Editor: Jason Hong ■ Carnegie Mellon University ■ jasonh@cs.cmu.edu

HotMobile 2012

Igor Pernek, Mostafa Abdulla Uddin, and Jack Fernando Bravo-Torres

The 13th International Workshop on Mobile Computing Systems and Applications (HotMobile 2012) took place late February in San Diego. The workshop, sponsored by ACM Sigmobile, brought together 87 researchers from Asia, Europe, and North and South America. This year, only 14 papers were accepted out of 69 submissions, making the workshop's acceptance rate the lowest in its 13-year history.

The two-day workshop started with a keynote address by Robert Gilmore, Qualcomm's vice president of engineering, followed by six paper sessions, a demo and poster session, and a doctoral consortium session, where doctoral students could get feedback on their dissertation research. Also, for the first time, the workshop held a short "primer" before the demo and poster session, so authors could introduce their work before the actual session.

THE FUTURE OF MOBILE EXPERIENCES

In his keynote, "The Future of Mobile Experiences," Gilmore pointed out that the mobile platform is already the biggest platform in the history of mankind, and it's expected to grow significantly during the next decade. To accommodate such growth, parallel improvements in networks and processors will be crucial. Although future processors will be able to achieve higher frequencies with lower energy consumption, we must improve the capacity of future networks through more cells, a broader frequency spectrum, and network-traffic optimization.

Gilmore also highlighted some of the technologies Qualcomm has developed, such as Snapdragon, an all-in-one mobile processor that enhances mobile devices' performance and battery life. He also shared his vision of femtocell technology, discussing how it will change the future of mobile computing and communication.

The mobile platform is already the biggest platform in the history of mankind, and it's expected to grow significantly during the next decade.

MOBILE ADVERTISING

The first session focused on ad-supported mobile applications. "Mitigating the True Cost of Advertisement-Supported 'Free' Mobile Applications," presented by Khan Azeem from Singapore Management University, proposed a new ad-delivery framework called Cameo. The authors showed that the majority of network traffic in free mobile apps is caused by ads, which must be served in real time based on the current context. This can stress the capacity of wireless networks and be costly for users on metered data plans. To mitigate the stress from such non-essential traffic, the Cameo framework uses predictive profiling of a user's context to anticipate ads to be served. It then modulates the delivery mechanism to enable effective mobile advertising at considerably lower costs.

Ilias Leontiadis, from the University of Cambridge, presented the second paper, "Don't Kill My Ads! Balancing Privacy in an Ad-Supported Mobile Application Market." In the paper, the authors examine the problem of privacy protection for mobile phones. By analyzing the current business model for mobile application ads, the authors showed that wide adoption of more rigorous privacy-protection mechanisms could lead to the collapse of the ad-driven mobile application market. To address this, they designed a market-aware privacy-protection framework that aims to balance the developer's revenue and user's privacy. The proposed framework is based on establishing a feedback control loop that adjusts the level of privacy protection on mobile phones in response to ad-generated revenue.

INFRASTRUCTURE AND TOOLS SUPPORT

Two papers were presented on infrastructure and tools support. John Rula, from Northwestern University, presented, "Crowd (Soft) Control: Moving Beyond Opportunistic Sensing," which proposes an alternative approach to attaining coverage without scale by controlling the movement of participants. The approach can help control the temporal and spatial movements of mobile users by leveraging the built-in incentives of location-based gaming and social applications. By pairing

Published by the IEEE CS ■ 1536-1268/12/$31.00 © 2012 IEEE

network services with these location-based apps, Crowd Soft Control lets researchers use an application's incentives to control the movement of participating users, increasing the effectiveness and efficiency of the associated network service.

The second paper, "Code in the Air: Simplifying Sensing and Coordination Task on Smartphones," was presented by Lennin Ravindranath from MIT's Computer Science and Artificial Intelligence Laboratory. The Code in the Air (CITA) system eases the development and running of tasking applications on smartphones. It provides a task execution framework to automatically distribute and coordinate tasks; energy-efficient modules to infer and compose user activities; and a push communication service for mobile devices that overcomes some shortcomings in existing push services.

VISION: A LOOK INTO THE FUTURE

Two papers described how we can use mobile technology to address challenges of the future. "An Amulet for Trustworthy Wearable mHealth," presented by David Kotz from Dartmouth College, described an easy-to-use mHealth architecture that provides strong security and privacy guarantees. The authors propose an omnipresent and trustworthy wrist-worn platform that can authenticate its wearer and establish secure communication channels between a mobile gateway and different wearable sensors. However, several attendees questioned the appropriateness of placing the device on the wrist. Because watches are usually worn as a fashion statement, the wrist placement might not be discrete and thus could make people with health deficits easier to recognize. The authors responded that the wrist-watch design was picked specifically because wearing a watch is socially acceptable. They also countered that that benefits of such a device would outweigh its disadvantages.

Nigel Davies from Lancester University presented the next paper, "Sixth Sense Transport: Challenges in Supporting Flexible Time Travel," which focused on designing a mobile platform that helps travelers enjoy a more flexible schedule when travelling. Davies presented challenges in trying to develop a system that provides mobile users insight into future mobility patterns and supports new forms of opportunistic travel sharing. He described three scenarios in which the system proposed could be of use—in tourism, daily travel activities, and logistics. Most comments focused on the problem of system adoption. The main suggestion was to study the system in a particular domain and tackle a specific problem, such as missed package deliveries in the logistics domain.

A novel scheme for efficient long-range neighbor discovery leverages the clustering of nodes and the radio heterogeneity of mobile devices.

MOBILE APPLICATIONS

This session focused on innovative mobile applications that use smartphone sensors. "WalkSafe: A Pedestrian Safety App for Mobile Phone Users Who Walk and Talk While Crossing Roads," presented by Tianyu Wang from Dartmouth College, presents the design, implementation, and evaluation of the WalkSafe application. WalkSafe captures the video stream from the backside camera of the mobile phone and uses image recognition algorithms to alert the users about approaching vehicles. To save energy, the vehicle detection process is triggered only during an ongoing phone call. Wang also presented a demo video of the application during the presentation. After the presentation, interesting questions were raised about the approach's feasibility, but overall, the idea was well-received.

Seongwon Han from UCLA presented, "EyeGuardian: A Framework of Eye Tracking and Blink Detection for Mobile Device Users." Due to lower resolutions and smaller screen sizes, using mobile devices requires more user concentration. This can lead to a lower eye-blinking rate, which can cause Computer Vision Syndrome (CVS). To alleviate the CVS risk, EyeGuardian provides a mobile framework for eye tracking and blink detection that can detect a low blink rate and provide feedback and recommendations. Han explained that conventional methods for detecting blink rates are inappropriate for mobile devices, because the location of the eyes is harder to predict due to higher movement dynamics when using a mobile device. Using accelerometers built into mobile devices could be an effective method for eye location prediction.

After the presentation, Mahadev Satyanarayanan of Carnegie Mellon University pointed out some possible directions of mobile applications research and opened up a discussion on how we might use mobile devices to improve human behavior.

SECURITY AND NETWORKING

Emmanuel Owusu of Carnegie Mellon University opened this session with "ACCessory: Password Inference using Accelerometers on Smartphones." He showed that accelerometer readings can leak sequences of text entered on a smartphone touchscreen keyboard and proposed a set of counter measures that could be taken to defend against accelerometer-based side-channel attacks. He argued that this kind of information leakage is especially risky, because accelerometers are perceived as benign and don't require special privileges on current smartphone operating systems.

In the next talk, Eyal Zohar from Technion presented, "Celleration: Loss-Resilient Traffic Redundancy Elimination for Cellular Data."

CONFERENCES

The paper proposes a novel gateway-to-mobile Traffic Redundancy Elimination system, specifically tailored for the new generation of data-intensive cellular networks. The authors conducted extensive experiments to show that their approach, which eliminates individual clients' redundant traffic across the cellular network, achieves significant bandwidth savings compared to existing solutions.

The last paper of the session, presented by Robin Kravets from the University of Illinois at Urbana-Champaign, was "United We Find: Enabling Mobile Devices to Cooperate for Efficient Neighbor Discovery." The authors proposed a novel scheme for efficient long-range neighbor discovery that leverages the clustering of nodes as well as the radio heterogeneity of mobile devices to reduce power consumption. The basic idea of their approach is that coordination over a low-power, short-range radio, such as Bluetooth or Zigbee, can help clustered nodes to distribute the load of high-power, long-range scanning.

At the end of the session, Rajesh Krishna Balan of Singapore Management University pointed out that mobile security mechanisms are becoming increasingly more obtrusive with features similar to random smartphone vibration during keyboard input, which was proposed to counter the problem of leaking key presses through accelerometers. He started a discussion on how far the community is willing to go with similar security measures and possible alternative approaches.

INDOOR LOCALIZATION

This last session focused on addressing the challenges of indoor localization. Souvik Sen from Duke University presented, "SpinLoc: Spin Once to Know Your Location." He discussed the feasibility of using wireless signal attenuation between an access point (AP) and mobile device to reveal the AP's direction. The basic idea is to measure the *energy of the direct path* (EDP) at different rotational angles. Then, each AP's direction is estimated by exploiting the fact that the EDP between an AP and the mobile device is always lower when there's a body in between the two. After determining the direction of multiple APs, triangulation can be used to define the user's precise location. The authors illustrated the validity of this approach for different scenarios.

In the next talk, Jie Xiong of University College London presented, "Towards Fine-grained Radio-Based Indoor Location." He mentioned that there's a clear trend of using an increasing number of APs and multiple antennas in each AP to meet the increasing demands for wireless bandwidth. This trend motivated the authors to explore the indoor localization problem from a new perspective. Their ArrayTrack indoor localization system uses the angle-of-arrival technique to determine the client's location. The system's median accuracy is approximately 25 cm for stationary clients in indoor environments.

Mobile security mechanisms are becoming increasingly more obtrusive with features similar to random smartphone vibration during keyboard input.

Zheng Sun of Carnegie Mellon University presented the last paper of this session. In "Polaris: Getting Accurate Indoor Orientation for Mobile Devices Using Ubiquitous Visual Patterns on Ceilings," Sun argued that compass readings of smartphones are highly error prone in indoor environments because of metallic objects. The proposed Polaris system can capture a picture of the ceiling to determine the user's orientation with respect to the building. The system assumes that the buildings are rectangular, and uses Google Maps to acquire the buildings' orientation.

In the discussion session, the question of whether indoor localization is still an open problem was raised. Some attendees viewed this as a closed problem, but others argued that existing indoor localization solutions aren't scalable to larger areas so there's still room for improvement.

DOCTORAL CONSORTIUM

This year's doctoral consortium was organized by Eyal de Lara of the University of Toronto. It provided a platform for us, as PhD students, to present our dissertation research and receive feedback.

Igor Pernek presented his research on exploiting off-the-shelf smartphones to promote a healthier lifestyle through strength training. His work defines requirements for wearable exercising assistants; presents a feasibility study of different sensing modalities and sensor set-ups for capturing strength-training exercises; and researches activity- and pattern-recognition algorithms to infer high-level strength-training features such as the type, number, and correctness of performed exercises. He's also developing a strength-training data-collection and analysis platform to allow opportunistic use of smartphones for capturing strength-training data. The platform will provide different functionalities to stakeholders, such as those doing the exercises as well as personal trainers and researchers.

Jack Fernando Bravo-Torres presented, "Supporting More Efficient Communications in Vehicular Ad Hoc Networks with New Constructs Based on a Virtualization Layer." This work aims to study the potential advantages of having virtual mobile nodes (VMNs) that move within vehicular networks to plan, design, and deploy new communications services. Also, to minimize the participation

of vehicles in many decision-making processes and reduce overhead, the author proposes two new constructs: virtual virtual circuits (VVCs) and virtualized circuits (VLCs). Mimicking the use of traditional virtual circuits in datagram networks, a VVC is a route over a sequence of VMNs for delivering packets between two communicating endpoints. On the other hand, a VLC is a predetermined route over a sequence of VMNs that can be used to deliver the packets of one or several communications. For this, the definition of VLCs will be driven by traffic models that capture the major flows of traffic in a city, the peak traffic hours, and so on. Moreover, the author aims to integrate additional information sources, such as transport-sharing policies, as well as advances in automated transportation, fleet management, and real-time traffic information systems.

Mostafa Uddin presented his ongoing research work, "Audio Wi-Fi: Augmented Wi-Fi Network with Audio Channel for Smart Devices," supervised by Tamer Nadeem (Old Dominion University). This work focuses on using audio as an additional control channel with Wi-Fi to enhance Wi-Fi network performance. Wi-Fi still faces some perturbation in terms of channel use, fairness issues, and energy consumption. Uddin discussed several challenges of using Wi-Fi as an augmented channel with audio channel and solutions. Gaetano Borriello from the University of Washington suggested using the positive characteristics of the acoustic channel rather than focusing on the negative characteristics of the audio.

Mobile devices are rapidly improving their communication interfaces, processing power, and sensing capabilities.

Mobile devices are rapidly improving their communication interfaces, processing power, and sensing capabilities, opening up the potential for new applications and attracting the interest of researchers from different communities. The HotMobile workshop is a perfect venue for presenting ground-breaking and out-of-the-box ideas that will help shape the future of the mobile community.

Igor Pernek is a PhD student at the University of Maribor. His research interests include mobile sensing, ubiquitous computing and applying machine learning techniques to support sport activities. Contact him at igor.pernek@uni-mb.si.

Mostafa Abdulla Uddin is a second-year PhD student in the Computer Science Department at Old Dominion University. His research interests include wireless networking and mobile computing research. Currently, he's focusing on developing audio channel assistive wireless networking for smart devices. Contact him at muddin@cs.odu.edu.

Jack Fernando Bravo Torres is pursuing a PhD in the Department of Telematics Engineering at the University of Vigo, Spain. He's also an assistant professor at the Salesian Polytechnic University, Ecuador. His research interests include the design and development of innovative services for vehicular ad hoc networks. Contact him at jbravo@ups.edu.ec.