



MobiSys 2013

Mariya Zheleva and Michael Nekrasov

The 11th International Conference on Mobile Systems, Applications, and Services (MobiSys 2013) was held June 25th through the 28th in Taipei, Taiwan. Hosted in Asia for the first time, the event attracted over 200 attendees, with 57 percent of the participants coming from Asian institutions (33 percent were from North America and 7 percent were from Europe). This three-day event featured 33 papers, 18 posters, 15 demos, and a first-time video track with 14 videos.

After 10 successful years of showcasing cutting-edge academic and industrial research in mobile technology, this year's MobiSys introduced a new panel, "The Untold Stories of Research." The panel facilitated a lively discussion on topics related to the life of research projects after they appear in an academic publication. In addition, MobiSys continued to bridge industry and research with its Industry Panel. This year's panel featured two Taiwan-based world leaders in the field of mobile communications—HTC and MediaTek—and provided insight into Taiwan's perspective on the future of mobile communications.

KEYNOTE: NATURAL USER INTERFACES

Patrick Baudisch, from the Hasso Plattner Institute, started the conference with an inspiring keynote reflecting on how the evolution of mobile technology is coupled with that of human-technology interfaces. Baudisch talked about natural user interfaces (NUIs) as the new era in hardware interfaces. NUIs won't require training—a

"one-year old [could] navigate these interfaces," said Baudisch.

The talk outlined how command-line interfaces (CLI) evolved into GUIs and how emerging touch technology has resulted in the leap to NUIs. According to Baudisch, "The main problem of CLI and GUI is the massive asymmetry between input and output." He explained that while a computer projects gigabytes of audio

Future interfaces will use visual as well as touch data to let users provide greater dimensions of input to the system.

and visual data to the user, the input is often limited to a singular (X, Y) coordinate of a mouse. In contrast, future interfaces will use visual as well as touch data to let users provide greater dimensions of input to the system, so they can interact with the system in a more natural way. Baudisch illustrated how touch leads to effective NUIs on mobile devices, because interactions between physical and digital occur in "a common Euclidian Space," where physical fingers directly manipulate digital objects.

Baudisch concluded his talk by exploring how the evolution of NUIs can empower remote and developing areas by reducing the prior knowledge required to operate advanced technology.

FROM SAFER VEHICLES TO SMARTER ADS

The first set of conference sessions covered a wide range of topics, from vehicular safety to privacy in mobile computing to more customized ads.

Vehicular Systems and Applications

The first session showcased papers on a variety of vehicle-centric topics, ranging from user interface design for vehicular mobile phone applications to detecting unsafe behavior on the road. All projects used mobile phones as behavior sensors. Kyungmin Lee presented "AMC: Verifying User Interface Properties for Vehicular Applications," discussing an approach to automated interface verification to determine whether mobile phone applications are a distraction to drivers.

"CarSafe App: Alerting Drowsy and Distracted Drivers Using Dual Cameras on Smartphones," described an effort to use front- and back-facing phone cameras to monitor road conditions and driver behavior. The main challenge was scheduling the two cameras, because current mobile devices only allow one camera to be active at a time.

Next, "CrowdAtlas: Self-Updating Maps for Cloud and Personal Use" talked about leveraging information from moving vehicles, such as taxis, to provide automated high-accuracy street mapping. The discussion focused on countries where street laws are often ignored and how crowd mapping might be an effective way to highlight differences between road policies and actual human use of the roads. Yingying

Chen presented the final paper in this session, “Sensing Vehicle Dynamics for Determining Driver Phone Use.” This work was concerned with determining a phone’s location in a car using the phone’s centripetal acceleration and using a second sensor in the car to apply phone-usage policies specific to the location of the phone relative to the vehicle’s driver.

Energy, Privacy, and Security

This session opened with “Optimizing Background Email Sync on Smartphones,” which identified factors that contribute to high power consumption in email synchronization on mobile phones and suggests ways to reduce such energy consumption.

“Energy Characterization and Optimization of Image Sensing toward Continuous Mobile Vision,” presented by Rice University’s Robert LiKamWa and awarded Best Paper, focused on optimizing applications that are aware of user’s actions via continuous mobile vision. The major limitation of continuous vision is the high energy consumption of commercial image sensors. This work showed how increasing the clock speed of image sensors, in conjunction with aggressive standby scheduling, can drastically reduce energy use. The discussion that followed showed excitement about the potential effect of this technology, with attendees asking questions such as whether this functionality would be accessible to developers through APIs. LiKamWa replied that such APIs should be available and that their implementation is feasible.

While the first two talks in this session were related to energy efficiency, the last two were concerned with security and privacy. First, Abhinav Parate presented “Leveraging Graphical Models to Improve Accuracy and Reduce Privacy Risks of Mobile Sensing,” a project focused on sensing an individual’s context in an energy-efficient and privacy-preserving way.

Finally, Yuvraj Agarwal presented “ProtectMyPrivacy: Detecting and

Mitigating Privacy Leaks on iOS Devices Using Crowdsourcing.” This work demonstrated an iOS application, currently used by more than 140,000 users, that offers fine-grained privacy policies on mobile devices. The app uses the crowd’s common decisions on privacy settings to automatically suggest policies for users. Victor Bahl expressed concerns that the user feedback might be biased, because the app requires users to jailbreak their phone, which caters to users who are more privacy conscious. Agarwal answered that both jailbroken and nonjailbroken users install from the same market of applications, so the feedback from the application can still benefit nonjailbroken users, who are perhaps less concerned about their privacy.

Advertisements and Search

The next session opened with “Smart-Ads: Bringing Contextual Ads to

The app uses the crowd’s common decisions on privacy settings to automatically suggest policies for users.

Mobile Apps,” a project that contextualized advertisements based on the application content a user is viewing.

The next project discussed was “CAMEO: A Middleware for Mobile Advertisement Delivery,” which focused on reducing resource consumption caused by mobile advertising by allowing offline ad caching. Using context prediction and “bartering” advertisement rights, CAMEO reduces resource consumption associated with advertisements.

The session concluded with “Scalable Crowd-Sourcing of Video from Mobile Devices,” a project that focused on using scalable techniques to process crowd-sourced video.

Poster and Demo Session

The first day came to a close with a session hosting 18 posters and 15 demos.

“Power Management Using Game State Detection on Android Smartphones,” by Benedikt Dietrich from TU Munich, received the Best Demo Award. He demonstrated an Android app that halves the power consumption from gaming using knowledge about game states.

Another compelling demo, “Visible Light Communications for Scooter Safety,” was presented by a group from National Taiwan University. Given that scooters are a popular form of transportation in Taiwan, the demo showed a cost-efficient approach for increasing driver safety using vehicle-to-vehicle communication. The approach varies the intensity of the scooter’s rear light to signify the driver’s actions.

One of us (Michael Nekrasov) received the Best Poster Award for “Instrumenting Thailand’s Coastline: Mobile Devices for Environmental and Disaster Monitoring.” The work explored how mobile computing could enable automated low-cost environmental and disaster monitoring in developing regions. The project deployed a mobile device in the Gulf of Thailand to detect changes in water quality that would affect the shell fish farmed by locals in the Gulf.

SOFTWARE, LOCATION, AND INTERFACE DESIGN

The next day of the conference started with a session on operating systems, software, and virtualization.

OS, Software, and Virtualization

This session started with “Just-in-Time Provisioning for Cyber Foraging.” The project used cloudlets to rapidly provision personal clouds on demand.

Next was “SIF: A Selective Instrumentation Framework for Mobile Applications,” which allowed automatic instrumentation of application binaries based on abstract rules.

Benjamin Davis from UC Davis gave an interesting presentation. His project, dubbed “RetroSkeleton: Retrofitting Android Apps,” provided a framework for retrofitting Android apps that



How to Reach Us

Writers

For detailed information on submitting articles, write for our Editorial Guidelines (pervasive@computer.org) or access www.computer.org/pervasive/author.htm.

Letters to the Editor

Send letters to

Brian Kirk, Lead Editor
IEEE Pervasive Computing
 10662 Los Vaqueros Circle
 Los Alamitos, CA 90720
pervasive@computer.org

Please provide an email address or daytime phone number with your letter.

On the Web

Access www.computer.org/pervasive for information about *IEEE Pervasive Computing*.

Subscription Change of Address

Send change-of-address requests for magazine subscriptions to address.change@ieee.org. Be sure to specify *IEEE Pervasive Computing*.

Membership Change of Address

Send change-of-address requests for the membership directory to directory.updates@computer.org.

Missing or Damaged Copies

If you are missing an issue or you received a damaged copy, contact membership@computer.org.

Reprints of Articles

For price information or to order reprints, send email to pervasive@computer.org or fax +1 714 821 4010.

Reprint Permission

To obtain permission to reprint an article, contact William Hagen, IEEE Copyrights and Trademarks Manager, at copyrights@ieee.org.

“empowers users to observe and control third party applications” without access to the source code. This tool would enable companies and power users to apply custom security policies and other modification to applications downloaded directly from an Android market.

The final presentation was on “SmartSynth: Synthesizing Smartphone Automation Scripts from Natural Language,” a project focused on smartphone automation scripts based on natural language descriptions.

Location, Indoors and Outdoors

The proposed location techniques leveraged a variety of frequency-modulation (FM), acoustic, GPS, and Wi-Fi signals. The session opened with “FM-Based Indoor Localization via Automatic Fingerprint DB Construction and Matching,” which leveraged FM signals for indoor localization without profiling ahead of time.

Up next was “High-Accuracy Differential Tracking of Low-Cost GPS Receivers,” by a group from Vanderbilt University. The work focused on localization based on relative location rather than absolute coordinates. The proposed system leveraged GPS to provide an order of magnitude improvement in localization accuracy in comparison with systems that use differences between absolute coordinates.

A paper by researchers at the University of Florida, called “Guoguo: Enabling Fine-grained Indoor Localization via Smartphone,” proposed localization based on acoustic signals. Also, “Avoiding Multipath to Revive Inbuilding WiFi Localization,” presented by Souvik Sen, leveraged physical layer information from Wi-Fi signals to perform indoor localization without the need of wardriving. The proposed system leveraged direct and reflected paths of Wi-Fi signals to estimate Channel State Information and perform localization.

Interface Design

The preceding session explored how audio and images can be used for

effective interface design on mobile platforms. The session opened with the “Spartacus: Spatially-Aware Interaction for Mobile Devices through Energy-Efficient Audio Sensing,” from researchers at Carnegie Mellon University. This work explored acoustic interaction between two devices to identify pointing gestures for initiating phone-to-phone interaction. Their work sparked a suggestion from the audience regarding the benefit of adding infrared sensors that were once common on PDAs to mobile devices.

“View it Right,” presented by Guobin Shen from Microsoft Research Asia, explored the visual distortion that users experience when looking at their mobile devices from slanted view angles—when placed on a desk, for example. This team’s approach used a wide-angle front-facing camera to detect the user’s face and correct the distortion.

The third talk of the session, “Screen-Pass: Secure Password Entry on Touchscreen Devices,” presented by Landon Cox from Duke University, presented a way to identify a verified keyboard on the screen using image analysis.

The session closed with a video shown by Jeng-Shiou Leu of National Taiwan University of Science and Technology. Leu, building on the Technology, Entertainment, Design (TED) talk on sixth sense technology from MIT, demonstrated a combination of portable projector (that displayed a GUI on nearby surfaces) and gestures to control household appliances.

The Untold Story

The second day concluded with a panel and audience discussion on the continued effect of research after publication. Jane Liu, of Academemica Sinica, presented a career’s worth of academic breakthroughs, comparing ideas that affected industry to those that were ignored by industry. She highlighted how even feasible technologies might not be scalable to mass market implementation.

Ashu Sabharwal, from Rice University, challenged professors to step down from the managerial role to “be part of the team,” working with graduate students on ideas. Lin Zhong, of Rice University, argued that implementation “ultimately should be driven by intellectual curiosity.”

Mahadev Satyanarayanan, from Carnegie Mellon University, shared his experience transforming research into commercial products. He believes that “there is no easy [fiscal] quarter” for companies, so they’re unlikely to take up external ideas because their resources are already invested in their own projects. Researchers thus must transform ideas into commercial products. This session sparked heated debate, with some, including Ashu Sabharwal, believing in a hybrid model bridging academia and industry, while others, such as Victor Bahl, claiming researchers can’t be in both worlds at the same time and should choose between academia and industry.

INDUSTRIAL PERSPECTIVES AND ASSORTED TOPICS

The last day of the conference started with the Industry Panel and ended with a session on “assorted topics” in mobile computing.

Industry Panel

This year, MobiSys welcomed Yucheun Kevin Jou, the chief technology officer at MediaTek, and Edward Chang, vice president of mobile device manufacturer HTC. Jou discussed the evolution of mobile applications and the hardware that runs them. He outlined the progress in multicore architecture and move toward heterogeneous computing that shares workload and memory space between the graphic and application processors on a common memory space for energy savings and performance optimization.

Chang outlined industry interests in activity classification from motion sensors and interest in Wi-Fi-based

indoor localization at both HTC and Google. Dialogue with the audience revealed industry focus toward microcontrollers for low-energy sensing, as well as interest in new forms of sensing, such as biosensing, for use in the developed world and developing regions.

Cellular and Wi-Fi

The cellular session addressed improved performance and accountability of wireless communication. The first presentation, “Accounting for Roaming Users on Mobile Data Access: Issues and Root Causes,” from UCLA researchers, addressed the fiscal accounting nature of cellular networks. This work showed how areas of no signal, weak signal, and handoffs between coverage zones and network types can cause over-accounting of total transmitted data by cellular providers, especially for applications using the User Datagram Protocol.

Dialogue with the audience revealed industry focus toward microcontrollers for low-energy sensing, as well as interest in new forms of sensing.

The second presentation, on “Comparison of Caching Strategies in Modern Cellular Backhaul Networks,” outlined a custom built monitoring system deployed on a Korean Cellular network for a week. The goal of this work was to propose effective caching techniques after careful traffic analysis. Although the proposed caching scheme was effective in reducing packet redundancy, members of the audience discussed how the push to encrypted communication protocols, such as HTTPS, might in the future prohibit inspecting the content of individual packets.

Jun Huang from Michigan State University presented “LEAD: Leveraging Protocol Signatures for Improving Wireless Link Performance,” which predicted bits based on previously sent packets to increase throughput in weak signal scenarios resulting in high packet loss.

The final talk in this session, “PROTEUS: Network Performance Forecast for Real-Time, Interactive Mobile Applications” presented by Qiang Xu from the University of Michigan, investigated how applications can make choices about communication by predicting future performance.

Behavior and Activity Recognition

This session explored a wide range of research into detecting human behavior, including action, speech, and mood. The session opened with a presentation by Heng-Tze Cheng from Carnegie Mellon University on “NuActiv: Recognizing Unseen New Activities Using Semantic Attribute-Based Learning.” Through training using component actions, this application could classify a new complex activity using its base components (for example, “pointing” was the sum of an arm raise and arm extension).

“MoodScope: Building a Mood Sensor from Smartphone Usage Patterns,” another paper presented by LiKamWa, delved into the realm of the psyche with a system for predicting the user’s current mood based on his or her interactions with a smartphone (application launches, frequency of email, phone calls, and so on).

The next paper, from the University of Virginia, “Auditeur: A Mobile-Cloud Service Platform for Acoustic Event Detection on Smartphones,” discussed how to classify audio behavior. Auditeur leveraged a hybrid model of cloud planning with on-device classification for detecting and classifying a variety of acoustic features, from a baby crying to a person snoring.

IEEE Computer Society Publications Office

10662 Los Vaqueros Circle
Los Alamitos, CA 90720

STAFF

Lead Editor

Brian Kirk

bkirk@computer.org

Content Editor

Shani Murray

Manager, Editorial Services

Jenny Stout

Assoc. Peer Review Manager

Hilda Carman

Publications Coordinator

pervasive@computer.org

Webmaster

Jennie Zhu-Mai

Contributor

Joan Taylor

Director, Products & Services

Evan Butterfield

Senior Manager, Editorial Services

Robin Baldwin

Senior Business Development Manager

Sandra Brown

Membership Development Manager

Cecelia Huffman

Senior Advertising Coordinator

Marian Anderson

manderson@computer.org

IEEE Pervasive Computing (ISSN 1536-1268) is published quarterly by the IEEE Computer Society. IEEE Headquarters, Three Park Ave., 17th Floor, New York, NY 10016-5997; IEEE Computer Society Publications Office, 10662 Los Vaqueros Circle, PO Box 3014, Los Alamitos, CA 90720-1314, phone +1 714 821 8380; IEEE Computer Society Headquarters, 2001 L St., Ste. 700, Washington, DC 20036. Subscription rates: IEEE Computer Society members get the lowest rates and choice of media option – US\$52 (print and online access). Go to www.computer.org/subscribe to order and for more information on other subscription prices. Nonmember rate: available on request. Back issues: \$20 for members and \$205 for nonmembers.

Postmaster: Send undelivered copies and address changes to IEEE Pervasive Computing, Membership Processing Dept., IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854-4141. Periodicals postage paid at New York, NY, and at additional mailing offices. Canadian GST #125634188. Canada Post Publications Mail Agreement Number 40013885. Return undeliverable Canadian addresses to PO Box 122, Niagara Falls, ON L2E 6S8. Printed in the USA.

IEEE prohibits discrimination, harassment and bullying: For more information, visit www.ieee.org/web/aboutus/whatis/policies/p9-26.html.

Focusing on the social aspect of acoustic sensing, “SocioPhone: Everyday Face-to-Face Interaction Monitoring Platform Using Multi-Phone Sensor Fusion,” presented by Youngki Lee from Singapore Management University, showed how the volume features of speech across multiple devices, combined with the human behavior of taking turns in conversations, can be used for real-time mobile detections of social dynamics in conversations.

Similarly, a video presented on behalf of Mary G. Baker of HP Labs used turn taking to identify the location based on the patterns of silence in a room. The video showed how an application could use these patterns to detect acoustic colocation with others, even across a spatially separated phone conversation.

Assorted Topics

The concluding session of MobiSys 2013 featured presentations in assorted areas of mobile systems. One of us (Mariya Zheleva) presented “Kwiizya: Local Cellular Network Services in Remote Areas.” The work focused on a real-world deployment, bringing low-cost reliable cellular technologies to rural Sub Saharan Africa. The deployment used unmodified cellular handsets to provide reliable voice and SMS communication to locals, overcoming challenges such as unreliable power and limited resources.

In a change of direction, the following talk, “AdRob: Examining the Landscape and Impact of Android Application Plagiarism,” by Clint Gibler and his colleagues from UC Davis, focused on the issue of mobile application plagiarism, both within and across android markets. Their system examined over a quarter million applications across 17 markets, grouping them based on market ID, key signature, and advertisement ID to generate a lower bound for the revenue loss due to plagiarism. In

addition to providing a publically accessible tool to check application plagiarism, the analysis showed that Android developers suffer “at least a 14 percent loss in revenue and a 10 percent loss in user base due to plagiarism.”

Following the thread of security, Jeremy Gummesson from University of Massachusetts, Amherst, presented the last talk of the conference, “EnGarde: Protecting the Mobile Phone from Malicious NFC Interactions.” Gummesson presented a stand-alone device that, when attached to the back of a mobile device equipped with NFC, could intercept, examine, and, if necessary, jam interactions with NFC chips. A novel aspect of this system was the device’s ability to function by harvesting up to 30 mW of energy from the NFC interface on the phone.

MobiSys 2014 is scheduled for 16–19 June in Bretton Woods, New Hampshire. For more information, see www.sigmobile.org/mobisys/2014. ■

Mariya Zheleva is a PhD student in the Department of Computer Science at the University of California, Santa Barbara. Contact her at mariya@cs.ucsb.edu.



Michael Nekrasov is a PhD student in the Department of Computer Science at the University of California, Santa Barbara. Contact him at mnekrasov@cs.ucsb.edu



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