Editor: Mary Baker ■ HP Labs ■ mary.baker@hp.com

HotMobile 2006: Mobile Computing Practitioners Interact

Maria R. Ebling

Like the First IEEE Workshop on Mobile Computing Systems and Applications, held in 1994, the 2006 workshop aimed to foster interaction between mobile computing practitioners. In keeping with this goal, we returned to a small, informal workshop with few papers but significant discussions. We accepted just nine papers, but we had two important group discussions, two exciting panels, and an insightful keynote address. To reflect these changes, the steering committee changed the nickname of this year's workshop to HotMobile 2006. Approximately 40 people attended the two-day event, held at the Semiahmoo resort in Blaine, Washington.

MOBILE PHONES ONLY?

The workshop opened with a debate on the statement: "Resolved: The mobile phone is the only device people will carry in the future."

The pro position was basically that cell phones are already ubiquitous. Gartner estimated that in 2005 the number of cell phones sold reached 780 million and that the number will hit 2.6 billion by 2009. In India and China, cell phones appear to be the primary computing device. Given such a high penetration, application developers will concentrate on these devices, particularly addressing the power and storage challenges.

The con position was basically that people use a variety of different devices, including cell phones, watches, PDAs, MP3 players, and laptops. Combining all these devices' functionality into a single device—a "Swiss army knife" approach—might result in a device that does many things, but none of them well. Because the price of single-use devices will significantly decrease, it might be more appropriate for users to carry specialized devices that have the right form factor and user interface for the task at hand (for example, the iPod). Furthermore, fashion often dictates what people carry. For instance, some people wear watches for reasons that have nothing to do with time (for example, aesthetics).

During the debate, some attendees ended up arguing not only for their assigned position but also occasionally for the other side. At the end, we took a vote: just nine people supported the proposition!

MOBILE PHONES, LOCALIZATION

Our first paper session considered mobile phones as appliances. John Barton (IBM Almaden Research Center) presented "Mobile Phones Will Become the Primary Personal Computing Devices." He argued that, owing to increasing storage and computing power, mobile phones will eventually replace PCs.

John Davis (IBM T.J. Watson Research Center) then presented "Supporting Mobile Business Workflow with Commune." He proposed a workflow management system that employs "mini-workflows," network-isolated components that can be offloaded onto mobile clients by leveraging Web services.

Our second session focused on localization. In "Indoor Localization Using Camera Phones," Nishkam Ravi (Rutgers Univ.) proposed an indoor-localization scheme based on camera phones worn as pendants. The camera phone automatically takes pictures and transmits them over GPRS (General Packet Radio Service) to a centralized server. The server localizes the user by matching the current picture to a database of preloaded pictures. The discussion following this presentation focused on issues such as training costs, accuracy, and whether the entire system could run on the phone.

In "Are GSM Phones THE Solution for Localization?" Alex Varshavsky (Univ. of Toronto) argued that localization using GSM phones might be sufficient for many location-aware applications. His research shows that, with GSM fingerprinting, it's possible to achieve 2–5 m median error indoors, perform room-level localization, and achieve 70–200 m median error outdoors. Moreover, the system can detect with high accuracy the places people visit.

HAS LOCALIZATION BEEN SOLVED?

Gaetano Borriello (Univ. of Washington) moderated a panel exploring whether localization is a solved problem. The panelists were Dieter Fox (Univ. of Washington), Mike Hazas (Lancaster Univ.), and Jeff Hightower (Intel Research Seattle). Borriello had each panelist respond to four questions.

Published by the IEEE CS and IEEE ComSoc ■ 1536-1268/06/$20.00 © 2006 IEEE

Cell phones are the location-aware platform of choice. We should focus all our attention to improving location systems on phones (accuracy, privacy, performance, etc.). There are no other viable platforms. If it doesn't work on a cell phone, why bother?

Fox answered that it doesn't matter what platform we use for our research because everything can be integrated on the phone eventually. Hazas pointed out that people might need to interact with more than one device; Hightower agreed with that statement.

Indoor location systems will piggyback whatever outdoor system becomes dominant. Can special-purpose indoor infrastructure ever be practical for deployment, or will location systems have to converge?

The general consensus was that the systems would converge. Hazas responded that indoor systems requiring special hardware are too expensive. Hightower added that, if two systems are required, they should be implemented on the same device so that users don't have to carry more than one device. Fox argued that special infrastructure isn't required because Wi-Fi signal strength information will be available in virtually any building.

Getting a coordinate is a solved problem. No more papers need to be published on the issue. Shouldn't research now only focus on what to do with the coordinates to solve real problems?

Fox and Hightower basically agreed with this statement. Hightower suggested that future research should focus on place detection, learning, and labeling; combining activity inference with location; and designing applications with location awareness. Fox pointed out that, although we're mostly done with coordinate-based localization, we still don't know how to combine location with activity recognition, how to learn and maintain personalized maps to predict a person's future location, or how to combine information from multiple people. Taking the opposing viewpoint, Hazas responded that we aren't done because we don't know how to deploy cost-effective, practical coordinate-based systems.

The only people who really care about location privacy are researchers, lawyers, and bloggers. When you get right down to it, regular people just don't care that much, so let's stop worrying about it. OK?

Fox responded that it all depends on context. He noted that most people are uncomfortable being tracked but that the elderly might accept it. Hazas commented that, although people think they care about privacy, they really don't. Hightower argued that privacy will always be an important design goal, adding that people are pragmatic, that privacy isn't an all-or-nothing issue, and that it isn't solely a technology issue. To him, the goal is to help people avoid socially awkward situations, to support clarity in interpersonal interactions, and to provide transparency and reciprocity.

In the following discussion, Nina Bhatti (HP Labs) commented that people ask for tracking of objects or beings that have no strong privacy requirement. However, Borriello reminded everyone about indirect information disclosure—that the tracking of one entity can result in another entity's loss of privacy.

Tapan Parikh (Univ. of Washington) suggested that, with such great potential for information abuse, the real question is whether we should disclose the information at all. Fox stressed the importance of supporting the localization system on the user's device in convincing people to use the system. Hightower argued that people are pragmatic and will "sell" their privacy for the right application. Hazas advised that, although privacy is important, we shouldn't be paranoid about it.

M. Satyanarayanan (Carnegie Mellon Univ.) then observed that the discussion had been focusing on absolute location but that relative location is also important. He expressed an interest in an application that won't show sensitive data if somebody is nearby and might be looking at the screen. Hazas reported that he has been working on relative location using ultrasound devices.

FUTURE MOBILE APPLICATIONS

On the workshop's second day, Frederick Kitson, the vice president and director of Motorola Labs' Applications Research Center, gave the keynote address. He shared his team's vision of future mobile applications. He noted that more than 70 percent of i-mode (a wireless Internet service) revenue comes from entertainment applications, such as music, sports, and games. One goal of Motorola's research is to drive seamless mobility: to simplify effort, satisfy full mobility, and amplify the user experience.

He described a vision of "cache and carry" that transparently "mobilizes" dynamic content. Currently, users consume only a fraction of the content that they pay for, partly because the content they capture isn't located where they want to consume it. He also described the Push to X technology (originally called iDEN). With this technology, you could, for example, dynamically add visual content to an existing audio connection.

Kitson also defined "ambient communication." Whereas communication is intentional and conscious today, communication might be unintentional or subconscious tomorrow. In this world, one person might send a message without knowing it and the other person might receive that message peripherally, or "ambiently." People will feel more connected in a less obtrusive manner and will have greater social awareness.

CONFERENCES

Kitson also presented some daunting statistics. In the US, fewer than 10 percent of WAP (Wireless Application Protocol) phone users actually use the browser. Furthermore, among those that do, 50 percent are lost with each click (that is, they give up what they were trying to do, presumably because it took too much effort). To address these challenges, Motorola has been working on the SCREEN3 system, which transmits data to idle cell phones in the background, with no noticeable effect on performance. The data is personalized and scrolls across the screen. If the user pushes any button, the scrolling stops. If the user clicks on something, the phone displays more content. To represent the amount of data displayed with each additional click, Kitson used a "bite, snack, meal, feast" analogy. The bite and snack are cached on the phone. As the user requests the snack, the meal is prefetched, and so on. Motorola has also considered using this model for content delivery of location-based services, focusing the content on nearby services.

CONSIDERING THE ISSUES

After the keynote address, five breakout groups reported on discussions they held the day before. The first team considered various networking technologies. They examined issues including the impact of community-based networks, the impact of having wireless connectivity at highway speeds, and the need to support further research on disconnected operation.

The second team considered application issues. They dealt with questions such as why we don't see more applications research, what mobile applications' future directions might be, what's needed for mobile applications research to succeed, and what characterizes good applications research.

The third team considered device symbiosis, which they defined as "two or more devices being combined, as peers offering independent value, for a task." They further clarified that device symbiosis happens opportunistically and isn't configured. They identified possible applications of device symbiosis and discussed the importance of standards to the field's success. Finally, they named three challenges facing device symbiosis: creating critical mass, developing standards, and optimizing user experience.

The fourth team considered privacy. They first discussed what makes mobile environments different. They then proposed three ways to help ensure privacy:

- symmetric privacy (focusing on full disclosure with a mandatory audit trail),
- aggressive user interfaces (focusing on informing the user about leaked information), and
- helping the user (focusing on using information from an aggressive interface to suggest ways to increase the user's privacy).

Smart phones combine the portability of cellular phones with the computational and networking power of PCs to offer rich functionality. But they face many vulnerabilities.

The final team considered cross-disciplinary research. They discussed such questions as how techniques from other fields apply to mobile computing, which techniques and fields are the most important, and on what fields mobile computing seems to be having the most impact.

PRESENCE SHARING, EXPLOITING MOBILITY

The third paper session examined how to find the right balance for users. Varun Marupadi (Duke Univ.) presented "Presence-Exchanges: Toward Sustainable Presence-Sharing." He introduced a trusted broker into presence-sharing applications so that misbehaving users may not learn about the presence of others without sharing their own identity.

Anthony Nicholson (Univ. of Michigan) presented "Exploiting Mobility for Key Establishment." He observed that most Internet traffic today is unencrypted and blamed this on the lack of available easy-to-use tools for users. He and his colleagues propose a model that establishes keys insecurely and then automatically confirms them by exchanging cryptographic hashes of the keys over many different paths. It does this by utilizing inherent user mobility and overlay networks.

SECURE MOBILE COMPUTING

Ramón Cáceres (IBM T.J. Watson Research Center) moderated a panel session on secure mobile computing. The panelists were Carl Ellison (Microsoft), Steve Gribble (Univ. of Washington), Helen Wang (Microsoft Research), and Jason Hong (Carnegie Mellon Univ.).

Cáceres explained why we should talk about mobile security, summarizing a number of disturbing security breaches. He then asked the panelists four questions:

- Can we achieve secure mobile computing anytime soon?
- Is security in mobile computing different from security in general computing?
- Can we build usable security and privacy functions into mobile environments?
- Will trusted computing hardware and virtual machines play a big role in secure mobile systems?

Ellison compared mobile computing platforms to 1980s PCs. They both support single users, have a handful of software providers, have low CPU power compared to "real" computers, have a small amount of memory, hunger for

features, and use tricks to achieve features in spite of their limitations. A significant difference is that mobile computing platforms have been networked from day one and aren't physically protected via isolation. Consequently, mobile computing platforms potentially face even worse security problems than '80s PCs did.

Gribble opened by saying "Hold on a minute—we still haven't figured out secure *nonmobile* computing!" He pointed at spyware, phishing, worms, denial-of-service attacks, and flawed software as supporting examples. He identified some wide-open issues that have nothing to do with mobility, such as giving users a conceptual model of security and enabling safe sharing in a hostile environment. He argued that mobile devices exacerbate security issues but also provide some opportunities. For example, mobile devices might let us combine the device's physical context with digital security by requiring the user to touch the device to authorize a communication.

Wang talked about threats to smart phones. They combine the portability of cellular phones with the computational and networking power of PCs to offer rich functionality. But they face many vulnerabilities, including attacks from the Internet (for example, worms, viruses, and Trojan horses), infections from the desktop through synchronization (for example, compromise one and you can compromise both), and peer attacks.

Hong started by summarizing his opinion: "Outlook not so good." He believes that secure mobile computing faces significant challenges, one of which is usability. According to Hong, approximately 20 percent of Wi-Fi access points are returned because people couldn't figure out how to make them work. He guesstimated that about 80 percent of Wi-Fi access points aren't secured. He argued that we need invisible security models that are extremely easy to use. He finished with an observation about economic issues, pointing out that the solution's cost can't exceed the problem's cost.

John Barton commented that, in the physical world, we have police to fight criminals, but in cyberspace, everyone has to fight them. He wondered why cyberpolice can't fight the bad guys. The panel responded that we still carry keys to our front doors; that is, things aren't so different between the physical world and cyberspace.

John Davis asked why the credit card model (limiting a cardholder's liability for a stolen card to, say, $50) wouldn't work. The panel responded that whether it would work depends on your loss. If the loss is financial, then the credit card model might be sufficient. If the loss is of a secret or a life, the credit card model is inappropriate.

Another audience member argued that in cyberspace, an attacker expects to not be discovered, as opposed to the real world, where he or she expects to be discovered. The panel responded that about 60 percent of murders and 20 percent of burglaries are resolved. No one knows the percentage for cybercrimes because they're simply not reported.

Someone else observed that we need to know the threats we're defending against. Some open questions include how to figure out what they are, what the possible defenses are, and what the trade-offs are.

MAKING CONNECTIONS

Our final paper session focused on making connections. Richard Gass (Intel Research) presented "Measurements of In-Motion 802.11 Networking." This research measured a commodity laptop's ability to communicate with 802.11 access points while in a car traveling at 5 to 75 mph. The findings reveal that a significant amount of data can be pushed through the wireless link. However, the performance suffers owing to application-related problems, such as protocols with handshaking and long round-trip times.

In "Connection Time for Strange Devices," John Barton related experiences connecting small mobile computers to other computers. He showed that the benefit of connecting phones to larger displays and keyboards might outweigh the burden of making the connection.

Feedback about the workshop indicates that people enjoyed the return to the informal, highly interactive format. That success came because of the hard work of many individuals, including the members of the program, organizing, and steering committees. We thank them all for their efforts! Special thanks to Fred Kitson for his keynote address. For more details on the workshop, see the proceedings[1] and http://research.ihost.com/wmcsa2006.

HotMobile 2007 will run from 26 to 27 February in Tucson, Arizona (for more information, see www.cs.toronto.edu/hotmobile2007). Nina Bhatti and Eyal de Lara (Univ. of Toronto), the general chair and program chair, plan to keep the same format. We hope to see you (and your hot research) there!

ACKNOWLEDGMENTS

This article is based on notes by Tapan Parikh and Alex Varshavsky and on "HotMobile 2006: Workshop on Mobile Computing Systems and Applications Digest of Proceedings," by Maria R. Ebling, which appeared in *Proceedings of the 7th IEEE Workshop on Mobile Computing Systems and Applications*.

REFERENCE

1. *Proc. 7th IEEE Workshop Mobile Computing Systems and Applications* (WMCSA 06), IEEE Press, 2006.

Maria R. Ebling is a research staff member at the IBM T.J. Watson Research Center, where she manages the Privacy-Enabled Context Technologies Department. She was the HotMobile 2006 program chair. Contact her at ebling@us.ibm.com.