



## DEV 14 and HotMobile 15

### SYMPOSIUM ON COMPUTING FOR DEVELOPMENT

Nicola Dell and Trevor Perrier,  
University of Washington

The Fifth ACM Symposium on Computing for Development (DEV-5) was held on 5–6 December 2014 in San Jose, California. In its fifth edition, DEV continued its progress to becoming the premier venue to present original and innovative work on the applications, technologies, architectures, and protocols for computing in developing regions. In particular, the conference highlighted research results from contexts where conventional computing solutions are often inappropriate due to limited resources and a variety of contextual factors, including limited funds, language barriers, illiteracy, and the limited availability of power and bandwidth. Focusing on innovative technical solutions to these diverse contextual, infrastructural, and user challenges, the DEV-5 conference fostered exchange between computer scientists, engineers, and others interested in the use of information and communication technologies for development.

### FROM INNOVATION TO IMPACT

The opening keynote was an inspiring presentation by Edward Miguel, a professor of environmental and resource economics and faculty director of the Center for Effective Global Action at the University of California, Berkeley. Miguel focused on the intersection of economics, technology, and development, arguing that rigorous experimental research in development economics

### EDITOR'S INTRO

This issue we have not one but two conference report columns for your reading pleasure! We report first on ACM DEV 2014, the 5th annual Symposium on Computing for Development. ACM DEV attracts an interdisciplinary audience of computer scientists working to develop technologies, architectures, protocols, and applications for computing in developing regions. Our second report is on the ACM HotMobile 2015 workshop, held in February of this year. HotMobile is focused on mobile applications, systems, and environments, particularly for new directions or nontraditional approaches in mobile computing. We hope you enjoy our expanded column this issue!

— Elizabeth Belding

has had a large-scale impact on developing countries.

He discussed several compelling examples of randomized control trials in development economics that have led to policy change. One demonstrated the effectiveness of child deworming programs in Kenyan schools, while another focused on convincing farmers in India to grow genetically engineered (and flood resistant) rice. An interesting discussion ensued, in which Miguel described several challenges, such as trying to convince farmers to gamble their entire crop on a new type of rice. He also fielded a number of interesting questions, including the extent to which research results from one country, such as Kenya, could be expected to seamlessly translate to other countries and different scenarios. The key takeaway was that innovative technologies must be backed by rigorous, large-scale evaluations to have the kind of impact that leads to policy change.

### DEV-5 PAPER SESSIONS

The main conference program comprised four paper sessions, describing diverse research on a variety of topics,

including networking, HCI, accessibility, and mobile computing.

### Applications

The first session began with an interesting presentation, “FlashPatch: Spreading Software Updates over Flash Drives in Under-Connected Regions,” by Henry Corrigan-Gibbs of Stanford University and Jay Chen of New York University in Abu Dhabi. In developing countries, many computer users—and even Internet cafes—lack the bandwidth necessary to download and install critical software updates. As a result, software quickly becomes out of date and vulnerable to viruses and malware. To overcome this problem, Corrigan-Gibbs and Chen designed and built a system that distributes software updates and antivirus protection by piggybacking on the extensive flow of flash drives between computers in rural villages. Once installed, the system copies and distributes updates automatically and thus doesn’t require any changes in user behavior. Findings from a nine-month field deployment of the system showed that for 30 percent of the machines using the software, FlashPatch provided antivirus

protection in addition to software updates without imposing any tangible burdens on system owners.

The next talk, by Arun Ramanujapuram and Anup Akkihal of Logistimo, targeted a different application domain: “Improving Performance of Rural Supply Chains Using Phones: Reducing Information Asymmetry to Improve Stock Availability in Low-Resource Environments.” They present the design and implementation of a mobile-phone-based platform that aims to improve stock availability in primary health clinics using a digital “bulletin board” that captures the demand and supply of goods in real time from any location using low-end mobile phones. The bulletin board then broadcasts this information to vendors and managers upstream in the supply chain. The authors evaluated the system during an impressive 14-month study with 29 health clinics in India. After the introduction of the bulletin board, the availability of eight different vaccines showed steady improvement, leveling out at nearly 99 percent for all vaccines measured. In addition, the response time to replenish an out-of-stock vaccine decreased from 14 days to five days per vaccine over a four-month period—an impressive 64 percent increase in responsiveness. Since obtaining these results, the bulletin board system has been scaled to over 400 public health facilities across four states in India.

The final talk in the session described “UrJar: A Lighting Solution Using Discarded Laptop Batteries.” This innovative device, developed by Vikas Chandan and his colleagues at IBM Research in India, aims to transform discarded laptop batteries from e-waste into a valuable energy source that powers low-energy DC devices. In particular, the device can provide lighting to users who lack access to reliable electricity. To motivate their design, the authors conducted a formative study, examining the current practices of retailers and households without grid electricity in

Bangalore, India. They found that 66 percent used battery power, 20 percent used fossil fuel, and 14 percent used solar power for lighting. After developing the device, the authors conducted a preliminary user evaluation by giving UrJar prototypes to five street-side shops. These shop owners said that they used the device and valued the long duration of backup power that it provided. The ultimate goal is for UrJar to channel e-waste toward the alleviation of energy poverty, thus simultaneously providing a sustainable solution for both problems.

#### **Infrastructure for Diverse Users**

The second paper session consisted of three talks that aim to improve infrastructure in low-resource environments.

### **The ultimate goal is for UrJar to channel e-waste toward the alleviation of energy poverty, thus providing a sustainable solution for both problems.**

The first talk, “A Feasibility Study of an In-the-Wild Experimental Public Access WiFi Network,” by Arjuna Sathiseelan of the University of Cambridge and his colleagues, discussed the technical and social constraints of providing Internet access, free at the point of use, by sharing existing broadband subscribers’ connections. A system that used routers running a custom build of OpenWRT, broadcasting an open SSID, was deployed for seven months in a low-income neighborhood in a medium-sized British city. Data from the deployment showed that the system could provide economical Internet access to many people currently digitally disenfranchised.

The second talk was “Mapping Induced Residential Demand for Electricity in Kenya,” by Douglas Fabini and his colleagues at IBM Research,

Nairobi. The paper draws on socioeconomic, demographic, geospatial, and domain relevant data to capture the hypothetical demand that would exist in communities if access to electricity services were available. To demonstrate the applicability of their approach, the authors applied their model to Kenya on a fine geographic basis. Together with spatially explicit cost models for generation, transmission, and distribution, the induced demand predictions could be used to evaluate various technology options, business models, and tariff structures, or to guide the development of a public-sector electrification program.

The session closed with “Rescue Base Station,” presented by Kurtis Heimerl of UC Berkeley on behalf of Ibrahim Ghaznavi (Information Technology University, Lahore) and his colleagues. The key idea presented was to create a drop-in, solar-powered compatible, open source GSM communication system that could be used in the event of large-scale disruption to traditional communication systems due to disasters such as earthquakes, floods, tornadoes, and hurricanes. In addition to connecting people through conventional GSM services—such as phone calls, SMS communication, and smartphone features—the proposed system also has the potential to offer disaster relief services, such as intelligent call routing, SMS broadcast alerts, and emergency short-codes through which a victim could contact available doctors, firefighters, police, and rescue workers.

#### **Design for Resource-Constrained Environments**

The second day of the conference began with a session of three talks on designing for resource-constrained environments. Two of the talks targeted educational settings. Aditya Vashistha of the University of Washington presented, “Educational Content Creation and Sharing by Low-Income Visually Impaired People in India.” This was

one of the few truly qualitative contributions at the conference, describing a fascinating study that examined the educational landscape in low-income, visually impaired communities in rural India. The study found that, despite a clear preference for braille, most participants experienced severely limited access to braille materials. Instead, these communities appear to have established an informal network of peer-produced audio content that is shared via Bluetooth, memory cards, and CDs. The talk concluded by providing recommendations that could help to strengthen educational content production, consumption, and sharing practices within low-income, visually impaired communities.

Staying in the educational domain, Malolan Chetlur of IBM Research India and his colleagues presented “EduPaL: Enabling Blended Learning in Resource Constrained Environments.” This talk proposed a blended learning system that uses a low-cost USB flash drive as a portable learning platform for students with intermittent Internet access. The system consists of the EduPaL client, installed on a USB drive, enabling students to download and consume educational content. The system tracks students’ activities and lets them take notes and ask questions. These activities are then used to produce visualizations that help teachers identify students or topics that need more focus in the classroom. Findings from an exploratory study with 30 students in an Indian university showed that students used EduPaL at different times of the day and in different locations, such as in their hostel, home, and lab, thus enabling learning without the need for a continuous connection.

Switching gears, the final talk in the session, by Carlos Rey-Moreno of the University of the Western Cape described “Co-Designing a Billing System for Voice Services in Rural South Africa: Lessons Learned.” The paper reflected on the co-design of a billing

system for a community network in rural South Africa, focusing on the legal, financial, social, and technical feasibility and constraints of different billing methods. The value of the co-design approach described by Rey-Moreno is its sensitivity to context and its emphasis on promoting active community participation to ensure the project’s sustainability. Additionally, the process revealed factors embedded in the provision of voice services by mobile network operators in South Africa that prevent economically poor and illiterate users from fully benefiting from these services.

#### Location and Mobile Phones

The final session of the conference consisted of two talks that focused on location. The first was “PlaceMap: Discovering Human Places of Interest

### **Modi argued that new innovations will have more impact if the technologies are given away for free.**

Using Low-Energy Location Mobile Phones,” by Kuldeep Yadav of Xerox Research Centre, India and his colleagues from IIT Delhi. PlaceMap aims to discover places and routes using only GSM information. It stores all discovered places and routes and uses them to build spatiotemporal mobility profiles for users. The system was verified through an offline evaluation of the PlaceMap algorithms on two real-world datasets: self-collected data from 62 participants over a four-week period in India, and data from 38 participants at Nokia’s Mobile Data Challenge in Switzerland. Findings show that PlaceMap discovered up to 81 percent of places correctly.

The last talk of the conference, by Joshua Blumenstock of the University of Washington, presented “Probabilistic Inference of Unknown Locations: Exploiting Collective Behavior When

Individual Data Is Scarce.” The paper describes a method for predicting the approximate location of a mobile phone subscriber in contexts where the signal generated by each individual might be intermittent, but the collective population generates a large amount of data. The method works well when, for example, an individual isn’t consistently active on the network or when the phone is switched off. The model uses a nonparametric approach to probabilistically interpolate locations, and it has the advantage of associating a confidence with each prediction. Through tests on a large dataset of mobile phone records from Afghanistan, the authors show that the model can correctly predict a subscriber’s location in 76 to 95 percent of cases, and that, on average, the predicted location is off by only 0.2–1.9 kilometers.

#### **SERVICE DELIVERY IN THE DEVELOPING WORLD**

The conference closed with a thought-provoking talk by Vijay Modi, a professor of mechanical engineering and a faculty member at the Earth Institute, Columbia University. In his talk, Modi described an innovative metering system for shared solar energy that was deployed by his team in Mali. The system allowed people to purchase solar electricity using prepaid scratch cards similar to mobile airtime. Modi went on to describe lessons learned from the system, heavily emphasizing the need for solutions to be lightweight and highly adaptable to different or changing contexts. For example, the solar energy system should be designed so that it’s easy to add more solar panels as the community demand for energy grows. He also advocated strongly for a stronger focus on demand rather than supply, for energy and water service delivery, and for a deeper understanding of consumers’ relationship with service providers. Finally, Modi argued that new innovations will have more impact if the technologies are given away for free, so that they can be manufactured

cheaply and locally in the communities where they're needed.

### LIGHTNING TALKS

Taking advantage of the Silicon Valley location, the conference included a session of “lightning talks” by Bay Area-based organizations that have been making an impact through technology in developing countries. One organization, Potential Energy, designs, constructs, and distributes stoves that aim to save money and lives by using half the fuel, cooking twice as fast, and emitting half the smoke of an open fire ([www.potentialenergy.org](http://www.potentialenergy.org)). Premise is an app that gives users a simple list of tasks, such as taking photos of specific locations or finding particular products ([www.premise.com](http://www.premise.com)). When completing tasks, users earn points that they can redeem for mobile airtime or other rewards. Finally, LaborVoices provides a platform that global workers can use to share feedback on their safety and working conditions ([www.laborvoices.com](http://www.laborvoices.com)). Corporate brands and suppliers can then identify and solve problems before they become urgent.

### POSTER SESSION

DEV-5 included a poster session of nine posters showcasing an exciting array of works in progress and smaller research contributions. The poster session provided a great opportunity for conference attendees to talk to researchers in person, ask questions, and form collaborative ties. Several of the posters were presented by talented undergraduate researchers: Joy Ming from Harvard demonstrated, “Printing Paper Technology for Development,” and Krittika D’Silva from the University of Washington presented, “A Mobile Application for Interactive Voice Forums: Design and Pilot Deployment in Rural India.” In general, the posters represented a wide array of diverse research projects, including work on citizen engagement, social media, education, empowerment, energy, and networking, all of which have the common

goal of improving lives in underserved communities.

The conference also included several social receptions that encouraged attendees to discuss research trends and trajectories in the field of Computing for Development. The first reception, sponsored by Facebook, took place the night before the conference began, while the second was on the evening of the first day. One particularly interesting series of discussions among attendees revolved around how one of the primary strengths of the DEV community is that it is truly diverse and interdisciplinary in nature, both in the research that is presented and in the makeup of the conference attendees. Unlike conferences that focus on a specific subdomain of computer science, DEV draws participation from top researchers in systems, networking, security, HCI, assistive technology, mobile and ubiquitous computing, and others, and provides an unique opportunity for researchers in all of these domains to collaborate and exchange ideas. Each of these areas of expertise contributes uniquely to the community, and the exchanges between attendees have the potential to further widen the field and enable exciting new research that could positively impact the lives of people in underserved communities.

**Nicola Dell** is a PhD candidate at the *University of Washington*. Contact her at [nixdell@cs.washington.edu](mailto:nixdell@cs.washington.edu).



**Trevor Perrier** is a PhD student at the *University of Washington*. Contact him at [tperrier@cs.washington.edu](mailto:tperrier@cs.washington.edu).



### HOTMOBILE 2015

*Brad Campbell, Noah Klugman, and Thomas Zachariah*  
*University of Michigan*

On 12–13 February, 96 attendees came together in Santa Fe, New Mexico for the 16th Workshop on Mobile Computing Systems and Applications (ACM HotMobile 2015). The program consisted of 23 paper presentations, 18 posters, and 10 demos. To lead off the conference, Mark Corner, CTO of Fiksu and an associate professor at the University of Massachusetts Amherst, gave the keynote address, highlighting his experiences transitioning from a successful academic career to a successful startup.

### PURSuing A STARTUP AS AN ACADEMIC

Corner described his views on the two paths—academic versus startup—while touching on the good and the bad from each career. He started by describing the benefits of academia, including “exciting” moments such as receiving a grant or getting a paper accepted, which provide acute satisfaction to faculty and students. He praised academia for being able to solve important problems that have limited commercial viability but are still worth addressing. He also noted that academic institutions are great at producing the next generation of PhD graduates and future academics. However, Corner also addressed the drawbacks of an academic career, noting that the popular argument, “You can do whatever you want,” doesn’t always hold in the face of Institutional Review Board requirements, funding availability, and legal and social limitations. He touched on the complexities of finding grants and questioned the reach of academic research, asking if anyone even cares, given how few projects reach mainstream adoption.

Drawing from his experiences at the mobile-ad technology startup Fiksu, Corner then addressed the pros and cons

## CONFERENCES

of pursuing a technology-based startup. First, he noted the positives: starting a company only takes a week, there's a tremendous existing infrastructure of online tools and services for building the company, and watching the company grow is particularly rewarding. Second, he described the negatives: it's extremely difficult, unlikely, and stressful to create a successful company; the job requires managing nontechnology aspects and interacting with nonscientists; and it's often difficult to manage (including hiring and firing) employees.

Incorporating insight from both of his careers, Corner made note of three reasons he believes an academic might be the best fit for a startup. First, academics excel at defining problems, constructing a solution, and arguing for that solution. Second, they evaluate those solutions with targeted experiments—a critical task when growing a company. Third, teaching material in an area cultivates more comprehensive knowledge of the topic than simply taking the course. Although academic research might not transition well to a successful company, Corner argues that the people behind the research are well suited for building startups.

### HOTMOBILE PAPER SESSIONS

There were eight sessions, each of which covered a different topic in the mobile computing research space, and all of which sparked much discussion both during the sessions and afterward in the hall.

#### Gesture

The lead-off session had three papers exploring gesture-based interactions with mobile devices. The first presenter was Kaori Ikematsu of Ochanomizu University, who spoke about “Memory Stones: An Intuitive Information Transfer Technique between Multi-Touch Computers.” Memory Stones are a technique for using the same multitouch gesture on two nearby touch-enabled devices to transfer content from one to the other. The touch gesture is converted to a polygon and

sent to a server along with the currently displayed content. The server then determines if a matching gesture-polygon occurs on another physically local device; if so, it transfers the content to the device. Addressing privacy concerns, Ikematsu provided evidence that the gesture is difficult to copy simply by watching.

The second paper on gestures, “Finger-Writing with Smartwatch: A Case for Finger and Hand Gesture Recognition Using Smartwatch,” presented by Parth H. Pathak of UC Davis, explored a new method for interacting with devices with small or no screens using the accelerometer and gyroscope common to smart watches to detect hand and finger gestures. Because the ten-

**Although academic research might not transition well to a successful company... the people behind the research are well suited for building startups.**

dons in human hands extend to the wrist, finger movements can be detected by a wrist-worn sensor, and the authors found that they could detect 94 percent of letters traced by a finger.

Nicola Dell, from the University of Washington, presented the third paper, “Mobile Touch-Free Interaction for Global Health.” Motivated by the needs of health clinics using smartphones to provide care, this work lets smartphone users interact with apps without touching the device, mitigating contamination concerns. Dell defined five gestures that can be detected by the front-facing camera and demonstrated that they can be effectively used to operate an app while only consuming 6 percent of the phone's battery per hour.

#### Internet of Things and Energy

The next session tackled problems emerging with the Internet of Things

and smartphone energy use. First, Angli Liu, from the University of Washington, presented “Retro-VLC: Enabling Battery-Free Visible Light Communication for Mobile and IoT Applications,” demonstrating a custom handheld device that uses backscatter visible light communications (VLC) for connectivity and solar panels for harvesting energy. The device leverages a photodiode, a retro-reflector, an LCD, and a microcontroller to provide two-way VLC communication within the small power budget provided by energy harvesting. This represents a significant advancement in VLC for indoor batteryless IoT devices.

Next, Thomas Zachariah, from the University of Michigan, presented “The Internet of Things Has a Gateway Problem,” which argues that the current system of connecting IoT devices to the Internet using application- and device-specific apps and individual gateways is fundamentally broken. Zachariah proposed a new model, where all smartphones act as transparent Bluetooth Low Energy gateways, which any device can use to connect to the Internet. This open gateway raises many questions concerning feasibility, privacy, and adoption if the system is to be effective.

The third presentation, “Reducing Energy Consumption of Alarm-Induced Wake-Ups on Android Smartphones,” was presented by Sewook Park of Yonsei University. Park describes a system that attempts to maximize the time smartphones spend in sleep mode by removing unnecessary alarm-induced wakeups. Noncritical alarms are deferred and executed when a time-critical alarm does occur. The system makes an attempt to auto determine alarm criticality by taking into account time criticality and the wakeup's effect on the system. Results from a user study revealed that energy savings ranged from 2.6 percent to 19.3 percent—the more apps installed on the system, the greater the energy savings.

The final presentation in this session, “Energy-Efficiency Comparison

of Mobile Platforms and Applications: A Quantitative Approach,” by Weisong Shi of Wayne State University, looks to improve smartphone battery performance by providing tools for comparing smartphone apps and devices in a fair and consistent way. They measure smartphone apps from a range of categories while ensuring that unmeasured variables are held constant in all trials. Their results provide insights for comparing the energy performance of apps and found that buffering streaming video was more energy expensive than consistent streaming.

### Cellular

To begin the session on topics related to cellular networks, Wenlu Hu of Carnegie Mellon University presented, “The Case for Offload Shaping,” which proposes performing on-device work or offload shaping to improve the performance of offloading computation to the cloud. Offload shaping can greatly reduce the offloaded workload, saving processing time, network bandwidth, and energy. An offload shaping API is proposed and evaluated using a video blur detection and removal example.

Next, Emir Halepovic of AT&T Research presented, “Can Accurate Predictions Improve Video Streaming in Cellular Networks?” The talk explored the use of prediction to improve video quality and user experience over highly variable bandwidth in cellular networks. It demonstrated that trivial prediction isn’t sufficient and prediction-based adaptation is key. The proof-of-concept implementation combines prediction with rate stabilization functions and can reduce the optimality gap to just 4 percent.

To end the session, Feng Lu of UC San Diego presented, “CQIC: Cross Layer Congestion Control for Cellular Networks.” He explained that TCP can’t keep up with the fast-changing nature of last-hop cellular links. The presentation suggested using channel-capacity estimation based on physical layer information to vastly improve congestion-con-

trol design. CQIC, a congestion-control variant of Google QUIC, eliminates slow start and bufferbloat, and performs significantly better than TCP.

### Network Monitoring and Imaging

The three papers in this session explored different applications of using mobile phones as sensors. The first paper of this session, “CrowdREM: Harnessing the Power of the Mobile Crowd for Flexible Wireless Network Monitoring,” was presented by Andreas Achtezehn of RWTH Aachen University. Achtezehn described a crowd-sourced cellular-network framework capable of monitoring large areas. In this system, the smartphone of each member of the crowd passively

**Their results provide insights for comparing the energy performance of apps and found that buffering streaming video was more energy expensive than consistent streaming.**

gathers RF signal-strength measurements, user-context information, and control traffic traces. Measurements are cross-correlated in the cloud. The authors performed a 1,000-hour study in which they measured the paging rate (a measure of mobile device discovery overhead), number of active users, and network load in Aachen, Germany.

Yibo Zhu, from the University of California, Santa Barbara, presented “60Ghz Mobile Imaging Radar.” This talk addressed the growing trend of smartphone apps that are enabled by small, low-cost sensors that can determine the position, shape, and surface material of nearby objects with high accuracy. An attractive approach for constructing such a sensor uses reflections of high-frequency RF signals, which can be generated by inexpensive radios. Furthermore, these reflections are

relatively immune to environmental factors. The authors describe their RF radar system, “Nightcrawler,” constructed of two 60GHz radios: one that transmits and one that receives the reflections from the transmitter. By taking advantage of user mobility, Nightcrawler can emulate a large aperture virtual antenna. A small amount (1 meter) of mobility lets Nightcrawler achieve 2 cm imaging accuracy from an object 8 meters away. Additionally, Nightcrawler can accurately determine the material of the object.

Tan Zhang of the University of Wisconsin-Madison presented the final paper of the session, “A Wireless Spectrum Analyzer in Your Pocket.” The work describes “Snoopy,” a low-cost RF frequency scanner constructed from an 802.11 NIC and a frequency translator that can sense from 30 MHz to 7.5 GHz. A prototype implementation of Snoopy is demonstrated to achieve low error rates (3–15 percent) for detecting primary signals at –90dBm. Additionally, Snoopy can achieve a median error of less than 4 dB in measuring power for most TV channels. Further work seeks to implement Snoopy as a hardware peripheral for off-the-shelf mobile devices.

### Location

This session focused on new techniques to locate individuals and devices. Shijia Pan of Carnegie Mellon University started the session with her presentation on “Indoor Person Identification through Footstep Induced Structural Vibration.” Pan made the observation that each person’s gait induces identifiable structural vibration patterns. She introduced a system that senses floor vibrations and correlates sensed patterns with individuals. This system is shown to achieve 83 percent identification accuracy in an experiment with five individuals walking down a 16 m hallway, one at a time. Future work will include identifying groups.

Next, Bo Han of AT&T Labs–Research discussed “Human Assisted Position Using Textual Signs.” This work presents an offline localization

system that leverages human users' perceptions of textual signs. A comprehensive sign database can be constructed and will be small enough to store locally (less than 31 kB for a London dataset). By asking a user to enter the text from any nonstandard street signs in view, the user can be localized based on the largest cluster of matches in the sign database. A speech-recognition library was described that also, due to its small size (less than 63 kB for a London dataset), can be accessed locally. Han and his colleagues evaluated the system using sign data from Google street view and OpenStreetMap, and the system accurately localized 90 percent of locations with four signs and 99 percent of locations with seven signs.

To wrap up the session, Jeremy Gummeson of HP Labs presented "Step-by-Step Detection of Personally Collocated Mobile Devices." This work demonstrates a lightweight method to determine device collocation through the comparison of accelerometer-based footstep signatures across devices. If one device detects steps and the other does not, then it's possible the owner has forgotten a device and is walking away without it. Step counts are binned in one-second intervals to aid the comparison and to reject false positives. Compared against RSSI-based collocation techniques, this system achieves between 3 and 7.5 times lower latency. Additionally, this method is shown to use only 21  $\mu$ W of average power.

### Mobile Platforms

Three works exploring platform-level API changes came next. First, Anudipa Maiti of the University at Buffalo presented, "The Missing Numerator: Toward a Value Measure for Smartphone Apps," which looks beyond merely using energy consumption measurements to address the energy management of smartphone apps. The work aims to treat energy as a fundamental resource that can be multiplexed by an OS, much like the way

storage and computation are currently treated. Maiti proposed a method for measuring an app's "value" metric to be used for fairly distributing energy. A solution has yet to be found, but it's clear that the value must be representative of the temporal variance exhibited by apps.

Together, Ali Ben Ali and Jerry Ajay of the University of Buffalo gave an entertaining presentation on "Maybe We Should Enable More Uncertain Mobile App Programming," in an attempt to remove the strict requirements of certainty for app developers when it comes to the mobile environment on which the app is run. To do so, they introduce a new language construct, "maybe," allowing for a liberat-

### Taking inspiration from a paper shredder... the authors splice and subsample audio samples so words can no longer be identified.

ing development approach based on structured uncertainty. A server-based scoring system resolves uncertainty on each specific device, which can help provide different experiences catered to particular user groups and mobile platforms.

To conclude the session, Kyungmin Lee of the University of Michigan presented, "The Case for Operating System Management of User Attention." Lee discussed synchronization issues between app notifications and available user attention. To resolve this disconnect, he proposes that the mobile OS use inputs from on-board sensors—such as GPS, accelerometer, and microphone sensors—to extract user context and manage user attention as a resource. The solution can be modeled as a scheduling problem that considers the user's previous app usage and current activity, as well as the app's desired interactions.

### Learning

Both papers in this session applied machine-learning techniques to discover user context. Nicholas Lane of Bell Labs presented, "Can Deep Learning Revolutionize Mobile Sensing?" This work pushes for the use of deep learning models to reliably infer user behavior and context from sensor data. A low-power Deep Neural Network (DNN) inference engine is constructed to run on modern mobile hardware. This implementation is compared against Decision Trees and Gaussian Mixture Models, two commonly used inference engines. While running on a smartphone's DSP, the DNN demonstrates equivalent or better accuracy when determining context in three datasets. Additionally, the DNN has low overhead and can scale better than the other models with increasing numbers of inferred classes.

John Rula of Northwestern University presented "Mobile AD(D)." Mobile ads are priced, sold, and configured prior to their viewing. By predicting the exposure time a user will have with an ad, advertisers could make better decisions and ads could be better priced. This work analyzes the mobile device usage patterns of 37 individuals in two countries for two weeks and builds a model to predict exposure time based on temporal components, user activity, network connectivity, and type of application. A classifier was constructed that outperforms the mean error in over 90 percent of cases.

### Privacy

The final session of HotMobile 2015 contained two papers on protecting privacy in the cloud. The first paper, "Policy-Carrying Data: A Privacy Abstraction for Mobile Services," was presented by Stefan Saroiu of Microsoft Research. A scheme is introduced that encrypts data being sent to the cloud using ciphertext-based-attribute-based encryption (CP-ABE). The user selects the attributes used to encrypt the data to form a terms-of-service

(ToS) agreement. When data is uploaded to the cloud, a cloud provider must explicitly agree to each attribute in the ToS agreement for decryption. The cloud provider's agreement to the user-specified ToS is legally binding and, while there is no guarantee it will be followed, does increase the cost for any policy violations. An early version of this system exhibited minimal performance overhead.

The final talk was Sumeet Kumar, from Carnegie Mellon University, presenting, "Sound Shredding: Privacy Preserved Audio Sensing." This work attempts to remove the ability to recognize words from audio samples while preserving the less invasive abilities to perform context, gender, and identity detection. Taking inspiration from a paper shredder, which removes the ability to determine the words on the page but leaves information about the paper type, ink type, and other properties readily available, the authors splice and subsample audio samples so words can no longer be identified. The authors recovered gender and user identity of a sample after shredding but couldn't detect speech.

### IN MEMORIAM: GAETANO BORRIELLO

At the banquet on 12 February, all those attending HotMobile 2015 held up a glass to remember Gaetano Borriello, a friend and mentor to many in the room. Roy Want of Google presented a chronology of Borriello's career and shared stories from their many years of friendship and collaboration. He also shared letters of remembrance sent in from people such as the famed primatologist Jane Goodall and organizations such as the Red Cross. [For more information, see this issue's In Memoriam department.] Borriello's commitment to serving those in need, in part through his widely used Open Data Kit project, was emphasized and the success of this effort was celebrated. In addition, Gaetano was remembered

for his teaching. He was the Jerre D. Noe Professor of Computer Science and Engineering at the University of Washington, where he was awarded both the UW Distinguished Teaching Award and the UW Marsha L. Landolt Distinguished Graduate Mentor Award.

**At the banquet... all those attending HotMobile 2015 held up a glass to remember Gaetano Borriello, a friend and mentor to many in the room.**

The University of Washington has established the Gaetano Borriello Endowed Fellowship for Change (<https://www.washington.edu/giving/make-a-gift/?page=make&code=GBEDFL>). We hope that those reading this article will consider supporting this fellowship.

**Brad Campbell** is a fourth-year PhD student at the University of Michigan. Contact him at [bradjc@umich.edu](mailto:bradjc@umich.edu).



**Noah Klugman** is a first-year PhD student at the University of Michigan. Contact him at [nklugman@umich.edu](mailto:nklugman@umich.edu).



**Thomas Zachariah** is a second-year PhD student at the University of Michigan. Contact him at [tzachari@umich.edu](mailto:tzachari@umich.edu).



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



## Call for Articles

**IEEE Software** seeks practical, readable articles that will appeal to experts and nonexperts alike. The magazine aims to deliver reliable, useful, leading-edge information to software developers, engineers, and managers to help them stay on top of rapid technology change. Topics include requirements, design, construction, tools, project management, process improvement, maintenance, testing, education and training, quality, standards, and more.

Author guidelines:  
[www.computer.org/software/author.htm](http://www.computer.org/software/author.htm)  
 Further details: [software@computer.org](mailto:software@computer.org)  
[www.computer.org/software](http://www.computer.org/software)

IEEE  
**Software**