption. Discussion of the system focused on privacy concerns for people who appear in recorded query responses.

The final paper presented lessons learned in developing and deploying a mobile experimentation lab. Rajesh Krishna Balan (Singapore Management University) presented "LiveLabs: Building an In-Situ Real-Time Mobile Experimentation Lab." LiveLabs seeks to provide researchers with a live mobile testbed for large-scale deployment and evaluation while minimizing the overhead of initiating and maintaining such a testbed. Balan spent a considerable amount of time discussing challenges with this type of system, including administration, user retention, and the power cost of continuous sensing. As the system develops, future work seeks to grant researchers more finite control over experiments, including the length of duration and process branches based on feedback from results.

## PRIVACY

The privacy session included two papers that highlighted vulnerabilities to personal information on mobile devices. The first paper, "Information Leakage through Mobile Analytics Services," was presented by Terence Chen (University of New South Wales). Chen showed that personal profile information as well as application usage can be extracted from two major analytics services (Google Mobile App Analytics and Flurry). Attackers can spoof user device IDs to interact with the analytics services, and then extract the target's personal information from the services. Chen also showed that the researchers could alter the advertisements delivered to a given user.

The second paper, "Unveiling the Hidden Dangers of Public IP Addresses in 4G/LTE Cellular Data Networks," was presented by Wai Kay Long (National University of Singapore). The paper found that ISPs often hand out public IP addresses to mobile devices by default or with a simple change to the APN by the user. The consequences of having a device that is directly addressable on the Internet include multiple attacks against the devices. For example, data plan quotas can quickly be drained by unwanted traffic directed at the target. Also, denial-of-service attacks can simply overwhelm the victim with traffic from the Internet. Lastly, because the radio is a major energy consumer in a mobile device, sending periodic traffic to the device will cause the radio to remain in a high power state, thus quickly draining the battery. Long also discussed strategies for defending against such attacks using proxies, firewalls, and secondary IP addresses.

## NETWORKING

The network session spanned topics from software-defined radios to improved user experiences. Yongtae Park (Korea University) presented "Software Radio on Smartphones: Feasible?" Park's work evaluated current and projected smartphone architectures to determine the feasibility of supporting software-defined radios. Although current smartphone processing power is sufficient to support software-defined radios, limitations imposed by the real-time requirements of MAC/PHY protocols impose significant challenges. Park suggests that these challenges are due to variance in architecture and delays caused by SIMD instruction sets. Application-specific MAC/PHY protocols and architecture aware protocols can decouple software defined radios from the OS providers and help overcome architectural limitations and diversity.

Jeff Pang (AT&T) presented "Prometheus: Toward Quality-of-Experience Estimation for Mobile Apps from Passive Network Measurements," on behalf of Vaneet Aggarwal (AT&T). Cellular carriers have limited vantage and control of over-the-top (OTT) applications; thus it's difficult to allocate the network resources required to provide users of OTT applications with a high quality of experience (QoE). Using machine learning techniques to correlate QoE with network flow records, Prometheus provides transparency that lets carriers optimize configuration and provisioning of cellular networks for user applications.

The final paper in the networking session described Delphi, a transport-layer module for mobile network selection. Shuo Deng (MIT CSAIL) presented "All Your Network Are Belong To Us: A Transport Framework for Mobile Network Selection." Delphi combines network conditions, application requirements, network history, and device limitations to select an appropriate network. Additionally, Delphi integrates per-node and system-wide perspectives on the network to optimize network selection on an individual or utilitarian basis.

## ENERGY AND PERFORMANCE

The final workshop session included three papers. Wook Song (Seoul National University) kicked things off with "Reducing Energy Consumption of Smartphones Using User-Perceived Response Time Analysis," where the authors aim to reduce power consumption without negatively affecting the user experience. They explore user interactions and delineate subintervals between user-perceived and user-oblivious response times. Given this information, Song said that they could lower the CPU frequency once the user-perceived interval completes—thus preserving the user experience as well as offering energy savings.

Ashiwan Sivakumar presented "Cloud Is Not a Silver Bullet: A Case Study of Cloud-Based Mobile Browsing," which explores power consumption when using a cloud-based browser versus a

traditional device-based browser. Siva-kumar found that the prevailing thought of moving everything into the cloud isn't necessarily beneficial in all cases. Particularly, moving JavaScript into the cloud can prove detrimental when the pages include periodic updates or long-running JavaScript code. In those cases, a single large transfer would be more energy efficient. Sivakumar also showed that data compaction isn't always ben-eficial, because the time required makes the network hardware consuming more energy. There are several tradeoffs when considering computation placement for mobile devices.

Tarun Bansal presented the last paper; "RobinHood: Sharing the Happiness in a Wireless Jungle." The designed system leverages high-density wireless access points and a high-speed wired network to let multiple clients transmit simultaneously. Bansal explained that APs can use the high-speed backbone network to communicate with other APs to decode interfered packets by subtracting known packet samples, exchanged over the backbone, from interfered packet samples at the AP. The simulated system showed improve-ment over both omniscient TDMA and traditional 802.11.

A few themes were prevalent throughout. In the sensing-ori-ented presentations, group discussions questioned the appropriateness of using smartphones as sensors and the tradeoff between utility and relevance. Most pervasive was the discussion of scale, and how many of the challenges associ-ated with the accuracy of mobile sys-tems might disappear with significant scale. Also discussed was the tension between research contributing to a spe-cific application domain and research contributing to the field of computer sci-ence. Related was the challenge asserted by conference organizer, Matt Welsh (Google), to move the mobile research agenda away from the limitations of current mobile systems and look further into the future to when such limitations would no longer be as pertinent.

**Morgan Vigil** is a PhD student in the Computer Science Department at the University of California, Santa Barbara. Contact her at mvigil@cs.ucsb.edu.

**Paul Schmitt** is a PhD student in the Computer Science Department at the University of California, Santa Barbara. Contact him at pschmitt@cs.ucsb.edu.

cn Selected CS articles and columns are also available for free at http://ComputingNow.computer.org.