AN ARCHITECTURAL FRAMEWORK FOR HETEROGENEOUS NETWORKING

Glenford Mapp[a], David N. Cottingham[b], Fatema Shaikh[a], Pablo Vidales[c],
Leo Patanapongpibul[d], Javier Balioisian[e], Jon Crowcroft[b]

E-mail: g.mapp@mdx.ac.uk, david.cottingham@cl.cam.ac.uk, f.shaikh@mdx.ac.uk, pablo.vidales@telekom.de,
leo.p@cantab.net, javier.balioisian@ericsson.com, jon.crowcroft@cl.cam.ac.uk

[a] *School of Computing Science, Middlesex University, Middlesex NW4 4BT, UK*
[b] *Computer Laboratory, University of Cambridge, Cambridge CB3 0FD, UK*
[c] *Deutsche Telekom Laboratories, 10587 Berlin, Germany*
[d] *Samsung Electronics Research Institute, Staines, Middlesex TW18 4QE, UK*
[e] *Network Management Research Centre, Ericsson, Ireland*

Keywords: Heterogeneous Networking, Vertical Handover, OSI Model, Mobility

Abstract: The growth over the last decade in the use of wireless networking devices has been explosive. Soon many devices will have multiple network interfaces, each with very different characteristics. We believe that a framework that encapsulates the key challenges of heterogeneous networking is required. Like a map clearly helps one to plan a journey, a framework is needed to help us move forward in this unexplored area. The approach taken here is similar to the OSI model in which tightly defined layers are used to specify functionality, allowing a modular approach to the extension of systems and the interchange of their components, whilst providing a model that is more oriented to heterogeneity and mobility.

1 INTRODUCTION

Mobile devices with several interfaces such as new mobile phones and PDAs will soon become commonplace. We define the networking issues associated with this type of device as **heterogeneous networking** and the devices themselves as **hetnet** devices.

Heterogeneous networking poses many challenges in several areas. At the lowest levels, many new access technologies including 3G, WiMax and UltraWideBand (UWB) will be supported on hetnet devices. This expansion of the link and MAC layers has been accompanied by a contraction in the network layer as core networks use IP packets to facilitate key services including telephony, data, and multimedia.

One of the major capabilities of heterogeneous networking is that of **handover**. This is necessary because the networks that are currently available and/or their points of attachment may be changing as a mobile node changes its location. Another key issue that needs to be addressed because of heterogenous networking is that of Quality of Service (QoS) in peripheral networks. This is because different wireless networks have varying QoS so vertical handover not only affects the point of attachment but also the QoS of the link as seen by other entities higher up the protocol stack. This in turn affects the ability of the network and transport services to deliver effective performance since these systems must respond to changes in

QoS in the available channels.

At the higher layers how these QoS changes are dealt with and what can be done by the system to minimize their effects on applications must also be examined. Alternatively, future applications will be able, with the help of the system, to structure themselves to make use of QoS changes in wireless networks. In addition, new kinds of application environments which can facilitate new features such as personalisation, personal area networks and location-based services will be built.

In order to address the issues above, the authors believe that there needs to be a **framework** that encapsulates the key challenges of heterogeneous networking for mobile devices. Like a map clearly helps one to plan a journey, a framework is needed to help us move forward in this unexplored area. The approach taken here is similar to the OSI model in which layers are used to specify functionality, allowing a modular approach to the extension of systems and the interchange of their components such that different implemenations can be easily explored. The authors are implementing a prototype testbed to explore the proposed architecture.

The rest of the paper is structured as follows: Section 2 looks at the OSI model which has been used as a communication architecture. Section 3 describes the architectural framework. Section 4 describes the Cambridge Testbed; Section 5 describes work in

progress, while Section 6 describes the issues that are still to be investigated. The paper concludes with a summary in Section 7.

2 The OSI model

The most well-known framework in data communications is the OSI model (Zimmermann, 1988) developed by the CCITT. This is regarded as a reference model, and not an implementation plan, but it clearly delineates the functions of the layers to provide a framework for exchanging data between networked applications.

The authors believe that the OSI model cannot be considered the dominant model for heterogeneous networking. It is not wrong; it is simply inadequate for heterogeneous networking for several reasons. Firstly in the OSI model, the first three layers, the physical, link and network layers are concerned with the movement of packets between networks. The higher layers of the OSI model – the transport, session, presentation and application layers – are designed to deal with end-to-end issues between application processes. Such a horizontal separation between networking and End-to-End Issues is no longer sustainable in heterogeneous networks. There are several reasons for this.

In the OSI model, the network is essentially used to simply forward packets to their relevant destination. However, in heterogeneous networking, there are new functions that the network must also support. One of these is vertical handover. The point here is that vertical handover requires frequent and intimate communication between the mobile host and the network which cannot be simply incorporated into the OSI model. To go further, vertical handover also involves the reconfiguration of certain parameters in the network, such as allowing for the reservation of resources to ensure quality of service in the receiving network. Such a reality is difficult to model in the OSI context.

The other major observation is that the OSI model works well when the characteristics of the networks at the edge are not overly dissimilar from the core network. The early Internet had Ethernet or Token ring systems which were basically wired networks capable of several megabits per second between endpoints. However, looking at network trends since that time, the core network and end systems have taken different evolutionary paths. The core network is being made faster with the use of technologies such as MPLS and single-mode fibre optics, while the peripheral systems are rapidly becoming dominated by the emergence of wireless technologies that have very different characteristics in terms of latency, bandwidth, availability and error distribution properties.

In the light of these observations, the authors believe a new framework is required to better reflect how tomorrow's heterogeneous networks should be constructed. This new model is is described in the following sections.

3 THE ARCHITECTURE

3.1 Layering: a Conceptual Framework

As outlined in the previous section, the OSI model does not act as a good conceptual framework for today's networks. A key point is that this does *not* invalidate a layered approach, but rather indicates that it must be updated. Such an approach must specify the hierarchy of functionalities, but not prescribe detailed interfaces between them. The ordering of the hierarchy is perhaps the most important consideration; this is evidenced by the difficulties that have been experienced in placing, for example, vertical handover functionality above the transport layer.

We propose an architecture that, like the OSI model, has seven layers, but that uses a novel hierarchy of functionalities (Figure 1). In our model, vertical handover, with input in turn from policy management, is placed *below* the network transport layer. Similarly, quality of service is given its own specific layer to separate it from both the application and the network transport modules. This contrasts with current approaches where QoS appears to be more of an "add-on" rather than an integral part of the network stack.

It is also important to realise that the layering paradigm does not restrict implementors to rigid modularisation. A framework is a useful concept, rather than a detailed design specification. Hence, whilst our model details the necessary *ordering* of functionality, it does not discount the possibility of cross-layer approaches. Indeed, as detailed in Section 4, policy management and vertical handover functionality might well be integrated into a single component. Also, higher layers may well need context not only from the layers directly beneath them in the hierarchy, but also those further down. This can be seen to be the case for the QoS layer, where hints from the network abstraction layer would be of great utility.

Hence, we emphasise the need for clear conceptual separation, whilst leaving open the possibilities for vertical integration and trans-layer interfacing in terms of implementation. In the following sections we proceed to describe each of these separate conceptual layers in turn, using a bottom-up approach.
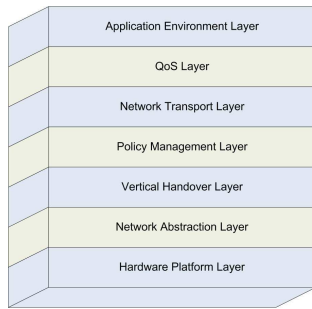
Figure 1: The architectural model

3.2 Hardware Platform Layer

This layer is used to define the hardware components and technologies required to support wireless networks. This layer defines several characteristics, for example, the electromagetic spectrum required for a given technology, the different modulation techniques that may be used, as well as the MAC algorithms for acquiring and reserving channels. It is recognised that individual systems may be totally incompatible with each other. Hence, the layer is composed of vertical sub-layers, with each sublayer representing a particular network, e.g. 3G, WLAN, WiMax, etc.

3.3 Network Abstraction Layer

The Network Abstraction Layer is used to support different networking technologies using a common interface. Different wireless device drivers will eventually be written to map onto this layer. The network abstraction layer has to do with controlling and maintaining the network on the mobile node. Recently, the IEEE convened the 802.21[1] working group to examine the possibility of standardising the interface to different wireless MACs. This work is relevant to the network abstraction layer of our architecture.

3.4 Vertical Handover Layer

Vertical handover is clearly a key component of heterogeneous networking (McNair and Zhu, 2004). There are two distinct approaches to vertical handover. The first is the network-controlled approach in which the network decides when and how the handover will occur. This means that there are mechanisms in the network that maintain all relevant information on the mobile host, including its location relative to different networks, their signal strengths at the mobile node's location, and its direction and speed. However, we do not believe this approach is scalable.

---

[1] http://www.ieee802.org/21/

As new wireless networks are added, the information about all the networks to which the host is currently attached becomes difficult to maintain. Such an approach also relies on operators sharing detailed information about their networks, a concept that is unlikely in the current cellular telecommunications environment.

The second approach is termed **client-based handover**. With this approach, handover is controlled by the mobile device. There are clear advantages to using this approach. Firstly the mobile node will keep up to date information about its network interfaces and therefore is in a superior position to decide when handover should occur. Secondly, the mobile node will also be aware of the state of its TCP connections and other higher level issues, and therefore these factors can be taken into account when making the decision to conduct a vertical handover. Finally, by removing handover functions from the core network, it should then be more cost effective to build increasingly robust core systems.

3.5 Policy Management Layer

A policy management system is needed to evaluate all the circumstances when handover should occur, taking into account various factors such as changes in coverage and signal strength, the state of the network, and the state of any transport connections associated with the mobile host.

There are two types of policy management which are being explored: reactive and proactive. A reactive policy depends entirely on notification from the network abstraction layer about the presence or absence of networks as the hetnet device changes its location. Almost all the policy management systems so far reported in the literature (e.g. POLIMAND (Aust et al., 2003)), including the most recent, PROTON (Vidales et al., 2005a), have been reactive.

Proactive policies attempt to acquire and use information about the likely coverage and signal strength before the mobile actually reaches a given location. A key parameter sought from proactive policies is the Time Before Vertical Handover or TBVH. Knowing the TBVH allows the higher layers of the protocol stack to make maximum use of channels that may soon be unavailable.

Proactive policies can be divided into two types: A proactive-knowledge-based policy mechanism makes explicit use of knowledge of the mobile node's location and knowledge about the types of coverage and signal strengths of each network that the hetnet device is likely to encounter at that location. These systems require enhanced location detection and monitoring routines as well as a mechanism for ensuring data on what networks are available at a given location is accurate and easily accessible. Such data must be

available in a format that allow it to be processed by hosts with limited resources, and transmitted over potentially low bandwidth networks without significant impact. A few systems (Laasonen et al., 2004)(Soh and Kim, 2004) have utilised data collected from previous journeys.

Another approach is a proactive-modelling approach in which a mathematical model is used to determine the TBVH based on simple geometric calculations. Though less accurate than the proactive knowledge approach, this approach is flexible and can be used in simulations as well as in real networks.

There is also a growing need to combine proactive and reactive policies, with a view to developing an architecture where it is possible to choose which policy to use in a given situation. Hence, when there is accurate information then a proactive policy may be used. However, when there is no coverage data or the data is unreliable, then the system can fall back to a reactive policy mechanism.

3.6 Network Transport Layer

This layer concerns functions that would normally be assigned to the network and transport layers of the OSI model. Hence this layer examines addressing, routing and transport issues in peripheral networks.

The current opinion is that all networks whether in the core or on the periphery should be using TCP/IP. This thinking has been reinforced by End-to-End arguments which have been used throughout the architectural discussions when the Internet was designed (Saltzer et al., 1984). The current evolution of the Internet questions some of these End-to-End arguments. As indicated previously, today the Internet is evolving into a very fast core network with mobile networks on the periphery. This means that characteristics of the core and the periphery are diverging in terms of latency, throughput and error profiles.

In the light of this, the assumption that TCP/IP should be used in peripheral networks for heterogeneous networking needs to be carefully re-examined. Firstly we should question whether IP should be used in such networks. An assumption of the current IP infrastructure is that every machine should have a globally unique IP address to use on the network. This has, however, been directly challenged by the success of Network Address Translation (NAT) techniques. In NAT, a private IP address space is employed in the peripheral network while only a few global IP addresses are used to actually communicate with other machines on the Internet. The NAT software then provides the translation between the global IP connection and the local machine with its private IP address. Because all datagrams must traverse the server performing the NAT, it provides a point in the network where incoming packets can be analysed and filtered as necessary.

In addition, it increases security by *not* making local machines visible on the Internet, thus reducing the potential for targeted security flaw exploits, and DoS attacks against specific machines.

The success of NAT – which is considered to be far from ideal by network purists, including the authors – questions the assumption that all machines should be assigned a globally unique IP address. NAT makes the case that IP addresses should be confined to moving data within the core network. In the peripheral networks some other form of local addressing may be used with the translation between the networks taking place at the local gateway. Such an approach, if deployed, will question the (often challenged) assumption that there are insufficient IPv4 addresses. Though the authors support the deployment of IPv6, the key requirement for its deployment – to provide an infinitely large global address space – needs to be re-examined in the light of new realities.

3.6.1 TCP – Found Wanting

It is clear that TCP is unsuitable for wireless networks (Meyer, 1999)(Xylomenos et al., 2001). This is primarily due to the fact that TCP interprets packet loss as exclusively due to congestion and reacts by substantially decreasing its send rate, and then employing its slow start mechanism. While such a conclusion may be valid for wired systems such as the core network, peripheral wireless systems continuously lose packets due to channel fading, interference, vertical handovers, and other related effects. Most of these transients are temporary and unrelated to network congestion.

There have been several attempts to modify TCP in the light of these effects, such as those described in (Balakrishnan et al., 1997) and (Chandra et al., 2003). Recently, there has been a move towards not modifying the TCP protocol engine but making it more responsive to temporary network outages (Scott and Mapp, 2003). While this is useful, a clear, generally applicable, and elegant solution has not been found.

3.7 The Case for Network Plurality and Application Conformity

The idea that a different networking infrastructure runs in peripheral networks brings with it many challenges. Most importantly, the ability to translate to different naming and addressing schemes as packets are transmitted through different networks. Some of issues are addressed in Plutarch (Crowcroft et al., 2003) by the introduction of **contexts** and **interstitial functions**. The key requirement is the development of a framework where different networks interwork to

the benefit of the applications running on those different networks.

However, in order to ensure that every application developed using TCP/IP will not be required to be rewritten, the issue of application conformity must be addressed. An interesting solution is to pursue the idea of TCP/IP not only being a real protocol suite in the core network but also a **protocol interface** in peripheral networks. This means that on end user devices in these networks, the TCP protocol engine forms an overlay above whatever networking protocol is actually being used on the network. Hence, applications can assume TCP behaviour whilst the actual protocol in use need not be TCP. We believe that this approach should allow network plurality to emerge whilst maintaining application conformance.

Though application conformance may be necessary in the short and medium terms, in the long term we believe that the we need to move to a situation where an application's transport requirements are not specified by the selection of a specific protocol, but by the definition of a QoS vector which specifies the transport requirements of the application, as suggested in (Mapp and Hodges, 1997).

3.8 Quality of Service Layer

Because the hetnet device will be using multiple networks, it will inevitably experience different qualities of service (Rodriguez et al., 2004). Therefore, a QoS layer is required to manage this variability. To cover this appropriately we now introduce the concepts of downward and upward QoS.

3.8.1 Downward QoS

The intention is to construct a series of mechanisms to handle the different qualities of service that are encountered in vertical handovers in heterogeneous networks. In order achieve this a mechanism is needed that bundles connections over the different available channels, which themselves may be varying. We term this **Downward QoS**, which will be required to support legacy applications.

3.8.2 Upward QoS

It is hoped that future applications will be able to react to QoS changes as the hetnet device changes its location. These changes will be provided via the QoS layer. This is termed **Upward QoS**. In this case, an event-based QoS signalling mechanism can be used to inform applications of QoS changes. Applications, when they begin executing, would be able specify routines that should be called in response to event notifications by the QoS layer. This is similar to the X Windows System (Scheifler and Get-

tys, 1986) in which clients of the X-Server can specify what routines should be called on being notified about certain events concerning windows they have created on the screen. In the networking community, the TRIGTRAN (Dawkins, 2003) and PILC (Karn, 2003) projects illustrate this paradigm, by lower layers providing hints to higher layers.

3.8.3 QoS-Aware Middleware

A hugely useful component in the support of Upward QoS is the development of a QoS-aware middleware platform for Distributed Environments. The concept is to take a well known environment such as CORBA and add support for Upward QoS capability natively to the architecture. This would allow distributed applications to work seamlessly over heterogeneous mobile networks.

3.9 The Application Environments Layer – A Toolkit Approach

Some application environments attempt to encapsulate several architectural layers (Niebert et al., 2004). These systems are therefore so large that compatibility with other systems is not really considered important. This is an issue if we desire true connectivity between all devices and applications.

An alternative approach is to adopt a **Toolkit** philosophy in which the principal goal is not to build a particular application environment but to specify the components of a toolkit which would aid different groups in building application environments. Hence there would be a degree of compatibility which would also encapsulate the functionality of the lower layers of the architecture.

4 ARCHITECTURE TESTBED

The original ideas expressed in the framework architecture have been explored by the development of a testbed in the University of Cambridge Computer Laboratory. The testbed is depicted in Figure 2. Unlike many other testbeds, our setup makes use of connections to commercial providers' networks, allowing us to evaluate the performance of Mobile IPv6 in realistic scenarios, i.e. production environments. The architecture of the testbed is outlined in greater detail in (Vidales et al., 2005b).

This advantage has enabled us to carry out research into both the performance of cellular networks for data transfer (Chakravorty et al., 2002), and also much work on the optimisation of vertical handovers using client-based techniques, e.g. by

RA caching (Vidales et al., 2003) and BU simulcasting (Chakravorty et al., 2004), as well as soft handovers.

Recent work has focused on the latency of TCP's adaptation to new networks on vertical handover (Cottingham and Vidales, 2005) and policy systems that sense the environment and take handover decisions based on such data. Future work will concentrate on the integration of new and emerging technologies, such as UMTS (3G) and WiMax.

In terms of policy management strategies for vertical handover, the original work implemented a reactive policy mechanism called PROTON (Vidales et al., 2004) which was written in PONDER, a policy specification language. PROTON was implemented as a three-level subsystem based on a handover execution layer, which performed the handover. A policy layer allowed policy rules to be specified and was the mechanism that reached a decision to implement handover. Finally PROTON possessed an input/output layer which captured relevant events and triggers and fed them into the policy layer.

The computational requirements of PROTON exceeded that of ordinary mobile devices, so an autonomic system was developed where PROTON ran on machines in the network but downloaded a finite-state machine to the mobile node which encapsulated these policies in a resource-efficient fashion.

5   CURRENT WORK

At present, work is being done to develop proactive policy mechanisms. At the University of Cambridge coverage maps of WLAN, GPRS and 3G networks within the City are being built. The idea is to be able to ascertain the coverage expected at particular locations. This work will then be used to develop algorithms that allow users to specify a constant or minimum bandwidth they wish to be provided with on a journey, for which the system will recommend an appropriate route.

A proactive system based on mathematical modelling is being pursued by the Networking Research Group at Middlesex University (Shaikh et al., 2006). This looks at simple ways of developing mathematical models of Time Before Vertical Handover for upward handover scenarios where a WLAN network is currently in range but will be out of range based on the velocity and trajectory of the mobile node, whereupon the mobile node will then perform a vertical handover to a UMTS system. The model introduces the concept of a Boundary Base Station (BBS) and determines the TBVH for the node leaving the wireless network at that BBS. The calculation of TBVH is done using parameters that can be easily determined by current net-

work deployments.

Work has also begun on examining End-to-End Transport issues. The approach taken is to first look at developing a flexible method of specifying networks, and the need for certain defined characteristics such as addressing and naming. Work is being carried out to develop a precise definition of a context as well as interstitial functions (Crowcroft et al., 2003).

In terms of the QoS layer, a mechanism to exploit downward QoS has already been specified. This involves the development of a Stream Bundle Management or SBM Layer (Shaikh et al., 2005). This layer attempts to map application QoS requirements to the available channels as the policy management layer indicates the availability of certain networks as well as the TBVH parameter. The major component of the SBM layer is a resolution matrix in which for each network there is a unique network id, a status indicator showing off or on, the average available bandwidth, the received signal strength, the time before vertical handover and the round trip time. The resolution matrix is used to map application traffic management requirements onto available network channels.

6   FUTURE WORK

Future work involves investigating the implemention of the upper levels of the architecture. We would like to explore issues in the upward QoS approach, examining how QoS events are generated from the lower layers. These efforts also need to be compatible with what is being done in the TRIGTRAN and PILC working groups.

In addition the use of these mechanisms to build a QoS-aware Middleware will also need to be explored. We are interested in starting with a version of omniORB[2], attempting to add the events mechanism into the CORBA architecture. In terms of application environments, we intend to port the testbed to a commonly used mobile operating environment such as Microsoft Windows Mobile 5 or the Symbian OS. Finally there is a need to devise and implement a security plane for this architecture which encompasses many layers of the framework.

7   CONCLUSION

This paper has proposed a new framework for heterogeneous networking. The framework allows much greater flexibility than previous models, whilst maintaining a clear separation between the different layers. It is fully recognised that there is still much to
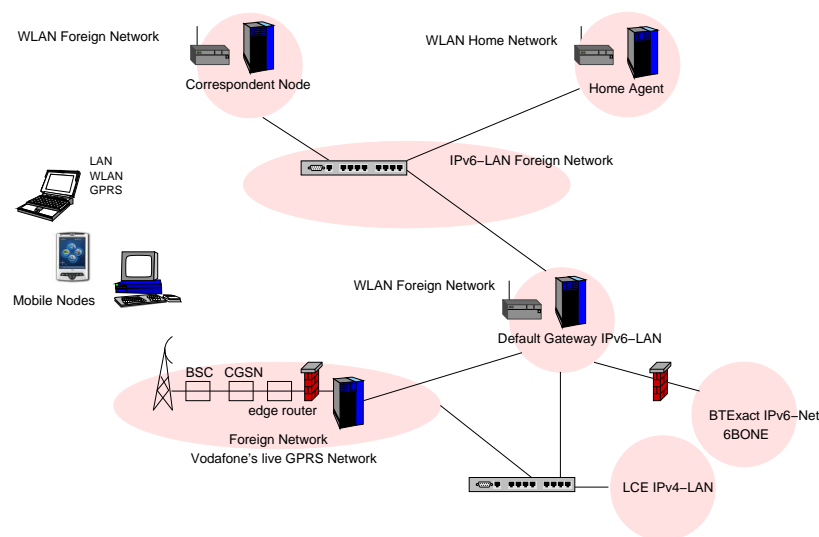
---

[2]http://omniorb.sourceforge.net

Figure 2: Structure of the Cambridge Wireless MIPv6 Testbed

do. Hence the authors would welcome feedback on the contents of this paper.

REFERENCES

Aust, S., Fikouras, N., Protel, D., Gorg, C., and Pampu, C. (2003). Policy based mobile ip handoff decision (polimand). Technical Report draft-iponair-dna-polimand-00.txt, Work in Progress, IETF.

Balakrishnan, H., Padmanabhan, V. N., Seshan, S., and Katz, R. (1997). A comparison of mechanisms for improving TCP performance over wireless links. *IEEE/ACM Trans. On Networking*, 5(6):756–769.

Chakravorty, R., Cartwright, J., and Pratt, I. (2002). Practical experience with TCP over GPRS. In *Proc. IEEE GLOBECOMM*.

Chakravorty, R., Vidales, P., Subramanian, K., Pratt, I., and Crowcroft, J. (2004). Performance issues with vertical handovers – experiences from GPRS cellular and WLAN hot-spots integration. In *Proc. IEEE PerCom*.

Chandra, D., Harris, R., and Shenoy, N. (2003). TCP performance for future IP-based wireless networks. In *Proc. IASTED Wireless and Optical Communications Conference*.

Cottingham, D. N. and Vidales, P. (2005). Is latency the real enemy in next generation networks? In *Proc. ICST CoNWiN*.

Crowcroft, J., Hand, S., Mortier, R., Roscoe, T., and Warfield, A. (2003). Plutarch: an argument for network pluralism. *SIGCOMM Comput. Commun. Rev.*, 33(4):258–266.

Dawkins, S. (2003). End-to-end, implicit 'link-up' notification. Technical Report draft-dawkins-trigtran-linkup01 (work in progress), IETF.

Karn, P. (2003). Advice for internet subnetwork designers. Technical Report draft-ietf-pilc-link-design-15 (work in progress), IETF.

Laasonen, K., Raento, M., and Toivonen, H. (2004). Adaptive on-device location recognition. *LNCS 3001*, pages 288–304.

Mapp, G. and Hodges, S. J. (1997). QoS-based transport. In *Proc. IFIP Workshop on Quality of Service*.

McNair, J. and Zhu, F. (2004). Vertical Handoffs in Fourth-Generation Multinetwork Environments. *IEEE Wireless Communications*, 11(5).

Meyer, M. (1999). TCP performance over GPRS. In *Proc. IEEE WCNC*, pages 1242–1252.

Niebert, N., Schieder, A., Abramowicz, H., Malmgren, G., Sachs, J., Horn, U., Prehofer, C., and Karl, H. (2004). Ambient Networks: An Architecture for Communication Networks Beyond 3G. *IEEE Wireless Communications*, 11(2).

Rodriguez, P., Chakravorty, R., Chesterfield, J., Pratt, I., and Banerjee, S. (2004). MAR: A Commuter Router Infrastructure for the Mobile Internet. In *Proceedings of the ACM Second Mobile Systems, Applications and Services Conferences (ACM Mobisys 2004)*.

Saltzer, J. H., Reed, D. P., and Clark, D. D. (1984). End-to-end arguments in system design. *ACM Trans. Comput. Syst.*, 2(4):277–288.

Scheifler, R. W. and Gettys, J. (1986). The X window system. *ACM Trans. Graph.*, 5(2):79–109.

Scott, J. and Mapp, G. E. (2003). Link layer-based TCP optimisation for disconnecting networks. *Computer Communications Review*, 33(5).

Shaikh, F., Lasebae, A., and Mapp, G. (2005). SBM layer for optimum management of co-existing telemedicine traffic streams under varying channel conditions in heterogeneous networks. In *Proc. MIT*.

Shaikh, F., Lasebae, A., and Mapp, G. (2006). Client-based SBM layer for predictive management of traffic flows in heterogeneous networks. In *Proc. IEEE ICTTA*.

Soh, W.-S. and Kim, H. (2004). Dynamic bandwidth reservation in cellular networks using road topology based mobility predictions. In *Proc. IEEE INFOCOM*.

Vidales, P., Baliosian, J., Serrat, J., Mapp, G., Stajano, F., and Hopper, A. (2005a). Autonomic system for mobility support in 4G networks. *IEEE Journal on Selected Areas in Communications*.

Vidales, P., Chakravorty, R., and Policroniades, C. (2004). PROTON: A policy-based solution for future 4G devices. In *Proc. IEEE International Workshop on Policies for Distributed Systems and Networks*.

Vidales, P., Mapp, G., Stajano, F., Crowcroft, J., and Bernardos, C. J. (2005b). A practical approach for 4g systems: Deployment of overlay networks. In *Proc. TRIDENTCOM*, pages 172–181.

Vidales, P., Patanapongpibul, L., and Chakravorty, R. (2003). Ubiquitous networking in heterogeneous environments. In *Proc. 8th International Workshop on Mobile Multimedia Communications*.

Xylomenos, G., Polyzos, G. C., Mähonen, P., and Saaranen, M. (2001). TCP performance issues over wireless links. *IEEE Communications Magazine*, 39(4):52–58.

Zimmermann, H. (1988). OSI reference model – the ISO model of architecture for open systems interconnection. *Innovations in Internetworking*, pages 2–9.