MobiSys 2005: The Third International Conference on Mobile Systems, Applications, and Services

Seattle, WA
June 5–8, 2005

Summarizer shepherd: Maria R. Ebling

Tristan Henderson of Dartmouth College conducted a study of the conference WLAN during the conference. Questions regarding the study can be directed to him at tristan@cs.dartmouth.edu.

KEYNOTE ADDRESS

Summarized by Himanshu Raj

- ***Technology's Future: A Crisis in Confidence?***

Rick Rashid, Microsoft Research

Rick Rashid addressed the growing unrest among the CS community about the future of the IT industry and opportunities for research. He then outlined certain trends emerging from ubiquitous and pervasive domains and suggested that we are actually just getting started on research that will impact society at its core.

A terabyte of storage can hold every conversation you take part in for your entire life. It can hold one picture for every minute of your life. It can hold a year of everything you can see, in full-motion video. The SenseCam project at Microsoft Research is examining what you could do with a terabyte of personal storage. They have devised a wearable data and image recorder. The device contains a camera and numerous sensors (e.g., accelerometer, passive IR, light level, temperature, images, etc.). Image capture is triggered by sensing some interesting change in environment (e.g., lighting levels). They are hoping to apply this technology to patients with moderate memory impairment and as an aid for caregivers of patients with severe memory loss. It can also be used for self-reflection.

Rashid also gave examples of ubiquitous I/O, which turns any surface into an interactive computing surface, and of streaming intelligence, such as skyserver.sdss.org and skyquery.net, which allow scientists to do data mining against a large number of databases and has resulted in the discovery of new astronomical phenomena.

Rashid believes the goal of ubiquitous computing should be to bridge the gap between the rich and the poor. More than three billion people pay a poverty premium for basic goods and services. They have little access to important amenities and make decisions affecting their livelihood (e.g., when to plant crops) based on incomplete information. Even in the developed world, people with incomes below $35,000 have a 70% chance of not having good online access. Wiring to every home in the whole world is not realistic. Rashid described the Mesh Networking Project, which is applying wireless networking technology and mesh networks to form a cooperative network to cover the last hop. These networks must be autonomic, in that they should be self-managing, self-healing, and inexpensive. Other projects looking at this problem include the TIER project at Berkeley (http://tier.cs.berkeley.edu/) and the Digital Gangetic Plains project at IIT Kanpur (http://www.iitk.ac.in/mladgp/).

In conclusion, there are ample opportunities for reinvigorating CS research. The emphasis must be on access to information rather than mere access to devices.

During the Q&A session, Ramón Cáceres (IBM Research) asked about how to sustain technology change in rural areas. Rashid responded that you must find the economic value behind the technology and that, unless the person providing the information access points (such as Internet kiosks) makes money, he will not keep providing the service. Mandayam

Raghunath, also of IBM Research, followed up by asking how to get rural areas started. Rashid responded that there have been a number of approaches. One that worked in parts of India and Bangladesh was a bank that supports microloans. Tapan Parikh (University of Washington) pointed out that the positive effects we foresee from technology have to do with the applications that emerge from those technologies. Rashid agreed that this was true and then discussed whether this would cause the field of computer science to lose its identity. He pointed out that this style of research can bring risks from a career and from a funding perspective, but he also pointed out that about 50% of the basic research investment made by Microsoft goes toward projects that funding agencies are not (yet) comfortable with. He expressed hope that funding agencies would eventually view these cross-disciplinary projects as legitimate research areas.

APPLICATIONS ON THE GO

Summarized by Hongwei Zhang

- ***A Systems Architecture for Ubiquitous Video***

  Neil J. McCurdy and William G. Griswold, University of California, San Diego

  Neil McCurdy presented the architecture of their ubiquitous video system, RealityFlythrough, which enables users to visually explore a remote area in real time. McCurdy started by discussing the potential applications of the system, which include disaster response, remote monitoring, remote navigation, and virtual shopping.

  McCurdy introduced the technical challenges: the low density of camera deployment, the continuous movement of both the object and the camera, and the need for reliable, live, and real-time data delivery. He then elaborated on how he dealt with these technical challenges. RealityFlythrough uses "motion" as a substitute for an infinite number of cameras; it uses dynamic path estimation with smoothened transition between scenes; and it uses dynamic archiving of high-quality images from different locations to reduce the density of camera deployment and the amount of image data that needs to be delivered. He concluded the talk with a discussion of their evaluation. Overall, they found the number of live cameras to be the bottleneck, not the total number of cameras.

  After the talk, McCurdy discussed with the audience issues such as how to fetch archival images, how to deal with outdated images, the impact of GPS precision on system performance, and the network challenges posed by the possible use of omnidirectional cameras.

  More information is available at http://peanutgallery.homeip.net/drupal/taxonomy/term/1.

- ***LiveMail: Personalized Avatars for Mobile Entertainment***

  Miran Mosmondor, Ericsson Nikola Tesla; Tomislav Kosutic, KATE-KOM; Igor S. Pandzic, Zagreb University

  Miran Mosmondor presented LiveMail, a prototype system for communicating personalized images between mobile devices such as cell phones. He briefly introduced the concepts and techniques of 3D modeling, face animation, and 3D graphics on mobile platforms, then described the client-server architecture of LiveMail. The client enables the user to customize the characteristics of face animation, and the server acts as the relay between the sender and the receiver. Mosmondor stressed the importance of performing face animation using parameter adaptation, instead of transmitting raw face images: by transmitting the parameters instead of images, LiveMail only requires a bandwidth of 0.3Kbps, which is available even in resource resource-constrained devices such as cell phones. Mosmondor pointed out that the server of LiveMail spent 98% of the time on parameter adaptation yet only 2% of the time on storage and retrieval.

  Mosmondor answered questions regarding the relative importance of bandwidth and processing capability, the necessity of 3D rather than 2D face presentation, and the relation of LiveMail to other efforts. Mosmondor stressed that LiveMail does not require high bandwidth so much as high processing capability, such that the server could be dropped out of the architecture if the client (e.g., cell phone or PDA) is fast enough.

- ***MediaAlert—A Broadcast Video Monitoring and Alerting System for Mobile Users***

  Bin Wei, Bernard Renger, Yih-Farn Chen, Rittwik Jana, Huale Huang, Lee Begeja, David Gibbon, Zhu Liu, and Behzad Shahraray, AT&T Labs—Research

  Bin Wei presented MediaAlert, a system for delivering TV news according to the interests of users. Wei discussed the techniques employed in MediaAlert: processing media and extracting descriptive abstractions of the media, and content repurposing to disseminate media to devices with distinct protocols and capabilities. He stressed the importance of aligning unsynchronized images and captions and the scalability of the system to support a large number of users and devices.

  Wei walked through a scenario where a user created an identity and topics of interest, and MediaAlert then analyzed and delivered related images to the user. Wei also pointed out that media processing takes longer than dissemination and noted the importance of dealing with false positives in alert.

  David Kotz asked whether the user's location is used in the selection of content. Wei replied that it is and that the news content itself is location-dependent, because different news channels provide different

content. Mark Corner asked whether closed captioning was important to this system. Wei replied that closed captioning is very important to the existing system and that this is an area that will evolve. Tristan Henderson asked whether they had done any usability testing. Wei said that they had only used the system within their own team and that more usability testing was necessary. Andre Hesse asked what happens once the user has received the maximum number of alerts for the day and then something really important happens. Wei clarified that the maximum number of alerts per day is for the scheduled alerts and that the system sends the most relevant clips first. Real-time alerts cover emergency events and are not subject to the maximum.

SHAKE 'EM, BUT DON'T CRACK 'EM

Summarized by Neil McCurdy

- ***Cracking the Bluetooth PIN***

  Yaniv Shaked and Avishai Wool, Tel Aviv University

  This paper describes the implementation of an attack on the Bluetooth security mechanism and received considerable press coverage prior to MobiSys.

  Shaked presented results which show that a four-digit PIN (the PIN size on many commodity devices) can be cracked in less than 0.3 sec. on an old Pentium III 450MHz computer, and in 0.06 sec. on a Pentium IV 3GHz HT computer. Because the attack uses brute force, the time to crack larger PINs increases by a factor of 10 for each additional digit used in the PIN.

  The second contribution of this paper is the re-pairing attack, which forces two Bluetooth devices to rerun the pairing process that had just been described. Because devices usually store the link keys indefinitely, the first attack only works during the short interval when devices first connect to each other. However, an attacker can simulate a lost key by impersonating one device and telling the other device that it has forgotten its key. This action causes the two devices to run the full pairing process the next time they communicate, thus opening the devices to the attack. Shaked suspects that custom hardware would be required, because it may be necessary to spoof the Bluetooth ID and because the timing of the attack is critical.

  Shaked concluded the talk by suggesting possible countermeasures to this attack. One consequence of the re-pairing attack is that the user is asked to re-enter the PIN. Users should be wary of such requests. Users should also enter their PINs as infrequently as possible to reduce the risk of an attacker eavesdropping on the pairing process. He also suggests that the hardware manufacturers use the 128-bit PINs that are allowed in the standard to make the brute-force attack less likely to succeed. Even when using a 128-bit PIN, though, a denial of service attack could still be performed by constantly running the re-pairing attack.

- ***Shake Them Up! A Movement-Based Pairing Protocol for CPU-Constrained Devices***

  Claude Castelluccia, INRIA, France and University of California, Irvine; Pars Mutaf, INRIA, France

  This talk complemented the Bluetooth cracking talk. Claude Castelluccia showed how secure authentication between sensor devices could be performed across a clear channel without using a PIN and without using anything as CPU-intensive as a public key. Other goals were: no preconfiguration, no extra cost, no special equipment, and protection from denial of service attacks.

  The strategy assumes that an eavesdropper cannot tell the source of a message (source anonymity). Assuming you have two devices, Alice and Bob, Alice can send a message to Bob that says, "I'm Alice." Bob knows this is true, since he's not Alice, so he and Alice can agree that the bit is 1. If Alice instead had said, "I'm Bob," Bob would know that this is false, so they would both agree on the bit being 0. Eve hears these messages, but does not know who was the source and therefore cannot determine the values of the bits.

  The challenge is to construct an environment where Eve cannot determine the source of a message. In this talk, Castelluccia discussed three approaches Eve could take to discover the source of a message. She could use timing information, signal power, or frequency. He then described ways that Alice and Bob could protect against each of these techniques, one of which required physically "shaking them up."

  Someone suggested that an attacker could intentionally create a collision when Alice sends a message to Bob, so Alice and Bob would not agree. Castelluccia pointed out that this would not work, since 802.11 is reliable: Alice would expect an ACK from Bob. Another questioner wondered whether cameras could be used to record the positions of the devices as they were being shaken. The question was mostly asked in jest, but Castelluccia pointed out that, even if such an attack were feasible, the devices could be obscured from view while being shaken.

MOBILE SERVICES

Summarized by Mike Blackstock

- ***Reincarnating PCs with Portable SoulPads***

  Ramón Cáceres, Casey Carter, Chandra Narayanaswami, and Mandayam Raghunath, IBM T.J. Watson Research Center

  Awarded Best Paper!

  Mandayam Raghunath presented SoulPad, a new approach to solving the problem of providing a personal and customized computing environment anywhere. When users

move locations, they want to suspend their work and then start up again exactly where they left off. In the SoulPad solution, the PC (body) is separated from the session state and bits (soul). SoulPad uses a pocket-sized USB drive that can be plugged into any PC to allow users to resume their personalized computing environment. The key enablers for SoulPad are fast, small, high-capacity USB 2.0 drives from mass market media players, auto-configuring OSes such as Knoppix, and mature virtualization technologies.

Although the virtual machine and swap space use an encrypted file system, SoulPad is still currently vulnerable to hardware and BIOS attacks. To address reliability, the system supports network backups, though local backup is also possible. To deal with hardware diversity, the system needs to keep many drivers up to date. Some older systems cannot boot from a USB dongle, so a small boot CD could be required.

Raghunath was asked whether there are other issues that need to be addressed to put SoulPads to use. One issue is that certain applications check the hardware and assume that they are pirated when they have been moved to a new PC. There are also legal issues that need to be addressed.

Further information is available at http://www.research.ibm.com/WearableComputing/SoulPad/soulpad.html.

- ***Slingshot: Deploying Stateful Services in Wireless Hotspots***

  Ya-Yunn Su and Jason Flinn, University of Michigan

  Ya-Yunn Su presented the Slingshot system, which alleviates the bottleneck associated with the back-haul connection to hotspots by replicating application state to surrogate computers closer to wireless access points. She presented the scenario where a user brings a Pocket PC into a coffee shop and uses VNC to execute applications on a server over the Internet. The bandwidth over the Internet is often limited to a T1 (1.5Mbps) and latency is high, leading to poor response to desktop interaction. In a manner similar to cyber-foraging, Slingshot replicates state from the remote (home) computing environment to a VM running on one or more surrogates. Because Slingshot only replicates the state, users can fall back to their home server in the event of surrogate failure. Once a replica is instantiated on one or more surrogates, a proxy broadcasts user interaction to all replicas and the home server. The applications tested are the VNC desktop and speech recognition, though Su only discussed the VNC application in her presentation.

  The system was evaluated to show that once the state was transferred, taking 27 minutes, the interaction performance was 2.6 times faster. Using a microdrive to store volatile state and chunk hashes sped things up considerably, requiring only 6 minutes: about 3 minutes to transfer state, another 3 to replay the logs. Overall the system improves performance for low-latency applications and hides surrogate failures by using replicas and broadcasting the interactions.

  Ramón Cáceres asked if the system relied on the applications running in a deterministic manner; for example, could a Slingshot application maintain communications from the outside world? Su responded that determinism is required and that outside communication is not yet supported. Another member of the audience asked whether requiring VMware on all surrogates might perhaps be too strong an assumption. Su responded that they do assume that the surrogate is already running VMware, but suggested that perhaps VMware itself could be transferred first in the future.

- ***DeltaCast: Efficient File Reconciliation in Wireless Broadcast Systems***

  Julian Chesterfield, University of Cambridge; Pablo Rodriguez, Microsoft Research

  Julian Chesterfield presented the DeltaCast system, which very efficiently reconciles two versions of a file between a server and any number of clients with any version of a file using a pure radio broadcast system. DeltaCast uses a hierarchical hashing scheme, combined with decomposable hashes and erasure coding for high efficiency. DeltaCast was compared with file download, flat hash, hierarchical hash schemes using Web pages, and binary data such as software upgrades. From this evaluation, it was determined that the number of hierarchy levels could be dynamic depending on the data type. Compared to hierarchical and single-layer hash systems, the time required to get the data required to update a file on a client is also much lower for Web pages. The amount of data downloaded is the same as in hierarchical hash schemes and less than in single-layer hashing. DeltaCast trades off decoding time, but the overall penalty is low. This system can be applied to not only broadcast radio networks but also IP multicast, overlay networks, and content distribution networks.

  One attendee asked why the authors did not consider the use of Turbocodes or other well-known hash algorithms. Apparently these hashes do not require the system to solve linear equations, but use iterative algorithms, so they should be faster. Chesterfield answered that they used a hash that was already available to them and understand that perhaps their system may not be optimal in this area. Maria Ebling asked whether the carousel size of hashes or data on a given channel was fixed. Chesterfield answered that any number of erasure codes can be generated and any of

these blocks/codes is useful in regenerating data and hash codes.

POSTERS AND DEMONSTRATIONS

Summarized by Denitsa Tilkidjieva

This session hosted 24 displays of researchers' work in the field of mobile systems, applications, and services. A wide range of applications were on display. We highlight just three of these here:

The pre-hospital patient care system (Hashmi et al.) consisted of a number of pulse sensors, attached to patients' fingers to monitor vital signs. It is most useful when multiple patients are in need of attention, because it can allow quick decisions at critical moments.

The inHand system for ubiquitous personalized interactive content (Bhatti et al.) was shown in both a poster and a demonstration. The researchers demonstrated the inHand device and how it can be used to gather user-customizable information about movies, events, and the like.

If you try to cook a turkey, a duck, and a chicken all in one, you will get a Turducken. Sorber and his colleagues borrowed this name for their system for hierarchical power management, consisting of a laptop, a stripped-down PDA, and a mote. The device always chooses the platform that performed the given task at the lowest energy cost, thus extending the lifetime of the device up to 10 times.

Photos of the posters and demonstrations can be found at http://www.mtholyoke.edu/~dntilkid/mobisys.

PLENARY SESSION

Summarized by Hongwei Zhang

- ***Staying Off the Hot Seat with Cool Mobile Systems***

Alfred Spector, IBM Software

Alfred Spector, chief technology officer at IBM Software, gave a visionary talk on the technological and societal implications of the development of mobile pervasive systems, with an emphasis on the robustness of these systems.

Spector first described the trend of mobile systems. On one hand, the ever-growing modality, decreasing form factor, declining cost, and exploding connectivity have been pushing the development of mobile systems. On the other hand, medical informatics and security applications are calling for mobile systems. For instance, the market for health care is about $1.5 trillion per year, which is about as large as the whole IT industry. Nevertheless, in most scenarios of mobile pervasive systems, we are envisioning amalgams of components. The systems are usually so complex that we cannot prove their correctness. In many cases, it is also impossible to anticipate what may go wrong in a system.

Given the complexity of pervasive systems, Spector argued that one key challenge of designing these systems is to guarantee their robustness, e.g., ease of use, evolution, QoS, reliability, security, and fitness to purpose.

After discussing the challenges posed by mobile pervasive systems, Spector analyzed why existing techniques and architectures do not satisfy the need for robustness. To demonstrate design for robustness, Spector described a hierarchical architecture based on a common trusted computing base, a secure hypervisor, and trusted virtual domains. Spector stressed the importance of adopting a top-down approach and the development of trustworthy capabilities (e.g., attestation, privacy services, and authentication).

In particular, Spector stressed the concept of "information provenance (InfoP)," by which the origin of information can be identified with proof. InfoP can provide the basis for law enforcement to play their role in improving the robustness of systems (i.e., punishing actions that are not allowed, such as spam, viruses, etc). But Spector also pointed out the challenges of InfoP: privacy, storage for massive data, digital signatures and CAs, and law enforcement across international boundaries.

Spector discussed the importance of the "science of design" in increasing system robustness. Besides the traditional time and space complexity, Spector specifically noted the need to manage the complexity of usage (i.e., user interface). To this end, he argued that we should pay attention to the meaning and measurement of robustness and of design methodologies. He also emphasized the importance of simple yet flexible architectures, as well as self-healing and self-optimization. Finally, he suggested that the technical community should work with the wider society to address a broader range of robustness issues.

SPEEDY WIRELESS

Summarized by Xiaoqiao (George) Meng

- ***Improving TCP Performance over Wireless Networks with Collaborative Multi-Homed Mobile Hosts***

Kyu-Han Kim and Kang G. Shin, University of Michigan

Kyu-Han Kim pointed out that wireless networks have capacity limitations. In practice, a group of nearby hosts may constitute a mobile collaborative community (MC2) where hosts share bandwidth and content. In such a community, each host is multi-homed and data is multiplexed to improve utilization. Accordingly, it is important to allow TCP to achieve high utilization of the aggregate bandwidth over multiple interfaces, but this presents several challenges, from requiring exact link-state information, to coping with dynamic wireless links, to handling out-of-order packet delivery, to controlling congestion.

To cope with these challenges, Kim introduced PRISM, Proxy-based Inverse Multiplexer. PRISM consists of adaptive scheduling, intelligent acknowledgment controlling, network-assisted fast recovery, and IMUX at the proxy's network layer. PRISM architecture has an adaptive scheduler (ADAS) to achieve cheap and adaptive fair-scheduling. PRISM has been implemented in Linux kernel 2.4.20 and netfilter. PRISM-IMUX is a filter at a network layer. The authors also devised a testbed, by which they found that PRISM delivers 95% of the aggregated bandwidth. In a setting with three heterogeneous mobile nodes, PRISM achieved more bandwidth than vanilla TCP.

David Kotz asked how members of the community find one another. Kim answered that they can use the service location protocol to find existing collaborative communities, but that forming communities and dealing with membership dynamics is an area for future work.

- ***Horde: Separating Network Striping Policy from Mechanism***

  Asfandyar Qureshi and John Guttag, MIT Computer Science and AI Laboratory

  Asfandyar Qureshi presented Horde, a middleware mechanism that allows multi-stream applications to communicate over multiple channels with widely varying latency and bandwidth. The authors were motivated to build Horde in order to support mobile telemedicine, specifically an application that allows doctors to examine patients in transit to the ER. This application requires the transmission of unidirectional video, bidirectional audio, and low-rate physiological data streams in real time from a moving ambulance. This system requires high throughput, low latency, and the ability to deal with vehicular motion in an urban area.

  Network striping in a WWAN environment presents substantial challenges, from coping with limited bandwidth, to dealing with applications having dissimilar needs, to dealing with dissimilar network channels with varying QoSes.

  Horde middleware provides a number of services, the most novel of which is QoS modulation. With QoS modulations, applications express stream QoS sensitivities, where the QoS is multi-dimensional. When an application sends data, it receives some utility from the consumption of its data at another host. Applications express QoS "objectives" which define QoS constraints on streams. In summary, the goal for Horde is to build a flexible network striping middleware for WWANs.

  David Kotz asked how many applications the authors have examined and how many more they plan to explore. Qureshi expressed concern that the interface may be too rich, in that you may be able to express more than is necessary. They are building the telemedicine application to gain more experience with the system.

- ***An Overlay MAC Layer for 802.11 Networks***

  Ananth Rao and Ion Stoica, University of California, Berkeley

  Multi-hop wireless networks are being considered for last-mile broadband connectivity. However, such networks are subject to issues of fairness. This work addresses this issue; specifically, the authors show how to prevent starvation of flows without changing the hardware.

  Ananth Rao first described a starvation problem with 802.11 MAC that they identified by using a testbed with six nodes. They found that the cause of this problem was interference and asymmetric carrier sense. Rao pointed out that existing solutions to this problem (e.g., contention-based MACs and TDMA) require hardware modifications but that the starvation problem can also be solved above the MAC layer by limiting how much data each host can send. The proposed solution is called Overlay MAC Layer (OML). OML uses readily available, inexpensive hardware, and can evolve to meet diverse application scenarios/requirements. Rao then described how to implement OML. He summarized the results of their evaluation on both a six-node testbed (based on Click from MIT) and a Qualnet network simulator. They found that the disparity between one-hop and four-hop flows was reduced but that the throughput on the one-hop flows was also greatly reduced.

  Himanshu Raj asked whether providing fairness reduced the overall throughput. Rao explained that this approach actually increases the throughput, though it is not shown in the graph because the graph does not account for starved flows. Ramón Cáceres pointed out that with an asymmetric link, one node would never be in the active list, but Rao clarified that a node in the middle can piggyback information about the other node so that it will be added to the active list. Richard Paine encouraged Rao to propose some of these changes to 802.11 because, although 802.11n addresses some of these issues, it would be useful to address the other questions as well. Brian Noble asked what happens if one of the nodes gets greedy, and Rao explained that they currently assume that the box is tamper-proof.

OPERATING SYSTEMS FOR SENSOR NETWORKS

Summarized by David Johnson

- ***Design and Implementation of a Single-System Image Operating System for Ad Hoc Networks***

  Hongzhou Liu, Tom Roeder, Kevin Walsh, Rimon Barr, and Emin Gün Sirer, Cornell University

  Hongzhou Liu presented MagnetOS, a distributed operating system designed for use in ad hoc networks. Liu observed that ad hoc

network applications are extremely difficult to program, even today. MagnetOS responds to this problem by combining all network nodes into a single, event-based virtual machine; this abstraction eases application development and increases network lifetime. In MagnetOS, synchronous and asynchronous events signal code execution by triggering event handlers. A static partitioning approach converts application class files into event handlers. Migration of event handlers provides improved energy efficiency and saving of computation if loss of power is imminent. MagnetOS provides several migration algorithms that are designed to minimize communication energy overhead.

Liu presented results from evaluating MagnetOS with an application called SenseNet, in which there are a number of fixed sensors and mobile data processing components. They compared a number of algorithms, three that required no communication overhead and two that did. The TopoCenter(1) algorithm, which moves objects using one-hop neighborhood knowledge from each communicating partner, increased system lifetime by a factor of 2.5.

Ahmad Al-Hammouri asked why they did not use aglets, since aglets provide a clean environment for mobile agents. Liu explained that the major complaint about mobile agents is security, because they can execute any code on any node. Doug Terry added that MagnetOS is using a different model; they are breaking their code onto different machines. Terry then asked whether they thought they would need multiple algorithms and adapt between them dynamically or whether one would be sufficient. Liu responded that they currently require the programmer to pick one and that their experience to date suggests that the two Topo algorithms perform well throughout the range.

- ***SOS: A Dynamic Operating System for Sensor Nodes***

  Chih-Chieh Han, Ram Kumar, Roy Shea, Eddie Kohler, and Mani Srivastava, University of California, Los Angeles

  Ram Kumar presented SOS, an operating system designed for use in sensor network applications. Because sensor networks often manifest in long-term deployments, individual nodes must be flexible to respond to remotely controlled changes. SOS is a modular, application-independent operating system that supports dynamic reprogramming via module updates and replacements. In contrast, TinyOS, the de facto sensor network operating system, produces a static binary composed of both system and application-level functionality and must be recompiled and replaced on each node to effect an upgrade. Another similar system, Maté, provides a virtual machine that can execute small code fragments distributed through the network. Application-level upgrades are possible, but interpreter upgrades must fall back on the TinyOS update system. SOS consists of a static kernel, which provides an abstraction of the hardware, and is installed on all nodes. Dynamically linked modules communicate with various kernel services and device drivers as necessary. Modules can register functions with the kernel, and potential callers may subscribe to provided functions. The kernel provides priority scheduling, dynamic memory and intra-module message passing, as well as safety features.

  Kumar discussed the results of an evaluation comparing TinyOS, SOS, and Maté. All ran Surge-like applications (Surge is a well-known multi-hop data-gathering application for TinyOS). They found that CPU activity was on average 1% higher in SOS than in TinyOS. Energy usage during code updates in SOS was an order of magnitude larger than in Maté, because Maté must update only a small chunk of bytecode. However, TinyOS update costs were 400 times as much as for SOS, because updates to TinyOS require the replacement of the entire system binary. Kumar observed that the important metric to observe when comparing update energy costs is the frequency of updates.

  Himanshu Raj asked what impact one module can have on another and whether they can crash one another. Kumar responded that the architecture provides no form of memory protection and that a wild pointer could corrupt the data space of another module. Unfortunately, there is no hardware support to correct this problem. Bhanu Pisupati asked about the programming model. Kumar responded that it is described in more detail in the paper, but that all modules are implemented as message handlers and listen on a particular port for messages intended for them. Jason Flinn found the idea of safety checks fairly interesting, but asked for clarification about how this really works. Kumar responded that static checks would require analyzing the entire source code and would not solve the problem. He clarified that the base stations maintain some information about the kind of modules present on the nodes and do some analysis at load time before sending a module.

LOCATION (HERE)

Summarized by Neil McCurdy

- ***A Relative Positioning System for Co-Located Mobile Devices***

  Mike Hazas, Christian Kray, Hans Gellersen, Henoc Agbota, and Gerd Kortuem, Lancaster University, U.K.; Albert Krohn, University of Karlsruhe, Germany

  Mike Hazas introduced Relate, a compelling approach to determining fine-grained relative locations and orientations between proximate devices. It uses ultrasound peer-to-peer sensing, giving the user relative location accuracy in the 10cm range and eliminating the need for the infrastructure-based

location support that is typically needed for such fine-grained accuracy. The Relate team created their own hardware—a USB dongle that has three ultrasound transmitters and sensors. Time of flight of the ultrasound signal determines the range, and the signal strength, as recorded by each of the three sensors, determines the angle of arrival.

To evaluate the system, Hazas described an experiment in which five laptops were placed at various positions on a 2.4 x 1.6m surface. One hundred iterations of this experiment were performed, with half of them ensuring that the dongles had line-of-sight to one another. With good line-of-sight, one can expect roughly 9 cm, 33° accuracy. With limited line of sight, one can expect approximately 11cm, 48° accuracy. Hazas closed by discussing some of the issues with this approach. First, he posited that hardware that is better equipped to do signal processing may be able to get the accuracy to as low as 2 or 3cm. He also pointed out that there is a limit to the number of devices that can be handled, since only one device can communicate at a time.

The Relate system was demonstrated on Monday night, and Hazas was asked how his system differed from Cricket. Relate dongles have three transducers, are optimized for co-planar calculations, and can calculate the angle of arrival to determine the orientation of the devices. The Relate system also does not require any calibration.

One questioner pointed out that location accuracy is more important as the devices get closer to one another. For example, 10cm accuracy may not be adequate when the devices are only 2cm apart. Hazas agreed. There was also a question about whether the signal processing could be done in software. Hazas did not see any reason why not.

- ***WALRUS: Wireless Acoustic Location with Room-Level Resolution Using Ultrasound***

  Gaetano Borriello, University of Washington and Intel Research; Alan Liu, Tony Offer, and Christopher Palistrant, University of Washington; Richard Sharp, Intel Research Cambridge, UK

  Gaetano Borriello presented a system that provides room-level granularity by using a combination of wireless and ultrasound. Borriello began by explaining the importance of room-level localization to usability. The goals of the system included low deployment cost, low support cost, and a system that was incrementally useful and deployable. The system should also approach 100% accuracy while maintaining privacy.

  Borriello described WALRUS as being inspired by lightning and thunder, with WiFi (through an access point broadcast of information) acting as the lightning and ultrasound (through commodity speakers) acting as the thunder. Clients receiving the WiFi packet keep the packet only if they also detect the sound. The system was tested in two environments under many different conditions (music playing, conversations ongoing, keys jangling, doors slammed, etc.), and the system proved largely immune to extraneous noise. Borriello concluded by discussing some of the limitations in the existing implementation and presenting a vision of future prototypes that might one day be located in a store near you: a wristwatch receiving device, with light bulb and air freshener transmitting devices.

  Lin Zhong asked whether they had considered listening to the ultrasound signal first and then the WiFi. Borriello responded that this approach might be perfectly reasonable, but that they haven't tried it. Ed Nightingale asked whether they had considered rooms that were changing. Borriello responded that you could put arbitrary information in the packets. David Kotz commented that one of the unique characteristics of this work was that they were using existing devices. He then asked why they decided to move toward custom hardware. Borriello replied that one thing to consider is how much optimization each device could have. Robert Hall asked whether they had considered the health implications of ultrasound. Borriello said that they are not boosting the speaker output and that one of the reasons to move toward custom hardware was to move further away from the audible range. Another member of the audience asked whether they had considered placing microphones in the room to measure the volume levels and then including those volume levels in the WiFi packets so that clients would know what volume to expect. Borriello responded that they had walked around the room and measured the volume levels, but they found that this was not important because their detection algorithm relied on a relative measure of strength.

WORK-IN-PROGRESS REPORTS (WIPS)

Summarized by Ya-Yunn Su

- ***CAM: Architecture for Automating Paper-Based Processes in the Developing World***

  Tapan Parikh, University of Washington

  Tapan Parikh said that in developing countries paper-based processes are inefficient, but cell phone penetration is growing, and most cell phones have built-in cameras. He proposed using such mobile phones to digitize manual, paper-driven processes. A user could take a picture of a paper document, transfer the data to a remote machine, and propagate it to the appropriate recipient. The receiver could then print the document.

  http://www.cs.washington.edu/homes/tapan

- ***Smart Attire–The Digital Diary***

  Tarek Abdelzaher, University of Virginia

  The author proposed building sensors into clothing to record user activity. The smart attire includes accelerometers, magnetometers, and temperature, sound, light, and GPS sensors. Possible applications are personal digital diaries, and medical monitors. The prototype system is a winter jacket with five motes and a GPS tracking device. The motes on the jacket record activities and upload the data when near an access station using 802.15.4. One example shows that the author can infer the user's activity (e.g., typing, walking) from the data collected by accelerometers. Another example shows the user's path or stillness by data from the GPS device. The user may use data mining techniques to understand personal life patterns (e.g., Is my social life declining?) and query the history to keep track of health or financial activities (e.g., When is the last time I visited my dentist?). The prototype system shows that power management is important for smart attire.

- ***Extracting a Mobility Model from Network Traces***

  Minkyong Kim, Dartmouth College

  Researchers need a realistic mobility model to simulate the effectiveness of new algorithms on wireless networks. There are two ways to generate network traces: syslog and GPS. Syslog data collected on access points contains client events, including time and action. It has the advantage of availability of large data sets, but it is often hard to estimate a user's location from an AP location, perhaps due to the device not being associated to the closest AP or perhaps because of some device-specific implementation. An alternative is GPS. Unfortunately, there is no GPS data set publicly available, and it does not work indoors. To address these limitations, Kim extracts the user's path from time and AP location from syslog data, extracts mobility characteristics from the user's path, and extracts speed/duration/pause to generate the mobility model.

- ***Emulab Unleashed! Into the WiFi and Mobile Wireless Dimensions***

  David Johnson, University of Utah

  David Johnson presented Emulab, a network emulation testbed. Due to the complex nature of wireless networks, simulation is not enough in testing and validating new research ideas. A real testbed like Emulab would be a better way. Emulab has added three wireless testbeds including a building-scale WiFi testbed, a fixed sensor net, and a mobile WiFi network. The WiFi testbed enables remote hardware reset functionality, since it is difficult to find the real location and reboot the machine manually. There are also fully programmable motes and mobile robots with a tracking system accurate to 1cm. Emulab can be remotely accessed and allows multiple users. It provides a realistic and repeatable wireless environment. Future work includes automated rechargeable stations and power monitoring. More information can be found at http://www.emulab.net.

- ***Content Management for Mobile and Pervasive Experiences***

  Nigel Davies, Lancaster University

  Nigel Davies emphasized the importance of content in a ubicom/pervasive computing environment based on experience in GUIDE deployment. Their solution was to assign many students to work on content and the user interface. Other projects, such as Can You See Me Now? (http://www.canyouseemenow.co.uk) and Equator (http://www.equator.ac.uk/) reached similar conclusions. In the e-Campus project, content can come from users, be automatically generated, etc. The challenge is how to manage multimedia content in future mobile and ubiquitous computing environments. More information on the GUIDE project is located at http://www.guide.lancs.ac.uk/.

- ***Invisible Agents***

  Nobuo Kawaguchi and Negishi Yuuya, Nagoya University

  Nobuo Kawaguchi pointed out that in a world with many computers (PC, laptop, PDAs, etc.) and smart appliances (music players, TiVo), we do not have a good interface to control and aggregate the interfaces of all the devices. They built invisible agents to solve this problem. The example is a conference room with a projector, multiple monitors, ceiling lights, and similar equipment. They can combine brightness with a human sensor to control the ceiling light, or a human sensor to activate the projector. On each of the target devices, they run a VNC server. The master device connects to all the target devices.

- ***A Social Networking Web Site for the Research Community***

  James Scott and Richard Sharp, Intel Cambridge

  James Scott said that the population of researchers is huge and growing. There are many types of relationships between researchers (e.g., as co-authors, work colleagues, conference attendees), but communities that overlap might also be isolated (e.g., SIGMOBILE, Pervasive/UbiComp, SIGCOMM). One useful feature of a social networking Web site would be an easy home-page generator containing basic information (e.g., bio, publications, photos). To prevent spamming, the Web site could be semi-exclusive or could follow Gmail's invitation-based model. The Web site would encourage people to register and verify their information.

- ***Crawdad—A Community Resource for Archiving Wireless Data at Dartmouth***

  David Kotz, Dartmouth College

  David Kotz pointed out that there is little real wireless network traffic available for researchers. They initiated Crawdad as a facility for storing data sets collected from real

wireless networks. They already had a campuswide wireless infrastructure for collecting traces. The challenges include developing a common data format, importing existing data, and anonymizing data. They hope to coordinate with other communities to develop network trace formats and tools. The tools and data can also be used in course projects. More information is available at http://crawdad.cs.dartmouth.edu/.

- ***Secure Mobile Architecture***

  Richard Paine, Boeing Technology

  Richard Paine proposed a secure mobile architecture that can cryptographically identify each packet. They can support mobility by transparently changing the address for the user and application. This framework improves their enterprise network by reducing operational cost and complexity. They use four techniques to achieve their goal:

  - 1. PKI: Public key infrastructure. Each client uses his/her badge for client authentication.
  - 2. HIP: Host mobility protocol. Communications are based on IPSec, therefore each host is identified by a security parameter index (SPI) rather than IP. Each host is further identified by a host identity tag (HIT), which is SHA-1 of the public key.
  - 3. NDS: Network directory services. The client goes through an identification process before using LDAP.
  - 4. LENS: Location-enabled network services: see http://www.opengroup.org/bookstore/catalog/select.tpl?text=secure+mobile+architecture.

- ***GSM War Drive***

  Mike Chen, Intel Seattle

  Mike Chen presented their goal of providing a playground for location-based services. Cell phones are the location devices people already carry every day. Some current location technologies are GPS and WiFi. GPS devices are not accurate and have limited coverage. WiFi can only be used on limited devices. Based on calculations from GSM tower signal strengths, the cell phone can provide accuracy within 30cm outdoors and 4m accuracy with a 1m grid. They plan to make the trace publicly available.

LOCATION (THERE)

Summarized by Mike Blackstock

- ***The Horus WLAN Location Determination System***

  Moustafa Youssef and Ashok Agrawala, University of Maryland

  Moustafa Youssef presented the Horus system used to determine indoor locations to an accuracy of less than 2m by using existing wireless LAN infrastructure. Horus, like other WLAN-based locating systems, uses APs as reference points and the observed signal strength to these APs to estimate distance via triangulation. However, when indoors, the observed signal strength readers can differ by 15dBm for a given distance. Like Radar, Horus uses a radio map to characterize the area to counter these effects. Unlike Radar, however, Horus is a probabilistic system rather than deterministic.

  The goals of this system were high accuracy, low computational requirements, energy efficiency, and scalability (both in number of users and in the area covered). The Horus techniques accounted for a 25% reduction in the average distance error. The Horus system is shown to have higher accuracy on average than Radar by more than 82% and better than 27% for the probabilistic system, and is more computationally efficient by an order of magnitude. The authors then applied the Horus ideas in Radar and showed that these ideas could reduce Radar's average distance error by more than 58% and decrease the worst-case error by more than 78%.

  Robert Harle asked what the test environments were like. Youssef indicated that it was a typical CS department with offline measurements taken at night and used during the day to capture environment variations in a typical deployment. Mark Corner asked how the new dynamic power control feature would affect location determination using the new WLAN systems. Youssef thought that their ongoing work in automatically generating the radio maps based on environment changes may be effective in dealing with this problem.

- ***Deploying and Evaluating a Location-Aware System***

  R.K. Harle and A. Hopper, University of Cambridge, U.K.

  Robert Harle presented their experience with deploying and using the Active Bat system at Cambridge. The Bat system is accurate to 3 to 5cm in three dimensions. The data on which this study is based was collected over a period of more than two years from a deployment that covers about 500 square meters in their new building.

  Harle found that the killer application for this type of system was allowing companies to learn more about how office building space is used, to encourage people to work more effectively. Surveys showed that privacy was not an issue, but that result may not extend beyond small communities such as theirs. A more meaningful study of such systems would require deployment in a corporate office or perhaps a hospital where there is more overlap in working hours and collaboration is more common.

  James Scott thought it was interesting that the corporate space usage was the best application rather than an application that benefited end users. Harle indicated that in fact one company expected to get a return on their investment in one year of data collection, without using any of the software that could

benefit the end user. Further, they were not interested in deploying end-user applications until they had reaped that ROI. Minkyong Kim asked whether providing more information to the end user regarding the accuracy of a location reading would increase system trust. Harle indicated that they experimented with a five-bar system ranking and found it useful, but that this ranking system did not help this issue significantly. Guanling Chen asked whether the high accuracy of the Bat system was necessary for the applications in this deployment. Harle indicated that although it was not needed for this type of application, there are classes of applications, such as those that use position clicking and pointing, that require it. He agreed that the most important thing is highly accurate room-level granularity.

- ***Accuracy Characterization for Metropolitan-Scale WiFi Localization***

  Yu-Chung Cheng, Intel Research Seattle and University of California, San Diego; Yatin Chawathe and Anthony LaMarca, Intel Research Seattle; John Krumm, Microsoft Research

  While an intern at Intel Research, Yu-Chung Cheng and his colleagues worked to characterize the accuracy of the Place Lab WLAN location determination systems for use outdoors. Unlike GPS, their system works in urban canyons and indoor environments, and it relies on more ubiquitous technology (check out the density map for Manhattan at http://www.wigle.net). Unlike other WLAN systems, their system requires less configuration time (1km^2 area/1 hour) using war driving. Although it is consequently less accurate (13–40m), this is not an issue for applications such as a location-based Web search.

  For their experiment they gathered data in three neighborhoods—downtown Seattle, an urban residential area, and an area in Kirkland (a less dense suburb where homes are 15 to 20 feet apart)—and compared their location estimates with GPS readings.

  They then applied three different algorithms. Their baseline tests found that the specific algorithm used did not matter much, though fingerprints performed poorly with only one AP. Interestingly, errors were higher in the more dense downtown area, probably owing to the fact that many APs are higher up and not contributing to making measurements more accurate. They concluded that it is possible to get acceptable accuracy of 13–20m in high-density areas, and around 40m in lower-density areas, with about 30 to 60 minutes of calibration for the neighborhood.

  David Kotz asked about the data corresponding to Figure 4, where the Y axis is labeled "% of Time"; Cheng clarified that it actually represents "% of records." Another member of the audience was asked whether the urban results were affected by the GPS noise in urban canyons. Cheng said that they countered this possible effect by checking that readings were correlated by three GPS units and only using those that were consistent. David Kotz then asked whether it would be possible to improve accuracy with the indoor Horus techniques. Youssef and Cheng discussed this possibility offline. Youssef later reported that they concluded that the Horus techniques would be useful when there is more information available for the localization algorithm. For example, when the number of APs per scan increases, Horus techniques can give a significant advantage. However, where you have just one AP per scan, there is not enough information to notice a difference between the techniques.

  The traces and source code for this paper are available at http://www.placelab.org/.

MORE POWER TO YOU

Summarized by Ram Kumar Rengaswamy

- ***Energy Efficiency of Handheld Computer Interfaces: Limits, Characterization, and Practice***

  Lin Zhong and Niraj K. Jha, Princeton University

  Lin Zhong noted that the role of the user interface has often been ignored in the design of energy-efficient systems. The speed of the human interaction is significantly lower than the speed of the computer. The computer ends up spending significant time waiting for user inputs. The energy consumed by the computer in these idle periods can be eliminated through better system design.

  Zhong and Jha compared two forms of user input, speech and handwriting, using energy efficiency as the metric. Energy efficiency is a combination of the speed of the input and its power efficiency. The experiments showed that speech is more energy efficient than handwriting. Zhong also designed a wireless wrist-watch to be used as a low-power, low-cost cache device. The cache device is a slave to the main computer. It collects user input while the main computer is put to sleep. This results in system-level power savings.

  The results evoked a lot of interest from the audience. Brian Noble raised an interesting point about the usage scenarios of the input techniques. Listening and writing are concurrent operations while speaking and listening are not. Hence, from a holistic viewpoint, it might be more energy efficient to listen and write than to listen and speak.

■ ***Turducken: Hierarchical Power Management for Mobile Devices***

Jacob Sorber, Nilanjan Banerjee, and Mark D. Corner, University of Massachusetts; Sami Rollins, Mount Holyoke College

Jacob Sorber explained that the main principle of hierarchical power management is to pick the most energy-efficient component of the system for a task. The challenge is in partitioning a given task into a set of subtasks and assigning each to the most efficient component. Such an approach automatically maximizes system lifetime. It is desirable to have minimum user intervention in such a system.

The authors developed the Turducken system for hierarchical power management in laptops. Turducken consists of a laptop attached to a PDA and a mote sensor node. The role of the mote is to maintain clock synchronization with a time server. The laptop and the PDA derive their clock from the mote upon their wakeup. The PDA is responsible for caching Web pages and waking up the laptop to display the pages once they are fully loaded. The user interacts directly with the laptop. The laptop responds to user queries (e.g., email retrieval). Turducken significantly lowers average power consumption compared to conventional systems.

The audience provided some very interesting comments and suggestions. The power supply design of such a system was discussed. It was concluded that it would be most energy efficient to have separate batteries for each system component. Usage of the lower-tier components during the transition interval of a high-tier device from one state to another was considered.