

Workshop on Mobile Computing Systems and Applications December 1994

Digest of Proceedings

M. Satyanarayanan
Program Chairman

1. Introduction

The goal of this two-day meeting was to foster interaction between active workers in mobile computing, with a view toward cross-fertilization of ideas. Given the youth of the field, such interactions could have substantial impact on its future direction. In keeping with this goal, the conference organizers chose to have a small, informal workshop rather than a larger and more formal conference. The workshop was sponsored by the IEEE Computer Society Technical Committee on Operating Systems, in cooperation with ACM SIGOPS and USENIX.

The workshop was held on Thursday and Friday December 8-9 1994 at the Dream Inn in Santa Cruz. The weather was beautiful and the oceanside locale spectacular --- alas, it is not clear whether these helped or hindered the workshop, since many longing looks were evident on the faces of participants as they gazed out of the windows! The General Chair, Darrell Long, had done an excellent job of selecting the workshop site and setting the stage for the workshop. He was assisted in local arrangements by two student volunteers, Jim Cummiskey and Chane Fullmer.

What follows is a summary of the discussions that took place during the workshop. It is based on notes taken by four student volunteers (Peter Grillo, C.K. Toh, Adrian Friday, and N. Asokan). They did an excellent job of taking detailed and complete notes. Any errors or omissions in this document are certainly my responsibility, not theirs.

This digest is intended to be a supplement to the papers in the proceedings, not a substitute. Rather than producing a verbatim transcript, I've tried to focus on those interactions that seemed most insightful, controversial or evoked most response from the audience. Such a report must, by its very nature, be subjective. I've tried to be as objective as possible, but I'm sure there are places where my personal biases show through. My apologies in advance if you attended the workshop and your favorite comment, question or discussion isn't mentioned here.

2. Thursday, December 8

2.1. Models and Methodology

The theme of the opening session, chaired by Randy Katz, was the identification of novel ways of thinking about mobile computing and using these viewpoints to derive system structures. Doug Terry of Xerox PARC presented the first paper, on the architecture of the Bayou system. The problem addressed by this work is the maintenance of consistency in shared, replicated data repositories updated by mobile hosts. Bayou's model of consistency is reminiscent of that of Grapevine, built nearly fifteen years ago. Updates by a mobile host to a particular repository site are tentative, until those updates are received by the primary site responsible for the data in question. Updates are propagated to all data sites in an epidemic, or rumor-mongering fashion and may become visible to other mobile hosts even before final validation by the primary site. Since secondary data sites may not receive all updates in the same order that the primary site finally chooses to order them, the state of the data on secondary sites may differ and tentative updates that have already been applied may have to be rolled back and reapplied after other incoming updates. The system is careful, however, to always make it clear to users which data is derived from tentative updates and which from permanent.

The questions after the presentation addressed two areas: clarification of the consistency guarantees, and the rate of convergence. Doug indicated that a tentative update may be rolled back and reapplied at a given secondary site at any time until that site has heard what the final "commit" ordering is that the primary site has chosen for the update. He also indicated that the anti-entropy mechanism responsible for update propagation may be executed many times for each update (if there are many servers). But the update procedure converges as long as no host remains partitioned forever.

The second talk, by Arup Mukherjee, made the case that existing work on mobility focused on computation and communication, to the exclusion of control. His thesis was that a rich taxonomy of applications emerges when control is given due prominence, and that the taxonomy offers valuable insights into structuring applications to function effectively under the constraints of mobility. In particular, Class 7

applications (in his 7-element taxonomy) were currently under-represented but offered many advantages in a mobile environment. The questions after the talk addressed two issues. One question was whether real applications could be mapped as cleanly into the taxonomy as the speaker claimed. Arup replied that complex applications are often composed of subsystems in classes distinct from that of the parent. He added that it is the task of the system builder to examine an application at a level of granularity relevant to the issues being considered. The other question was really an observation that Class 7 applications demand mobile hosts to have substantial computing resources; something like an Infopad will not suffice.

Perhaps the most controversial item in the workshop was the talk entitled "Are Disks in the Air Just Pie in the Sky?", given by Mike Franklin. The approach is to use a network as a rotating information medium by periodically retransmitting the entire contents of databases. The central idea behind this work is to superimpose multiple disks spinning at different speeds on the broadcast medium in order to support non-uniform data access. Rather than fetching data on demand, clients continuously listen to the transmissions and cache information of interest to them. This approach is especially valuable when the network has asymmetric bandwidth. The flurry of questions at the end of the talk covered many aspects of the work. Satya pointed out that networks in mobile environments tend to be unreliable: how can you depend on broadcast data when mobile? Mike agreed that this was a problem, but that it could be addressed by prefetching critical data. Mary Baker observed that broadcast may not be supported by some mobile networks. Karin Petersen warned that receiving data costs energy; it is therefore an illusion to believe that the broadcast approach comes for free. Mike stated that for some important applications, such as advanced traffic information systems, battery power is not a concern, and that given sufficient demand, there is no reason why lower-power mechanisms for monitoring the broadcast could not be developed.

2.2. File Systems

Lily Mummert presented the first talk in this session, chaired by Peter Honeyman. Her talk focused on techniques to cope with the performance and reliability of mobile networks. The techniques spanned three areas: deferring update propagation during periods of low bandwidth, opportunistically using high bandwidth when available, and the use of an abstraction called "dynamic sets" to reduce network latency during search. In the question period, Peter Honeyman asked how log replay is actually performed during trickle discharge. Lily answered that the replay occurs as a set of iterations on small parts of the log. Terri Watson pointed out that applications had to be changed in order to use dynamic sets, and that they have to be able to tolerate the reordering of requests implicit in the use of dynamic sets.

The second talk, on shrinking a replay log using peephole optimization in a postprocessing step, was presented by Larry Huston of the Little Work project. This approach is in contrast to that of Coda, which applies optimizations incrementally. The primary advantage of the Little Work approach is that the optimization code is a separable component; hence it is easy to apply to multiple file systems. Peter Honeyman asked what timestamps files received; Larry replied that they received the replay time, rather than the true modification time, because this allows programs like "make" to work correctly. Jay Kistler asked what the asymptotic performance complexity of this optimization technique was. Larry answered that it was $O(N^2)$ worst case, but that the running time in practice was quite acceptable. In response to a question from Terri Watson, Larry said that operation reordering was essential for up to 60% of the optimizations they were able to achieve.

Predictive caching was the topic of the third talk in this session, presented by Geoff Kuenning of the Ficus project at UCLA. The goal of this work is to reduce the burden on users of specifying files to be hoarded in anticipation of disconnection. The system uses a list of observed file references and a set of clustering algorithms to construct a plausible mapping of those references into distinct tasks. Hoarding is then performed on tasks rather than individual files. In the question period, Lily Mummert pointed out that multiprogramming would complicate the clustering analysis, since the observed stream of file references would be the union of two or more distinct tasks. Geoff agreed that this was the case, but said that clustering analysis could be refined to distinguish between one primary task and a number of secondary ones, a common scenario in single-user multiprogramming environments. Jay Kistler asked how much simulation of the proposed scheme had been performed. Geoff replied that he preferred results from real use to simulation results.

2.3. Wiring the Campus

In this first panel of the workshop, moderator Rich Wolff began by observing that the title of the panel was only a loose characterization of the work represented in it. Each of the participants then gave a brief summary of their work.

Abhaya Asthana described the design of a shopping environment with wireless connectivity for each shopping cart. Vince Russo gave an overview of the deployment of a wireless network at Purdue University, using an ATM backbone switch to cope with en masse movement of many users, such as will occur between classes. Mary Baker reported on a new project, called MosquitoNet, to increase connectivity when switching a host between wired and wireless communication on and around the Stanford campus. Since all three projects are at a very early stage, there were no war stories to report. The ensuing discussion focused on two major issues, both relating to the campus wireless projects.

The first issue was whether truly "mobile" computing, in the sense of people computing while walking across campus, was either likely or desirable. Many members of the audience felt that a more likely scenario involved students using their portable computers in each classroom, library, etc. but not while they were walking. For this scenario, all one needs are network outlets at each desk in a classroom; wireless coverage is not necessary. Satya pointed out, however, that truly mobile applications do exist. For example, experiments are in progress at UC Santa Barbara to allow visually handicapped people to navigate on campus using portable computers to sense current location and to give directions with voice synthesis.

The second issue was the impact of campus mobile computing on social mores and etiquette. For example, how does one prevent electronic cheating such as by students passing zephyr messages to each other during an exam? Even with a perfectly honest population, there are issues such as whether it is acceptable for a person with noisy keyboard to disrupt a lecture, or to intrude upon a discussion.

2.4. Application Frameworks

In the session after lunch, chaired by Dan Duchamp, four papers were presented. Each of these papers focused on a broad class of mobile applications, and described a paradigm or set of techniques applicable to that class.

The paper on teleporting, presented by Frazer Bennett, reported on experience with using a system that allows the display of an application to follow a user around as he moves, leaving program execution at the original site. This ability is especially convenient when combined with an active badge system that tracks user location. Questions from Peter Honeyman and James Kempf probed the limitations of this approach. In particular, they were concerned that hiding display changes from applications would render them unable to adapt correctly to changes in display size or color characteristics. Frazer agreed that this approach would indeed be inadvisable for applications that were tightly coupled to specific display characteristics. Dan Duchamp asked how ambiguities, such as the presence of two displays in the same room, were resolved. Frazer replied the user is iterated through the choices of display and can pick one. In response to a question from Karin Petersen, Frazer said that it was not possible at present to allow selective movement of windows. There was also a flurry of questions and heated discussion on issues of privacy and security.

The next talk, by Roy Want, described work at Xerox PARC on making ParcTab applications sensitive to the current physical location of the user. David Steere asked what the security consequences of losing a Tab were. Roy and Karin Petersen explained that each ParcTab was associated with a user, and that loss of a Tab was as serious as losing a key, though some additional security could be provided via a PIN code. Doug Terry added that the privileges of a Tab could be

easily revoked by killing the proxy server associated with it. Randy Katz requested details of the infra-red communication mechanism used by the Tabs. Roy said that the typical bandwidth was 19.2Kbps, although bandwidths up to 1Mb/s were possible.

Terri Watson then described her experience with designing applications for wireless computing. The theme of her talk was that developers should exploit application-specific knowledge to address mobile resource constraints. In certain cases, it is desirable to offer alternative actions to the user, allowing them to make performance versus cost decisions. Geoff Kuenning asked whether it is realistic to expect all existing applications to be rewritten according to this philosophy. Terri replied that the highest payoff applications would be rewritten regardless of effort involved, and that the total number of viable applications in a mobile environment were limited.

The final talk in the session was by Kenjiro Cho, reporting on the use of group communication primitives for mobile computing. The talk closely followed the paper, with emphasis on establishing that the performance overhead of this approach was indeed acceptable. In response to David Steere's question about behavior during network partitions, Kenjiro explained that ISIS only supports group communication in the majority partition. C.Toh asked whether clients needed to explicitly select a new primary server during partitions; Kenjiro replied that this selection was subsumed by ISIS.

2.5. Exploiting Mobility Commercially

Many participants have told me that this panel, representing industry's perspective on mobile computing, was the most exciting part of the workshop. Amal Shaheen of IBM Austin, the moderator of the panel, posed four questions for the panelists:

1. Is there money to be made in mobile computing?
2. What are the characteristics of successful mobile applications?
3. What is the impact of mobility?
4. What are the merits of a client-only approach versus one that requires modifications to both clients and servers?

After posing the questions, Amal gave her answers to them. She was confident that there is a lot of money to be made in mobile hardware, but felt that there is no data to decide whether the same was true of software. The trick will be to find out what the users expect and deliver something more than that expectation. She felt that packaging and ease of use were important characteristics of a successful application, Lotus Notes being a good example. Transparency can only go so far: things like conflicts and cache misses during disconnections are impossible to hide. Finally, Amal observed that it is logistically much simpler to provide support entirely

at the client end. Server-end changes render existing servers incompatible, and are thus much less attractive. This remains true even when server changes offer substantial functionality or performance gains.

Murray Mazer (now at the OSF Research Institute) spoke next, and reported on his experience with mobile computing at Digital Equipment. He observed that a broad range of people in the computer industry (ranging from Bill Gates and market analysts to real users) believe that there is a market for mobile computing. He therefore believes that there is definitely money to be made in it. He then pointed out that mobility will not be the differentiating factor in the future; rather, it will be the norm. Exactly when this will happen depends on when the infrastructure for mobility becomes widespread. Regarding applications, Murray observed that users are intolerant of bad interfaces. They will not go through poor interfaces to get to the cute functionality as we implementors might. They hate poor performance and unannounced missing functionality. Hence we should strive to make the user-visible components easy to use; this, in turn, requires us to manage complexity in applications and services. He expressed the belief that people will pay for valuable functionality; for example, cellular phones are popular even though their use is expensive. Rather than focusing on vertical applications, which is today's market, he suggested that remote information access was going to be the fastest growing and key class of applications. Finally, Murray argued for making quality of service more explicit in applications: be more careful in setting user expectations, and allow users to make explicit tradeoffs of cost and performance as far as possible.

The third panelist, Bill Fitler of Lotus, reported on his experience with the CcMail and Notes products. He first pointed out that there was definitely money to be made in mobile computing, and that the popularity of these two products was proof. He emphasized that total transparency was never going to be possible, and that users were not expecting it anyway. Mobility results in a very harsh environment for applications, and they often fail in serious ways under these circumstances. Bill also noted that support for mobility is much like support for fault tolerance: it has to be built-in and cannot be added on later.

Dorota Huizinga was the next panelist, speaking on behalf of herself and her collaborator Ken Heflinger of AST Research. She began by noting that their work had been inspired by Coda, and that they had persisted in their efforts to implement disconnected operation in DOS in spite of the fact that their measurements of write-sharing in the AST environment were significantly higher than those reported for Coda. For the same reasons that the previous panelists had cited, their work was an entirely client-side implementation with no server changes. Dorota noted that many of the implementation challenges they faced had nothing to do with mobility; rather, they were caused by the memory addressing limitations of DOS. Finally, she expressed the belief that

there was money to be made in mobile computing, but admitted that she was unable to substantiate this belief with specific data.

The next panelist, James Kempf of Sun Microsystems, was very brief. His primary message was that mobile computing applications would benefit greatly from widespread support for a special language that would allow applications to download code easily. In response to Peter Honeyman's prompt about Telescript, he agreed that the language should not be proprietary.

The last panelist, Bob O'Hara from Microsoft, was confident that there was money to be made in mobile computing. He observed that there were three portable computers in his presence right there at the workshop: a laptop, a pager, and a watch which was a joint product of Microsoft and Timex that could download his schedule from software running on a PC. Peter Honeyman asked whether we were likely to see body implants, to which Bob replied that it didn't matter whether the hardware was worn on the outside or the inside. Regarding transparency, Bob was of the opinion that it was important because it was the key to allowing third party software developers to write applications easily. Barry Leiner asked how he hoped to hide limitations of the network for applications like video, to which Bob replied that he had not given this class of applications serious thought. On the matter of mobile applications, Bob observed that vertically integrated applications like appointment books tended to be the most successful.

The rest of the panel session consisted of a number of discussions spanning the range of topics touched upon by the panelists. Amal, Bill Fitler, and Satya engaged in a heated discussion about the level of abstraction at which support for mobility should be provided. Amal argued that the support should be at the file system level, because all applications could benefit from it. Bill countered that providing the support at a higher level (such as the Lotus Notes application) allowed more information to be used for conflict resolution. Satya pointed out that this need not be an "either/or" situation: Coda provides support at the file system level, but allows application-specific resolvers to be transparently invoked upon detection of a conflict.

A second topic of discussion was on the issue of usability. Satya observed that the harder one worked to mask the ugly characteristics of a mobile environment, the more difficult it was to explain to naive users what had gone wrong when the masking was no longer feasible. The panelists agreed that this was indeed a difficult problem. Murray Mazer and Bill Fitler gave simple examples of how errors could be presented to users in meaningful and easily-understood ways, but everyone agreed that these merely scratched the surface of a difficult problem.

Marvin Theimer warned panelists not to place so much trust in marketing surveys. After all, pen-based computing had

been predicted to be a major market but there are no signs of it taking off yet. He then offered the opinion that entertainment (including games such as multi-user Doom!) would be the driving force of mobile computing. If this turns out to be true, he observed, the entertainment industry might pay for the cost of the mobile infrastructure. There was substantial disagreement on this conjecture. Many among the audience and panelists considered it unlikely that entertainment would pave the way for other mobile computing applications.

Dan Duchamp directed the panelists' attention to a different topic: academic research on mobile computing has focused on Unix, while industry is almost exclusively focused on Windows/DOS. Dan asked whether this was a healthy dichotomy, and whether academic research should switch to Windows/DOS. Bob O'Hara replied that the industry approach could be characterized as "Small steps for small minds." With the passage of time, the Windows family is getting to be more like Unix. Further, visitors from universities do contribute their Unix biases to industry. Hence Bob advised academia against giving up on Unix, but not to forget about desktop systems such as Windows.

The last few minutes of the panel session were spent on a potpourri of topics ranging from cellular telephones to a revisit of the importance of entertainment. But the long day and the aroma of hors-d'oeuvres from the next room sapped the vigor of the discussions. The panel and the day ended on a quiet note.

3. Thursday Evening: Exhibits

A set of exhibits from industry and universities, organized by Peter Honeyman, was displayed concurrently with the reception at the end of the first day. Peter had done an admirable job of ensuring that the exhibits were not mere marketing glitter but had something insightful to offer to the participants of the workshop. There were six exhibits, of which two were commercial products and four were research prototypes. Each is briefly described below.

3.1. IBM Mobile FileSync

Amal Shaheen and Tom Porcaro of IBM Austin demonstrated a new IBM product, Mobile FileSync, that has been bundled with Lanserver 4.0 for OS/2. Inspired by Coda, but differing considerably in its detailed design, this product supports disconnected file access in OS/2. The support is entirely at the client end, with no changes required to existing servers. The current version of Mobile FileSync provides support for hoarding, as well as for step-by-step reintegration via an interactive process. An important aspect of the implementation is that it is layered entirely above the file system switch. As a result, the support for disconnected operation works with any file system below the switch. The exhibit involved two IBM ThinkPad laptops on an infrared wireless LAN.

3.2. Lotus Notes and CCMail Mobile

These two popular products from Lotus were demonstrated by Bill Fidler. They are both examples of vertically integrated applications that originated in LAN networks but have been extended to mobile environments.

Notes is relevant to mobile computing because of its replication model. A client can connect to the network and obtain a replica from a server. Once a replica is downloaded, it can be used "off-line" (i.e., disconnected from its server). Considerable effort is made to hide whether you are on-line or off-line, but user control is possible via a sequence of menus. There is a full scripting language for creating filters so that only desired information is collected from the server in any given connection.

CCMail Mobile looks identical to the LAN version, with the addition of one new menu which deals with all the mobile aspects. This allows a user to send, receive and move messages back and forth between a mobile client and a server. The system allows you to set up default usage locations, and to associate those locations with attributes such as modem type and dial prefixes. A variety of communication mechanisms, including over 150 modem types, are supported. These can be tried in order rolling over from one to the next to discover a communication mechanism that works at the current location. Scheduling functions exist to allow the user to contact the server at startup, closedown or user specified intervals. Filters can be constructed to select messages based on criteria such as size and priority.

3.3. PARC Tab

Norman Adams from Xerox PARC demonstrated the PARC Tab hardware that has been used in a variety of experimental projects. The Tab has a small, graphics-capable screen, 128KB of memory, and an infrared transeiver. The infrastructure at PARC consists of room-sized cells equipped with infrared transeivers. Each Tab has a server process running on its behalf on a workstation on the wired network. Applications on a Tab can be implemented as Tcl scripts that are executed on the server, or as standalone programs with surrogate processes on the server.

The demo consisted of two cells and a Sparcstation functioning as server. The concept illustrated by the demo was that of "proximate selection". One example consists of a user walking into a cell, and selecting "forward call" on his Tab: his phone calls are automatically forwarded to the room he is in. Another example consists of an application to list available printers, with nearest first: when the user walks to a different room, the display automatically changes.

3.4. Wit

Wit is a research prototype built by Terri Watson of the University of Washington. The client hardware consists of infrared transceivers developed for the Xerox PARC Tab project and stock HP 1000LX palmtops. Largely unmodified

PARC Tab code implements low-level transport.

The software system consists of two components: a network-side proxy and a palmtop system that extends the DOS environment to support multiple active applications through windowing, user threads, and network connections. Tcl interpreters in both components serve as the primary application programming interface. Application functionality is partitioned between the proxy and palmtop by dynamically defining and executing new Tcl functions on the remote side, with the goal of reducing both bandwidth consumption and user-perceived latency. Terri observed that Tcl treats all data as strings, which can complicate applications' use of non-ASCII data.

3.5. Marine Maintenance Assistant

Arup Mukherjee, of the VuMan project at Carnegie Mellon University, demonstrated a wearable computer. It consisted of a small computer with a Private Eye display and a hand held controller with three buttons. The software on the machine was customized for a specific application, that of access to documentation for maintenance tasks. The demonstrated version of the system was 80C186-based, but later versions of the system will be 386-based.

3.6. Teleporting

Frazer Bennett showed a brief video to illustrate his earlier talk on Teleporting. The video showed people wandering around, pressing their active badge buttons and having X displays migrate to their current location.

4. Friday, December 9

4.1. Networks & Protocols

The first session of the second day was devoted to the topic of networking and protocol issues in mobile computing. Ramon Caceres chaired the session, in lieu of Krishan Sabnani who was unable to attend due to a personal emergency.

In the first paper, Raj Yavatkar examined the problem of end-to-end TCP adaptation in mobile environments. He observed that such communication often involves a short wireless segment and a much longer LAN or WAN segment. Standard TCP code fails to recognize the very different reliability characteristics of these two segments, resulting in unsatisfactory performance. Raj described a solution in which intermediary code allows TCP to independently adapt to the characteristics of the two segments. His solution provides substantially better performance, while preserving complete upward compatibility with existing clients and servers. Barry Leiner asked whether the goal of not changing TCP was a valid one. TCP was designed with certain link-level characteristics in mind, and if those cannot be met, it is better to redesign TCP. Raj disagreed, saying that preserving TCP unchanged as far as possible has enormous practical value. Further, the 10-12% packet loss typical of wireless segments is far too high; improving link-level reliability is essential.

Later, Barry asked whether it was appropriate to consider the proposed scheme end-to-end, because the intermediate code violates the end-to-end reliability semantics of TCP. The session chair shared the same concern and seconded Barry's comment. Raj replied that the situation was no different from that of a gateway.

Nigel Davies presented the second paper, describing experience with a mobile application for the electric utilities in the UK. The goal is to help linemen collaborate effectively with each other and with control rooms. The system developed for this requires each lineman to have a laptop with support for wide-area wireless communication. The client software includes collaborative tools for displaying and editing maps and provides users with feedback on the quality of the underlying communications network. Initial trials with the system have been conducted, and wider deployment is expected. In response to a question from Mary Baker, Nigel said that the main feedback from users was that they wanted the client software to look and feel more like Windows. Users also wanted the collaboration software to better distinguish input from different users. Barry Leiner asked where information about network quality was obtained, and whether the TCP stack was bypassed in doing so. Nigel replied that TCP was not used, and that the custom-built RPC layer provided an interface for applications to obtain information about network quality.

The third paper in this session addressed the problem of wireless communication between mobile hosts in locations where there are no base stations or other mobile infrastructure. Dave Johnson described a protocol in which the hosts themselves serve as forwarding agents and thus constitute an impromptu mobile infrastructure. There was heated discussion over whether a user would like his machine's cycles to be used for routing someone else's packets. Dave observed that this was the price of membership in an ad hoc mobile network. Terri Watson asked if signal strength could be modified under program control; Dave replied that current wireless hardware does not permit this.

Finally, Allen Lao of UC Berkeley presented a paper on a video transport protocol for wireless networks. The novel feature of this protocol is its ability to dynamically adapt the bandwidth required to the current content of the video. Specifically, video segments with a large amount of motion can be rendered in a lossy manner without noticeable degradation of picture quality. Allen noted that this is the opposite of MPEG, where segments with lots of motion tend to result in higher bandwidth requirements. Jim Kempf asked how this worked with video conferencing, where lip sync is important. Allen replied that this could be handled by using low resolution for the mouth area, but ensuring that it was sampled frequently. The talk ended with a brief video demonstrating the concepts.

4.2. Accessing the World-Wide Web

Over a short period of time, the World-Wide Web has acquired star status as an information repository. This session, chaired by Jay Kistler, focused on the topic of accessing the Web from mobile clients.

The session began with Joel Bartlett describing his experience with implementing a Web browser on an Apple Newton, communicating via a low-bandwidth wireless link. The talk was accompanied by a video demonstration. Joel observed that his strategy of partitioning applications so that the CPU-intensive processing occurred on powerful servers was critical to good performance; a German team that had also done a PDA implementation of Mosaic had gotten only 10% of Joel's performance. Terri Watson pointed out that Joel's strategy was consistent with the design philosophy she had espoused earlier in the workshop. Frans Kaashoek inquired about prefetching, and Joel replied that the next PDA screen was prefetched. Jim Kempf and Bob O'Hara asked about the client-server protocol and the server hardware. Joel replied that the protocol was custom-designed, and that the server was a DEC 5000 with a MIPS R3000 processor.

The second talk, by Josh Tauber, described a very different approach to mobile access of the Web. Web documents are now programs in Tcl/Tk that are executed at the client by an interpreter that enforces safety. At present the system works on IBM ThinkPad clients and Sparcstation servers over a 2 Mb/s WaveLAN wireless link. Terri Watson questioned the basic assumption of the approach: namely, that authors would be willing to write programs rather than documents. She observed that a major reason for the success of the Web was the simplicity of the HTML format. Josh replied that authoring tools would help in this process, and that development of such tools was essential for the success of this approach. Terri then expressed skepticism about the portability of the approach: different versions of each document would be necessary, to allow different types of clients and interpreters. Karin Petersen agreed, saying that translation between HTML and the client side filter was necessary; this would obviate the need for authors to foresee all possible client configurations. Josh observed that interface discovery techniques could be used to help. Murray Mazer suggested that for every type of mobile entity, there be a server agent that could perform appropriate translation. Finally, Marvin Theimer proposed that attention be focused on defining a standard PDA interface, rather than supporting heterogeneity.

In the final paper of this session, Geoff Voelker described a publish/subscribe approach to contextual behavior in Web documents. In this approach, active documents subscribe to some variables; these variables are periodically updated by agents. A change in a subscribed variable causes a document to be reloaded on a client. Josh Tauber asked how this would scale, since every time a variable changes the corresponding agent has to inform all subscribers. Geoff replied that the

work done at Xerox by Schilit and Theimer on using multicast to limit update traffic was relevant here. Darrel Long asked if subscription is transitive, and was told that it was not. There was an extended debate about the meaning of "go back to the previous document", when the information used to generate the previous document might no longer be available. Jim Cumiskey observed that the previous document would still be in the local cache. Finally, in response to a question from David Steere, Geoff said that not much thought had yet been given to the issue of security.

4.3. Privacy & Anonymity

The first session after lunch was a panel on the topic of privacy and anonymity, chaired by Marvin Theimer. One of the panelists, Amir Herzberg, was unable to attend due a personal emergency.

The two panelists in attendance, Didier Samfat and N. Asokan, had independently addressed the same problem: that of ensuring the privacy and anonymity of a mobile user when he is far from the certifying agents he normally uses. Both approaches were based on using public key encryption to design authentication protocols in such a way that the mobile user's identity is not revealed to unintended parties. Didier presented a taxonomy of anonymity requirements and presented a solution based on the use of one-time aliases in authentication protocols in place of the user's real identity. Asokan advocated the notion of "limited disclosure of information" (regarding the user's real identity) to obtain practical anonymity.

The first important issue raised was "is this a relevant issue?" Amal Shaheen asked whether it is even desirable to provide anonymity. Asokan replied that it is a policy issue and the goal is to be able to provide the mechanisms necessary to make anonymity possible if it is desired.

Another question that was debated at length involved the reliability of such an approach. A failed home site, or intermediary, would leave the mobile user with no means of obtaining services. Asokan responded that it was no different from today, if a store's attempt to validate a VISA card failed. In addition, Didier observed that the standard reliability measures such as duplication of servers within a domain are taken to ensure that such essential services are highly available.

Josh Tauber asked who would pay for the cost of establishing and maintaining intermediaries. Didier replied that this would most likely be done by businesses, but may also involve customer payment. There was then an extended discussion about encryption, and the need to make its use more widespread.

The final topic of discussion involved possible abuse of location information. Marvin Theimer gave the example of a person whose path regularly goes past a pornographic store. While various conclusions may be drawn from this raw data,

an entirely innocent explanation is possible: the person may merely work next door. Many other similar examples were discussed by the panelists and the audience. After several examples of how location information can be abused, a majority of the audience was convinced that protecting privacy was indeed important. The examples revealed that mobility increases the possibility of abuse in two distinct ways: first, by permitting the perpetrators to work unsupervised in remote locations; second, by providing new types of information that can be abused. For example, if an insurance company obtains the cellular phone records of a customer it may be able to determine that he often exceeds the speed limit and should therefore have his rates increased.

4.4. Panel: Agenda for Developers & Researchers

The final part of the workshop was an opportunity for each participant to reflect on what he or she had heard over the previous sessions, and to brainstorm with a small group on four questions:

1. Where would you like the field of mobile computing to be in 5 years?
2. What can individual researchers do to influence the field?
3. What can industry do to make mobile computing profitable?
4. What are the 3 most important problems (technical or otherwise) to be solved for mobile computing to advance?

The participants were divided into five groups. Each group had a leader, whose primary responsibilities were to facilitate discussion and to bring the group back in time for the final panel session. The groups and their leaders were: "seal" (Terri Watson), "otter" (Murray Mazer), "whale" (Bob O'Hara), "dolphin" (Mary Baker), and "sealion" (Doug Terry). The groups had about an hour and a half to brainstorm, and many of them chose to hold their breakout meeting outdoors. After the breakout session, we reconvened and one member of each group reported its conclusions in the final panel session.

Ken Heflinger, representing "seal" group, spoke first. His group believed that in 5 years there would be powerful but affordable PDAs and that we would be living in a predominantly paperless world. The main thing researchers can do to influence mobile computing is to figure out how to make key technologies cheap, to work on fundamental technologies, and to combine things in useful and interesting ways. Regarding profitability, this group believed that using entertainment to whet people's appetites and exploiting advertising on PDAs as a source of revenue were two promising paths to paying for the mobile infrastructure. The three critical problems to be solved were perceived as: (a) realizing the global infrastructure (wireless everywhere, at low cost), (b) the development of software to deal with

heterogeneity, and (c) battery power limitations.

The next panelist was David Steere, representing the "otter" group. David said that his group was divided into two camps: Peter Honeyman (whose views were too colorful to be mentioned in a respectable publication!), and everyone else (whose views are reported here). The group felt that vertical applications (such as the one described by Nigel Davies earlier in the workshop) and mobile infrastructure would be pervasive in five years. The main things individuals could do to influence the field would be to help develop the infrastructure, demonstrate the feasibility of mobile applications, and help understand consumer needs. The perceived obstacles to commercial exploitation of mobility were: the development of vertical applications, the need to provide interoperability across a wide range of platforms, and the need for a wireless communication infrastructure.

Bob O'Hara then summarized the deliberations of the "whale" group. In five years, this group believed that the global wireless infrastructure would be deployed, that the killer application for mobile computing would have been discovered, and that many different mobile devices and gadgets would exist. Individual researchers could help influence the field by trying to use and deploy new applications; university-based researchers were viewed to be particularly well-positioned to contribute to the development of wireless testbeds. This group drew a blank on the issue of profitability. The three big obstacles foreseen by this group were: (a) imprecise disconnection semantics, (b) absence of a ubiquitous infrastructure, and (c) running out of radio spectrum (especially for small companies that can't bid high at FCC auctions). One of the group members, Barbara Liskov, made two additional observations. First, she observed that the problems of mobile computing were specializations of the problems that researchers have been addressing for many years in distributed computing. Second, she observed that mobile computing may require substantial revision to the fundamental primitives of distributed computing, such as RPC.

The next panelist was Barry Leiner of the "dolphin" group. In five years, this group expected mobile hardware to have advanced to the point where a "desktop in a pocket" would be a reality --- this would allow full sized screens and keyboards to be effectively "rolled up" for portability. But little global improvement was expected in the area of networks: they were expected to be sporadically available, of variable bandwidth, reliability, and heterogeneity. Integrated information access, the perceived "killer application" for mobile computing, would be pervasive. The group felt that the most effective way for individuals to influence the field would be via prototypes that opened users' eyes to new possibilities. Barry reiterated the point made earlier by Barbara Liskov, that many of the problems of mobile computing were really refinements of problems already encountered in distributed computing. He also reported that his group believed that a cooperative

approach, like the Internet, was the best way for industry to build an affordable mobile infrastructure and thus maximize profits. Finally, the three most pressing problems in mobile computing were perceived to be (a) power management (b) scale-related issues and (c) user-perceived complexity in dealing with enriched service abstractions in a resource-poor context.

The last panelist to speak was Jim Rees of the "sealion" group. In five years, this group expected wide coverage via high-bandwidth wireless communication, electronic commerce, and interoperability via open services. The group felt that the most effective way for individuals to influence the field was by developing better abstractions and metaphors for mobile adaptability, and by obtaining a better understanding of trade-offs. Regarding the question of profitability, the group had four suggestions. First, exploit cheap hardware. Second, treat mobility as a premium and charge higher for mobile services. Third, ensure easy access to mobile computing facilities. And, finally, support electronic commerce in mobile environments. In the opinion of this group, the biggest challenges facing mobile computing were: (a) the absence of a mobile infrastructure and facilities for billing (b) the need for adaptability and (c) the need for consistency (so that the mobile world is not drastically different from the desktop world).

After the reports by the panelists, the floor was opened for general discussion. Three major topics emerged. One discussion, involving John Saldanha, David Steere, Karin Petersen and Mary Baker, focused on the question of what kind of devices it made sense to render mobile. No real consensus emerged, and it was apparent that people held quite a divergence of opinions on this. Another discussion, involving Barry Leiner, Doug Terry, Peter Honeyman and Amal Shaheen, explored the claim that mobile computing was merely a special case of distributed computing. The consensus that developed was that many of the problems of mobile computing were indeed subsumed by distributed computing; but there are important differences. For example, location transparency is often a goal in distributed computing, whereas location awareness is a requirement in many mobile applications. The third discussion, between David Steere and Barry Leiner, examined the role of the entertainment industry in mobile computing. David expressed the view that we should bet on game manufacturers rather than computer manufacturers being the driving force behind mobility. Barry disagreed, and said that history showed that the entertainment industry was a follower not a leader. But its ability to create a giant market for communication and cheap hardware could be exploited to advantage.

5. Final Thoughts

The feedback I have received from many participants indicates that the workshop was quite a success. Many attendees felt that they had learned a lot at the workshop, and

were also able to contribute effectively to it. They were especially appreciative of the informal format, the small size of the audience, and the quality of the presentations and discussions. They confirmed that many thought-provoking discussions and ideas arose during the workshop. Quite a few of them inquired whether there would be a follow-on workshop in a year or two.

Success does not, of course, come by accident. Many people worked hard behind the scenes to ensure it. Crucial to success were the efforts of my colleagues on the program committee: Dan Duchamp, Peter Honeyman, Randy Katz, Jay Kistler, Krishan Sabnani, Amal Shaheen, Marvin Theimer, and Rich Wolff. They did an excellent job of reviewing and selecting papers on a tight schedule. They also did a great job of chairing the workshop sessions --- keeping things moving on time, but encouraging discussions.

Darrell Long, the General Chair, and the other organizers (Richard Golding, Peter Honeyman and Luis-Felipe Cabrera) must also be complimented for their efforts in putting together a high-quality event. My secretary, Marge Profeta, helped me in numerous aspects of the workshop. But, in the final analysis, it was the level of participation and enthusiasm exhibited by the attendees that made this such a productive and enjoyable workshop.

Obtaining the Proceedings

Copies of the full proceedings of this workshop will be available from the IEEE Computer Society after late March 1995. Its complete title is "Proceedings of the Workshop on Mobile Computing Systems and Applications", and its order number is PR06345. The publisher can be contacted via email at cs.book@computer.org, via fax at (714)-821-4641, and via phone at 1-(800)-CS-BOOKS.