6th IEEE Workshop on Mobile Computing Systems and Applications (WMCSA 2004)

English Lake District, UK
December 2–3, 2004
http://wmcsa2004.lancs.ac.uk

Summarized by Adrian Friday, Mary Baker, Sachin Goyal, and Fahd Al Bin Ali

The WMCSA series of workshops was started in 1994 to explore a new and extremely active emerging area of research centered on the advent of truly portable devices and wireless telecommunications. Explicitly seeking to create an open and interactive atmosphere, the workshop attendance was capped at 70; all attendees had to submit either a long or a short paper to ensure that all participants were active in the field, and the program provided frequent opportunities for interaction, including panel sessions and breakout groups along with more formal paper presentations.

Ten years later, the field now mature and having spawned a whole series of high-quality conferences and workshops, we witnessed a still thriving research community, although arguably now working on different, more wide-ranging problems. The workshop, now in its sixth generation, still retains many of the core attributes that galvanized a community a decade ago: small numbers, rigorous peer review ensuring that we accept the best-quality work, and plenty of breaks to allow people to engage naturally with one another.

Opening Plenary Discussion

It was with this legacy in mind that we set out to open the workshop in a reflective mood—*what impact had we, as a community, had over the last decade?* Rather than a traditional keynote address, we began the workshop as we hoped it would continue, in discussion. Despite months of careful planning, we decided the evening before that the opening panel was likely to be too didactic and instill too much of a presentation culture for the remainder of the workshop. The opening panel was replaced with a group discussion involving the entire plenary. Each delegate was asked to join a small group of four to five nearby participants and to reflect on the following questions:

1. What impact have we had? Has any of the work we've done been adopted and found useful (successes)?
2. What hasn't been picked up, what wasn't useful? And what stopped its adoption?
3. Are there lessons we can learn from this? And where should we go over the next 10 years?

What ensued was a lively and interactive session, which informal feedback suggests was both one of the most enjoyable aspects of the workshop and an excellent ice breaker. The sessions certainly enabled all to participate, from first-time WMCSA attendees to battle-hardened veterans of WMCSA. The feedback from the groups was extremely varied, encompassing the full spectrum from (subject to artistic license) "It's all been an enormous success," to the somewhat dour "Our ideas have been almost completely ignored." We try to capture some of the key insights from their reportage below.

SUCCESSES

- Mobile computing is commonplace. We have mobile phones, laptops, instant messaging, and an extensive array of communications options, including wide-area data connectivity and the now ubiquitous 802.11 hotspot (others wryly opined that this was driven by market forces, not the research community!) Despite this, wide-area bandwidth is still insufficient or prohibitively expensive for many applications.
- The good news is that some of the ideas made it! Concepts from "mobile" file systems such as CODA (e.g., offline files, synchronization, and caching) made it into mainstream operating systems, and image distillation for mobile Web browsing is now in commercial products and standards.

FAILURES

- Location-based services and on-demand audio-video services still haven't happened. Arguably, we're still waiting for the killer application, but many believe that the reason we haven't found it is because it may not exist!
- Mobile applications are still fiendishly difficult to write, due to a lack of common programming abstractions and the many difficulties introduced by different standards and models of mobile device (even among families of the same device). Some also argued that research should only build enough to show the value of an idea, not build the whole system (this clearly depends on the nature of one's research).
- Mobile-user interfaces still don't adapt in response to underlying changes and still don't provide feedback to users (the "bars of connectivity" metaphor on mobile phones is possibly the only exception!).

LESSONS FOR THE FUTURE

- The surprising realization (from several groups), perhaps due to multidisciplinary

influences of ubiquitous computing or commercial drivers, was that we are still too technology-driven and need to remember the user. Applications should drive the technology, not vice versa. A top-down approach and user-centered design may lead to better-motivated and more acceptable systems research.

- There are no common platforms, and we are not quick to build on each other's work. There is little standardization and few accepted solutions. How, then, do we distill what we learn to make it more easily accessible to others? The platform researchers among us would doubtless like to see an open architecture, as we lack integration. However, many recognized the need for more collaboration, even within a single academic department.
- It is still very complex to build a complete mobile system, as it requires lots of design decisions, and it is even harder to deploy one! Perhaps it is not surprising, then, that it is hard to measure or quantify which approach is best, as there are few commonalities between any two mobile systems.
- We may lack criteria for effectively evaluating our work. If we in this community can't say what technologies to choose for particular aspects of mobile computing systems, then it is unreasonable to expect others from outside our community to take the lessons away. Others agreed that the distillation of the last 10 years' work has yet to take place. Perhaps we need multiple studies on the same problem to access its true value and publish the "best practice" of the community.
- We also said that often there were many approaches in our work that weren't successful, and that this was valuable knowledge that needed to be captured. Some claimed the need to publish "what we failed to do," although there would be many cultural barriers to overcome in getting researchers to talk about their failures. A "failure track" at a conference was suggested.
- A final suggestion was that there was much potential in mobile gaming, though computer science researchers tend not to have much impact in this domain. This may be due to the difficulty in obtaining funding for work in this area.

Paper Track

The paper track was arranged into sessions of three papers each, every paper followed by brief clarification questions. The session would conclude with a final panel composed of all three speakers in which questions could be addressed, either tackling cross-cutting issues or involving in-depth questioning about a particular issue. Details of the sessions and copies of the slides from each presentation can be found online at http://wmcsa2004.lancs.ac.uk/programme.shtml.

SESSION I: APPLICATIONS

The first of these sessions focused on complete systems and applications: a system for relaying the biometric performance of athletes to a Web portal (in this case, cross-country skiing); a system for providing context-aware information to mobile devices relevant to nearby posters; and a wearable system for augmented reality search and collection tasks.

One interesting aspect of the biometric system was that the authors overcame weak connectivity to the participants through a motion prediction algorithm. When questioned about which aspects of the system users were happy with, they claimed that 80% were happy with the updates, yet many found the estimated value irritating—even though they could see on the television coverage that a particular skier had stopped, the estimation system was still updating his location! Curiously, the researchers had not thought to reflect to users when estimated data was being used.

The authors were collectively asked what were the most important lessons to take from designing systems for real people. They all agreed that ease of use was key. Perhaps following from the opening discussion, user studies were also seen as an important aspect. Questionnaires could be an effective way to reach many people but were potentially an imprecise instrument; focus groups were recommended as preferable for engaging users. It was remarked from the floor that it's important to build confidence before spending too much energy developing systems and that questionnaires might provide that important first feedback. In the past decade there has definitely been a cultural shift toward including the user.

SESSION II: LOCATION TRACES

The second session was about studying, tracking, and exploiting user mobility. The first paper presented a study of using access logs of email use while on the move to develop traces of user mobility patterns, the second concerned the potential for exploiting such mobility to offer opportunistic networking, and the last was on the usefulness of Cell-ID location systems in offering location-based services to mobile users (based on some deployment experience).

The first speaker was asked about whether those studied had privacy concerns and whether these would

be exacerbated in a longer-term study. He confirmed that they did, even for the month-long study conducted. Concerns were also voiced as to whether IP addresses taken from email access logs corresponded to locality since the advent of VPNs/mobile IP and to temporality, as people are inclined to leave their email client open even when they are not physically present (one assumes that this must depend very much on the individual).

The second speaker drew the unexpected conclusion that it was sometimes faster to use person-to-person delivery to reach the infrastructure than to communicate with it directly. One supposes this assumes that the node stays out of reach of the infrastructure directly.

One of the questioners mentioned that one of the hardest things to do is to find experimental data to work with. The speakers were collectively asked whether they'd be prepared to put their traces online. Mary Baker, the program chair, mentioned the "http://cmc.cs.dartmouth.edu/data/Dartmouth archive of wireless-network trace data" that aims to promote the collection and sharing of mobility trace data for the community.

When asked, "What can we learn from your traces to build better applications?" the first speaker remarked somewhat wryly, "Users are not as mobile as we thought."

SESSION III: CONTEXT AWARENESS

The third session was on the ever popular topic of "context awareness." The first paper was about coordinating contextually driven components through "the environment," inspired by the observations made by the French biologist Grassé on how social insects coordinate their actions using indirect communication through pheromones (stigmergy). The second paper focused on the deployment experience of three context-aware applications based on a local broadcast (beaconing) approach. The session closed with an interesting insight into the artifact-centric design method advocated by the paper's author, in which context determination is done using low-level embedded sensing.

One important remark from the floor about the temporal nature of context, and activity tracking in general, was that "before-ness" is not exact: one rarely switches between activities cleanly; activities overlap, may be concurrent, are often ill- or subjectively defined, and are hard to disambiguate from one another.

Questioners also pointed out possible flaws in the artifact-centric and stigmergic approaches, in that it is potentially difficult, or at least as yet unproven, how one can build more complex applications. Optimizing the total amount of bandwidth consumed and yet not losing important data (e.g., through low-level filtering) is also an important issue for the authors of all three papers to consider. The second author did point out that reference locality (if users request the same information in the same location) can be of benefit in helping to address this issue.

SESSION IV: AD HOC NETWORKS

Routing in ad hoc networks was the broad theme for the fourth session. The first paper focused on providing effective routing support for highly mobile nodes by predicting trajectories; the second was on detecting misbehaving nodes in DSR using WatchDog mechanisms; and the last, and somewhat topic outlier in this company, was on utilizing "Plan 9" and file abstractions for programming mobility-resilient pervasive applications.

One questioner pointed out astutely that the first authors' mobility simulations, although initially impressive, were assuming an epoch of 100 seconds, leading to a consistent and predictable direction for a node even after a mobile cell had been crossed, simply by virtue of the simulation values that were picked.

The second author was asked whether one could detect against coordinated attacks by groups of nodes. The speaker responded that there was little to do if you are surrounded by "all evil" nodes. The technique is robust, needing only one minute to identify nodes that are rapidly changing their identity to prevent detection.

The last authors' work received some criticism for not paying due deference to early work in UNIX (and on the streams abstraction in particular).

SESSION V: PERVASIVE TECHNOLOGIES

Fifth on the program was a session discussing mobile systems in a pervasive-computing context. The first author presented a classic Web proxy for mobile Web browsing, their thesis being that image fidelity was best adapted to user interaction shared across a community of users. The second paper was on how one might instantly personalize devices to achieve a consistent user experience when ownership is transitory. The final paper of the session was on using physical "toss & swing" movements of the mobile device to trigger information transfer between users.

Some concern was expressed to the first two authors as to the generality of their classification schemes, partitioning users into groups in the first case, and moving to more applications in the latter. It was opined that activity might form a better metric for clustering than the user. Philosophically, one questioner expressed doubt that devices would be shared between users in the future, since making them sufficiently tamper-resistant

would require costly additional trusted hardware. (Ed.: One should note, however, the integration of fingerprint readers and boot-time verifiers in some vendors' commodity mobile products.)

SESSION VI: MOBILE P2P AND SENSOR NETWORKS

The penultimate session concerned peer-to-peer and sensor networks. The first paper discussed adapting Gnutella's protocol to allow peer-to-peer information sharing in ad hoc networks (using Bluetooth). The second was on integrating the Pastry Distributed Hash Table (DHT) algorithm with Dynamic Source Routing (DSR) in ad hoc networks. The last paper presented an approach for the physical placement and location estimation in sensor networks based on "deployment order" and the identification of "landmarks" at known physical locations.

In the open forum, there was considerable focus on the overhead of using DHTs in ad hoc networks (specifically, as nodes join and leave and the impact of mobility-induced failures). The speaker admitted that there was more work to be done to evaluate this, but that mobility was largely handled by the standard DHT routing failure mechanisms (unless nodes joined and left too frequently, in which case data could be lost and there would be overhead in maintaining consistency).

SESSION VII: MIDDLEWARE MODELS

The final and closing session of the workshop focused on middleware architectures, addressing refreshingly familiar problems in mobile systems. The first paper described a cyber-foraging architecture whereby expensive computations are offloaded to more powerful compute servers to increase performance and reduce drain on battery life. The second concerned a "Service Data Object"–based replication platform for allowing mobile access to data-driven Web applications. The final paper of the workshop described a power-aware Web proxy for use in wireless LANs which schedules the delivery of Web content to maximize the time the wireless receivers can spend saving power during activity.

While the author admitted that adding such proxies increased both complexity and delay, in tests the system was able to save up to 50% power while browsing popular test Web pages.

Demonstration Session

The demonstration session was accompanied by a buffet dinner. Although it is impractical to summarize the many eloquent conversations we overheard, we summarize the demos:

- **Seamful game**, demonstrated by Marek Bell and Paul Tennent, University of Glasgow, is a game based on ad hoc networking of mobile devices, where users go out to acquire virtual coins located in physical space and must exploit wireless networking hotspots to "cash in their bounty"—making the seams of the otherwise invisible network a key to winning the game. Although there is a working prototype, we were not able to play it, as it was a cold, dark, and potentially perilous winter's evening outside.
- **SoulPad: Personalized Computing with Minimal Infrastructure**, demonstrated by Ramón Cáceres, IBM Research, is a system based on virtual machines, the ever ubiquitous and increasingly useful flash disk, and a self-configuring version of Linux. The system allows a user to checkpoint and restore their entire laptop configuration using a bootable flash disk. The trick of course is in their efficient integration.
- **Information Dissemination in Spontaneous Networks**, demonstrated by Andreas Heinemann, showed a peer-to-peer network designed to disseminate information (e.g., ads) from a shopping mall. Users get bonus points for passing ads to a person who eventually makes a purchase. At the workshop they showed an mp3 dissemination application where users can specify their music interests/choices and receive music files from users in their immediate environment.
- **Fuego Core**, was shown by Sasu Tarkoma, Helsinki Institute for Information Technology. Sasu's demo showed a GUI simulation of a smart environment, where events are delivered based on current location/context. This can be used to investigate users' requirements for such information.
- **\net**, a laptop-based demonstration of the system presented in the paper program by Gorka Guardiola, Universidad Rey Juan Carlos, showed how interfaces could be flexibly created and migrated across devices using file-system primitives in their modified Plan 9 OS.
- **Unscripted interlude:** There was an unscheduled demonstration of the latest (very small!) "particle" Smart-ITs from Tec-O and from LMU a Smart-IT interfaced to a number of small "phone size" LCD displays for embedded use in pervasive artifacts.

LOOKING FORWARD

As we witnessed at this year's workshop, after 10 years there is still very much a community doing mobile systems work. We look forward to hearing about this as it evolves over the next decade of WMCSA's future. The 2005 workshop will take place somewhere in the U.S., with Maria Ebling (IBM T.J. Watson Research Center) as the program chair and Anthony Joseph (University of California, Berkeley) as general chair.

We would like to take this opportunity to thank everyone who helped put WMCSA 2004 togethers. Thanks also to Fahd Al Bin Ali and Sachin Goyal (our student scribes) for volunteering to take the notes that made these reflections possi ble.