Report of the visit to the University of Helsinki*

Hugo Miranda
Universidade de Lisboa

June 16, 2004

1 Introduction

This document reports my visit to the University of Helsinki, between the 7th and the 11th of June 2004. The visit had two goals which I believe were fully achieved: to discuss how selfishness behavior, trust and related security issues in mobile networks were being addressed and to strength the ties between the research communities of both universities.

During the visit I had meetings with several researchers of the University of Helsinki, namely: Professor Kimmo Raatikainen, Dr's. Jukka Manner and Markku Kojo and PhD. students Davide Astuti, Oriana Riva and Simone Leggio. I have also visited the Helsinki Institute for Information Technology (HIIT) where I met the Fuego Core project team, whose project manager is Sasu Tarkoma. All these meetings resulted in fruitful discussions concerning the various aspects of mobile computing.

Report Overview Resulting from the myriad fields of expertise of the researchers in both University of Helsinki and HIIT, a number of topics, ranging from the future applications of wireless networks to link level issues, have been addressed. The following sections of this report present the principal conclusions for each topic addressed during the visit. The report terminates with some general comments and conclusions in section 9.

*This visit was supported by the Middleware for Network Eccentric and Mobile Applications (MiNEMA) programme of the European Science Foundation.

2 Wireless Network Topology and Scenarios

Apparently, search and rescue and military operations have ceased of being the major motivation engines for ad hoc networks. The research projects taking place at the University of Helsinki are focusing on scenarios with a small number of nodes (up to one hundred) and where nodes are separated by at most two or three hops. The first stage of one particular project is even considering that all nodes are in direct reach of each other. These are adequate topologies for small applications where nodes are close. Examples of these situations are house networking and conferences. The application of ad hoc networks on the classical scenarios is being questioned. In particular, the suitability of these networks to military operations is not clear.

Expectations concerning 4G communication systems are that it will partially contribute to the wireless network ubiquity by providing very good connectivity at some places and close to none otherwise. It is likely that available bandwidth suffers a substantial increase.

Interface between wired and wireless networks The wired and wireless networks are substantially different network environments. However, it is desirable that connectivity between both is as transparent as possible for the user. This problem crosses vertically the entire network composition model. It emerged during the meetings as a subject of the link layer, network, transport and middleware levels.

The vast majority of currently standard network protocols was developed with the wired network environment in mind. Therefore, this problem is frequently presented from a wired perspective, where is up to the wireless community to mitigate the problems emerging at the air interface. One of the most common approaches is to develop mediators that can adapt the wired protocols to the wireless environment. This approach has been successfully pursued in different projects and is likely to continue to be used in the future.

3 Link Layer

The link layer can provide valuable contributions to the performance of wireless networks. Practical examples can be observed with the evaluation results of the L-SACP layer, that intermediates the link layer and the IP protocol. The high error rates presented by the wireless networks are enough to justify that some low-level protocol provides link layer reliability, possibly at expenses of including Forward Error Control mechanisms. In addition, this layer masks the high error rate to the transport layer, thus substantially im-

proving TCP performance. Although this layer has been evaluated in the context of satellite communications, it is clear that the approach could also benefit Wireless Local Area Networks. Simulations also provided interesting results in what concerns reliable data delivery. For typical error bursts in satellite communication, almost 100% reliability can be achieved at the link layer if unacknowledged frames are retransmitted once with forward error encoding.

It is common that the performance of new protocols can not be evaluated with field tests, either because the use of the medium shows to be particularly expensive (as is the case for satellite links) or because it is not possible to reproduce large scale environments. Simulation and emulation platforms play a fundamental role in the performance evaluation of protocols. In this case, it is important that the simulated environment mirrors as close as possible the real environment. Furthermore, it is important that the simulation framework provides as much feedback as possible concerning the various aspects of the tests and allow for the representation of as many scenarios as possible. Examples are the duration of the error bursts or the selective drop of key frames, necessary to evaluate worst case scenarios. Seawind is a network emulator for wireless links that addresses these problems, developed at the University of Helsinki. During the visit, I was able to discuss the capabilities of the simulator and also to observe and interact with the various tools that condition its behavior and interpret its output.

4 Network Layer and Mobility Support

The issue of mobility support was raised several times during the visit. Different approaches are being pursued to mitigate the known problems of IP address changes resulting from hand-overs. The Host Identity Protocol (HIP) is appearing as a viable alternative to the successively delayed Mobile IPv6. HIP separates the networking IP address (which is redefined as a low level reference) from the host identity (addressable by the application). The host identity is immutable, independently of node's mobility. HIIT is actively cooperating in HIP definition and has implemented an open source version for the Linux Operating System.

HIP can also address micro-mobility problems by using intermediaries. However, it does not fully address the efficient hand-off of Personal Area Networks (PANs). The IETF's NEMO proposal is seen as an approach going in the right track. It is also likely that commercial applications for this protocol emerge rapidly. One pioneer application is the wireless service provided by Lufthansa in its airplanes.

It is widely accepted that the configuration of wireless networks still pose significant challenges for the regular user. To generalize its acceptance, networks must become more user friendly by implementing some self-configuration mechanisms. Nodes must automatically locate and install relevant services and self adapt to the networking environment, for example, by configuring its addresses. This is a challenging task that has not been properly addressed until now.

5 Quality of Service

As the differences between wired and wireless Local Area Networks begin to vanish, the issue of Quality of Service is likely to emerge also for the mobile environment. Both share the same prospective scenarios: video and/or audio broadcasts for campus, conferences, etc. Under this kind of public interest events, it will make sense that some centralized authority manages the reservation of some resources, independently of the medium supporting the data broadcast.

6 Transport Layer

It seems clear that the existent transport protocols fail to adequately address the problems arising from the wireless environment. It has been shown that when competing against protocols without congestion control, like UDP, TCP performance degrades significantly. While this problem can be considered negligible in the wired Internet (where TCP accounts for more than 90% of the traffic), it acquires a great relevance in the wireless setting where the traffic shaping is expected to be different from that on wired networks. TCP-*friendly* mechanisms are being proposed. DCCP for example purposes to unify congestion control independently of the underlying transport protocol. However, the adoption of this or other proposals will always need to be handled with caution: as is, the Internet infra-structure relies on the congestion control capabilities of the TCP.

Other sources of performance degradation for TCP have been identified. Noisy channels and high or variable latency are two examples. The University of Helsinki has addressed this problem by masking errors at the Link Layer level (see section 3) and by proposing enhancements to the TCP protocol (F-RTO) which show to be more adequate in some environments and under particular circumstances. In this respect, the impossibility of finding a universal combination of TCP parameters that would show to be ideal for

any environment was discussed. However, improvements should be added as long as they do not disturb TCP performance on the remaining situations.

The management of TCP connections consumes resources at both communicating endpoints. For each connection, nodes must keep track of sequence numbers and window sizes and handle several timers. The Blocks Extensible Exchange Protocol (BEEP) helps to mitigate this problem by allowing the definition of several virtual channels between two communication endpoints multiplexed over a single TCP connection. Each stream can be individually configured and all may share some common properties, like authentication. BEEP proved to be a useful protocol in wireless infrastructured networks when mobile devices are required to handle simultaneous connections for several services.

7 Middleware

The research on middleware at both the University of Helsinki and HIIT is focusing on the adaptation of existing protocols to the wireless environment. Prospective targets of these protocols are new applications that emerge from the possible mobility of users.

7.1 SIP and SLP

The Session Initiation Protocol (SIP) and the Service Location Protocol (SLP) are currently hot topics that are being addressed by two distinct projects, one for ad hoc and the other for infra-structured networks. Both protocols rely on a myriad of servers, which cooperatively provide the protocol's functionalities.

Ad hoc networks pose a more challenging environment for their implementation due to the lack of a supporting infra-structure able to host the servers described in the specifications. It is therefore required that the participants in the ad hoc network share the responsibility of hosting them. However, the transient availability of the nodes imposes a distributed architecture, where information is highly replicated. In this respect, some issues concerning the balance between the efficient use of resources and the availability of up-to-date information have been discussed.

7.2 SOAP

The Simple Object Access Protocol (SOAP) is one example of how efficiency can be differently understood by the wired and wireless communities. SOAP

relies on XML, a particularly resource inefficient protocol from a wireless point of view. Standard XML is particularly verbose and hard to parse. This results in longer frames and additional computational power. The Fuego Core team is actively cooperating on the standardization of binary XML that would allow for a more efficient use of the resources.

The SOAP interoperability characteristics open interesting possibilities for the implementation of additional services. This protocol is presently in use as the underlying framework for other services, like an Event Service.

7.3 Event Services

The dissemination of events for mobile users implies the implementation of a complex infra-structure that must tackle with several particularities of the environment. Examples are event filtering and notification.

An overlay network can ease the burden of managing the subscriptions for a large number of providers and users geographically dispersed.

The implementation of a push event notification model is particularly challenging if only the resources of the transport layer and above are available. Significant signaling is used between the mobile devices and the infra-structure in the GSM protocol. However, this resource is not available to the middleware. Relying on the transport protocol for signaling implies that some TCP connection must be permanently open between the device and the trusted event service.

The implications of this solution and the apparent lack of alternatives were discussed. Keeping one open TCP connection between the event server and each mobile device increases the burden on the event service, who will be required to permanently manage a large number of connections that will be idle most of the time. Furthermore, this model is only suitable when billing is per data unit like in GPRS in opposition to billing by connection time like in GSM. Alternative solutions like the use of other protocols for notification of the availability of events (e.g. SMS) or reversed client-server with the device permanently listening for connection requests are either commercially of from a security point-of-view unacceptable.

7.4 Replicated File Systems

The implementation of a replicated file system was also pursued in the scope of the Fuego Core project. Besides the difficulties that such a system must tackle, in particular the need to keep the consistency of different file replicas, this implementation must also handle the limited bandwidth available for the mobile devices.

Efficiency can be achieved by limiting the amount of information that needs to be transferred between the different hosts. The determinism of the reconciliation algorithm permits that only differences be exchanged between the hosts. This has been putted to practice with the implementation of an algorithm that allows for the reconciliation of XML documents, which, in most cases, does not require user intervention. The usefulness of synchronous updates was questioned. Apparently. an on-demand synchronization provides a quality of service sufficient for the user, in opposition to a model where changes are automatically notified by the interested parties immediately after one of the replicas has been updated.

7.5 Context-awareness

Context-awareness is a key issue for the development of adaptive applications. It is expected that context information will be provided by a myriad of heterogeneous sensors. This information will become available in different representation formats and its exact meaning will vary, depending of the information provided by other sensors and by the goal aimed by the application.

This was an issue where more questions were raised than answered. The advantages and suitability of an unifying middleware framework that gathers and uniforms the output of a disparate set of sensors is consensual. Issues arose on the format of the output to be provided by such a framework. Even focusing on the most cited example of context information, the location, it is not easy to find an appropriate range or vocabulary that can be suitable for all applications. Low level references like pairs of geographical coordinates as those provided by GPS receivers may not be always available or may not provide suitable information for the application, who will be required to map them on a large number of more coarse-grained locations of reference. On the other hand, it is not clear if the definition of an ontology will be able to encompass all the diversity of information that may be required by applications. Even the meaning of simple words like "near" can vary significantly from application to application (and maybe from user to user!). Differences can be found on the desired granularity (1 or 10 meters), time (if the user is near home but in a traffic jam) or scale (near the refrigerator and near home).

The level and "intelligence" of the services that middleware frameworks should provide to the application have also emerged in the scope of the discussion above. This issue is not exclusive of the context-awareness middleware frameworks. It became clear that the border between middleware services and application responsibilities is not completely defined.

7.6 Distributed systems communication protocols

The mobile environment is not adequately suited to the implementation of resource demanding distributed systems communication protocols like total order or virtual synchrony. However, it is undeniable that such properties are useful for some applications that will emerge on the mobile scenario as it approaches the wired quality of service. Resource efficient versions of these algorithms can not be developed without appropriate feedback from lower level layers. However, it is important the gathering of such information does not consume more resources than the savings it provides. Other issues that also need to be addressed in this field are the appropriate handling of intermittent connections and an adequate balancing of the administrative tasks amongst the participants.

8 Security

Several distinct problems concerning security have been raised on different meetings.

Because infra-structured networks may have access to some trusted third-party, the problem of authenticating other nodes emerges mostly in ad hoc networks. Authentication is an important problem as it will rule the trust between two participants in an ad hoc network. Before exchanging files with restricted content in an unfriendly environment, two members of the same company need to authenticate each other. Authentication will provide the means for the establishment of a secure channel with two important security properties: it will provide guarantees that the file is delivered to the intended destination and it will prevent intermediary nodes, which will probably be required to forward the file, from snooping its content. Ideally, authentication should be integrated with middleware services required for participants and services location like the SIP and SLP protocols referred in section 7.1. Current research is seeking some off-line method for key distribution. Target scenarios are considering that before joining the ad hoc network, the nodes will be provided with the keys required to authenticate third party nodes.

Authentication is far from being a problem exclusively from ad hoc networks. Also wired and infra-structured networks can suffer from this problem. A classical example is Mobile IPv6, which has been facing successive delays due to security problems.

Another issue that needs to be addressed in the security domain is trust. Once again, and assuming that the supporting infra-structure is reliable and

trustworthy, the problem is more relevant in ad hoc networks. Some protocols can be particularly sensitive to the misbehavior of the nodes in the ad hoc network. This problem appears at the routing level (if nodes refuse to forward messages for other participants) and again at middleware protocols, for example in SIP, SLP and coordinator-based protocols. The problem will become more and more relevant as the prospective scale of the network increases. For a network model where all nodes are in direct reach of each other, messages from unauthenticated nodes can be discarded. In general, authentication and trust will be orthogonal issues if we consider that several authentication domains will coexist in ad hoc networks.

Privacy is also a hot topic for wireless networks. The problem lies in preventing unauthorized third-party from accessing content that needs to be disclosed to some nodes. This is the case of the event filters described in section 7.3, which should be public only to the overlay network that provides the service.

9 Conclusions

The University of Helsinki has a strong and dynamic research group on wireless networks. Although the majority of the research is focused on middleware, a few of its members are also addressing other subjects like link, network and transport layers. Its members have an active participation in standardization bodies like IETF and W3C and the ability to bring together to research projects competing companies like Nokia and Ericsson or Elisa and Sonera.

During the visit I had the opportunity to meet with both established researchers and PhD. students. This mixture showed up to be adequate to consolidate background and seminal knowledge in the field and to discuss future directions, research options and state of the art. It was clear that mobile computing is a research field with much more open questions than definitive answers.

The visit created the ties required to initiate fruitful collaborations. Links between the work being presently addressed in middleware for ad hoc networks by Jukka Manner and Simone Leggio and my present field of research can be easily established and experiences are being shared. We are currently seeking performance and energy improvements for the dissemination of information in the SIP and SLP protocols based on probabilistic protocols. We expect to publish the outcome of this collaboration in some international conference by the end of the year.

Acknowledgments

The author wish to thank to Davide Astuti, Jukka Manner, Kimmo Raatikainen, Markku Kojo, Oriana Riva, Sasu Tarkoma, Simone Leggio and to the members of the Fuego Core project team for the warm reception and interesting discussions.