Ubiquitous Networking in Heterogeneous Environments

Pablo Vidales (1) and Leo Patanapongpibul (1) and Rajiv Chakravorty (2)
(1) Laboratory for Communication Engineering, University of Cambridge Engineering Department, U.K.
Email: pav25,lbp22@eng.cam.ac.uk
(2) Systems Research Group, University of Cambridge Computer Laboratory
Email: rajiv.chakravorty@cl.cam.ac.uk

Abstract: This paper discusses the challenges in networking posed by the emerging field of ubiquitous networking and the deployment of heterogenous wireless networks. It reveals three fundamental problems in future networking: network interface management, interface co-existance, *and* support for adaptive mobility. *It describes a testbed to conduct experiments in heterogeneous environments, including low and high mobility scenarios. It presents a scheme that improves handoff performance in an overlay network model. The paper concludes with a discussion of the work in progress, and the key issues that need to be addressed for achieving ubiquitous networking.*

1 Introduction

Mobile devices and deployment of new wireless access networks and services will continue to proliferate in the next few years. By 2005, it is expected that there will be a billion mobile communication devices worldwide; sixty percent of Internet access will be from mobile devices and half of the traffic on mobile networks will be data [1]. Market research indicates that capability to access services anywhere-anytime through different wireless access networks using a mobile communication device will attract significant attention.

Since 1991, Mark Weiser's vision [2] of *ubiquitous computing* – the deployment of environments where computing and communication capabilities are highly immersed with users – has evolved in number of forms: cellular phones, laptops, and PDAs with various wired and wireless access networks. However, networking capability in mobile scenarios is still limited, and ubiquitous networking is not yet a reality.

What would it be to have ubiquitous access to services on-the-move? We can imagine a world where John, who moves from his office in Cambridge, U.K. to the company headquarters in New York, travels through heterogeneous environments before reaching his destination. During his journey, John uses his mobile device to access different information services using a range of wireless technologies. A primary challenge for this to come true is that John needs to maintain networking capability all the time (while he is driving, on the train, and at the airport): *ubiquitous networking is a fundamental requirement in future environments.*

Of course, there are many issues to address before this becomes a reality. The example above simply demonstrates that evolution of ubiquitous networking will demand integration between diverse network systems, applications, and services. Seamless mobility in hybrid and integrated environments is an interesting and open research problem that needs to be addressed. However, for this to happen, it demands a seamless inter-network handoff mechanism to keep connectivity as devices move across environments, while still minimising any disruption to ongoing flows during switchovers. A mechanism that enables this has to exhibit a low handoff latency, incur little or no data loss (even in highly mobile environments), scale to large inter-networks, adapt to different environments, and act as a *sealant* between heterogeneous environments and technologies without compromising on key issues related to security and reliability.

Indeed, the complexity of mobile application services will increase exponentially with users' demand. Future Mobile Broadband Services [20] will encapsulate service areas such as Tele-Medicine, Intelligent Transport Systems (ITS), and Mobile Geographical Information Systems (mGIS) that will demand an *always-on* connectivity with high Quality of Service (QoS) requirements. To support these services, what is needed is a mechanism that can allow seamless mobility between technologies and environments.

A traditional and well-known approach to address this issue is the use of Mobile IPv4 (MIPv4) protocol (see [11,13]). The protocol can tackle both local-area and wide-area mobility across wired and wireless networks. In recent years, MIPv4 has emerged as a primary driver to enabling seamless roaming, mainly due to the continous expansion of the mobile Internet, and specifically due to its compatibility with IP. However, there exist certain limitations in MIPv4 that can be overcome using Mobile IPv6 (MIPv6), such as: no foreign agents, avoidance of the traditional problem linked to triangular routing and tunneling, and also security issues [13]. However, this still does not solve all problems. Movement detection mechanism in MIPv6 [3] is not particularly suited for inter-network handoffs [7]. Hence, certain improvements remain to be seen.

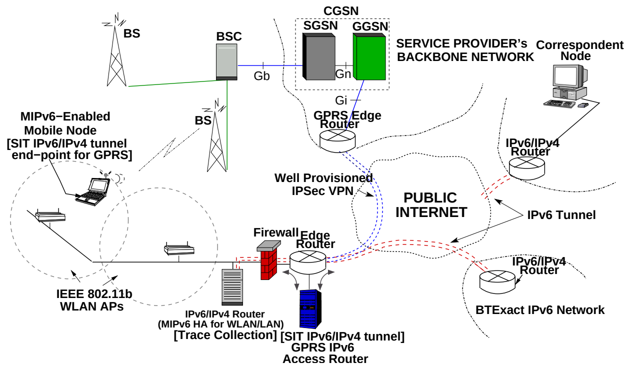

Figure 1: **MIPv6-based GPRS-WLAN Testbed Implementation.**

These include:

- **Network Interface Management.** In a multimode mobile device – device with multiple air interfaces to access different wireless networks – a single network interface cannot provide the connectivity (and desired QoS) at all times. Besides, the current version of MIPv6 also does not permit the assignment of the home address to multiple network interfaces in a mobile device. This means that management of multiple network interfaces will have to be controlled by a handoff scheme supported in the mobile device.

- **Interface Co-existence.** Multiple network interfaces can co-exist to improve QoS of the end-to-end connections and to minimise handoff effects in heterogeneous environments. A mobile device might benefit by using multiple interfaces, especially in overlap areas when it is about to handoff. It can make use of multiple links simultaneously to achieve better data rates and mitigate the effects of vertical handoffs. However, augmenting this feature into mobile devices is not simple, new rules will have to be defined and the MIPv6 protocol extended to support such handoffs.

- **Support for Adaptive Mobility.** Mobile hosts can have different sets of rules/policies that are constructed based on certain environment conditions. For example, mobility support in an office environment is not as critical as in a highly mobile enviroment such as when travelling in a train or a car. Hence, mobile clients should be able to decide *if* a transparent mobility management is required and to *what* extent.

We have implemented a testbed to conduct experiments to evaluate these issues, and several others endemic to mobility in heterogeneous environments. Our current investigation over the testbed involves the evaluation of a client-assisted handoff scheme. Furthermore, work-in-progress deals with some other network-supported and proxy-assisted enhancements.

2 Mobile-IPv6 based WLAN-GPRS Testbed Implementation

In this section, we will describe the initiatives we have taken for evaluating Mobile IPv6 based vertical and horizontal handoff schemes and other performance issues in heterogeneous spaces. As part of the Cambridge Open Mobile Systems (COMS) project [19], we have implemented a Mobile-IPv6 based GPRS-WLAN testbed (see figure 1).

The cellular GPRS network infrastructure currently in use with the testbed is a Vodafone UK's production GPRS network [7]. The WLAN access points (APs) are IEEE 802.11b APs located at different locations of the William Gates Building housing both University of Cambridge Computer Laboratory and Laboratory for Communication Engineering. The APs provide full in-building coverage to WLAN users. This setup gives us the opportunity to evaluate indoors spaces with nomadic users moving at low speeds. Additionally, access points with extended range antennas have been provisioned, to give an outdoor coverage for spatially co-located zones (e.g. an average extended coverage of 500 meters).

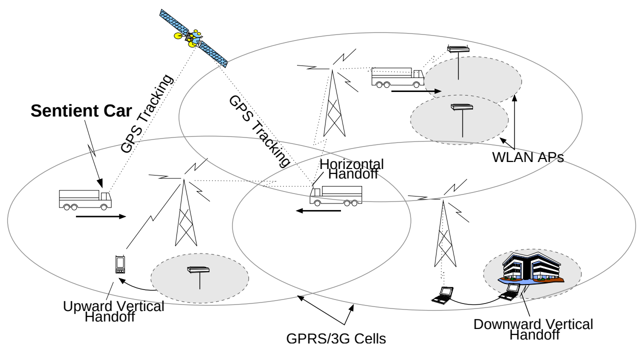

Figure 2: **The testbed recreates a highly mobile outdoor networking scenario.**

Further, the GPRS infrastructure offers base stations (BSs) that are directly linked to the SGSN (Serving GPRS Support Node) which is then connected to a GGSN (Gateway GPRS Support node). In the current Vodafone configuration, both SGSN and GGSN nodes are co-located in a single CGSN (Combined GPRS Support Node) [24]. A well provisioned virtual private network (VPN) connects the Lab network to that of the Vodafone's backbone via an IPSec tunnel over the public Internet. A separate "operator-type" RADIUS server is provisioned to authenticate GPRS mobile users/terminals and also assign IP addresses.

For user access to the wireless testbed, the mobile node (e.g. laptops, PDAs) connects to the local WLAN network, the node is also simultaneously connected to a GPRS phone/card using a bluetooth wireless link (or a serial point-to-point link). The mobile node is Mobile-IPv6 (MIPv6) enabled, which means we also have a MIPv6 home agent in our network (see figure 1) to serve all our mobile nodes. For thorough testing purposes, we have the latest GPRS-WLAN phones/cards from a number of manufacturers.

The GPRS-WLAN testbed, in addition to a number of nomadic clients, also has a Sentient Car that recreates a highly mobile scenario for outdoors (see figure 2). The Sentient Car is *context-aware*; it has a priori information of its location using a GPS (Global Positioning System) tracking system and its velocity using speed sensors. This information is then available to a computer terminal, which is also fully equipped with GSM/GPRS and WLAN connections. Using the Sentient Car (see figure 3), we are able to conduct many novel and complex experiments for understanding performance-critical issues related to mobility. Further, this enables us to uncover how well protocols and applications can scale over such heterogeneous environments.

3 A Client-Assisted Handoff Scheme

In the previous section we described the testbed that we will use for evaluating mobility-related performance issues in heterogeneous environments. In this section, we discuss the adaptation of a novel client-assisted handoff scheme that is used to improve handoff mechanism between IEEE 802.11b base stations [6] which can be suitably adapted for hybrid scenarios.

The scheme exposed in [6] allows a mobile client to control the handoff process, decide the best link based on signal strength, and to enable timely handoffs between different point of attachments for mobile hosts moving in MIPv6-enabled WLANs. The solution has been designed for the client-side, thus allowing the mobile host the freedom to make its own informed decision on the link to join. This is especially advantageous for a hetereogeneous environments where the network infrastructure is usually proprietary in nature.

Furthermore, for the network infrastructure to assist mobile hosts in heterogeneous handoffs, the cost factor, complexity in signalling, and network policies may not be favourable for an "all-IP" network paradigm. Some examples of network-assisted handoff improvements are shown in [23] and [15]. These proposals are for homogeneous environments, unfortunately, they cannot deal

Figure 3: **The Sentient Car for Context-Aware Networking.**

with the complexity posed by hybrid wireless scenarios.

We have implemented and tested the client-assisted handoff algorithm for vertical handovers after some modifications [5]. For horizontal environments, the handoff mechanism utilises link-layer beacons from nearby IEEE 802.11b wireless access points to detect available point of attachments. However, a number of factors would affect the handoff decision in the mobile client. Firstly, router advertisements (RAs) –part of the IPv6 neighbour discovery protocol– from nearby routers are trapped as soon as they arrive to the network stack, and cached in the mobile host to determine the best link (based on signal strength). Secondly, the mechanism also considers the TCP connection state to provide further intelligence to the mobile host during the handoff decision. This is rather important when the mobile host is located in a boundary of two overlapping cells within the same network.

However, for the mechanism to be used in a heterogeneous network environment, important considerations have to be made. Firstly, multiple network interfaces have to be taken into account. Furthermore, each interface has its own set of characteristics which will influence the handoff decision. Switching between different network technologies requires additional information to provide the mobile host with even more intelligence during the handoff procedure. Secondly, caching router advertisements works well when network cell boundaries are typically overlapped. This is particularly true for *upward* vertical handoff case (i.e. WLAN→GPRS). In an overlay model, RA caching has to occur for the network above the current one. For example, in the testbed presented, RAs from (virtually pervasive) GPRS can be cached *a priori*, avoiding the need to wait – or look up – during handoff (e.g. WLAN→GPRS) leading to complete elimination of the detection time in the new network (i.e. GPRS) – this implies an improvement in the overall handoff latency.

Additionally, TCP connection state can also be taken into account to execute or delay the handoff. For example, the mobile host can wait for more time in WLAN before handing off to GPRS if a *specific* flow is ongoing (because of security or cost reasons). Hence, in an overlay model TCP connection state is an important input for adaptive mobility.

In the implementation, the handoff module is hooked to the neighbour discovery module (ndisc.o) of the IPv6 stack [3]. RAs from overlayed networks are first received by the network stack, which are then passed to the handoff module. The handoff module checks if the RA corresponds to an upper layer network (i.e. GRPS) and caches it. Also, the module checks for RAs that have expired and makes the update using the most-recent RA.

Table 1: **Mean Detection Time with and without RA caching**

Scheme Used	Mean Detection Time
MIPv6 (Fixed RA: 300-400ms)	551.33ms
MIPv6 (with Handoff Module)	1.21ms

We have evaluated the performance of our handoff module for RA caching. In this test, we allowed RAs from GPRS to be cached and we force handoffs (i.e. WLAN-GPRS) periodically, by setting a timer in the handoff module, while simultaneously dowloading a file from a http server – representing the correspondent node. Table 1 shows the mean detection time in an inter-network handoff with and without the module. We can observe that using the module leads to almost complete elimination of the detection time. The reduction from 550 ms to about 1 ms on average is a substantial improvement in mean detection time, and hence, overall handoff latency. Note, however, this benefits only the handoff latency during *upward* vertical handoffs.

The introduction of multiple network interfaces to support mobility across heterogeneous networks can also have implications with power consumption in mobile devices. Other approaches may propose to avoid using multiple interfaces simultaneously. However, this is certainly a disadvantage since an important function of our client-assisted handoff mechanism – the router advertisement cache – is deemed useless. Nevertheless, our handoff mechanism can still minimise the effect of such disruptions, especially on TCP connections where TCP's retransmissions can be particularly distasteful to the overall user experience.

4 Work in Progress

The testbed we have implemented will be used to conduct experiments for evaluating schemes that improve vertical handoffs, transport layer enhancements, and schemes to aid MIPv6 over a fully-integrated GPRS-WLAN testbed for a number of mobility-related scenarios. We are currently evaluating the following issues:

- **Effective Handoff Control.** A mobile node may use different available schemes to detect when its link-level point of attachment has moved from one IP subnet to another. The client-based handoff scheme that we have implemented works well, and can be easily extended for overlay networks. However, there are issues that determine the effectiveness of a handoff scheme: a handoff process control, decision for the best link, and finally the handoff execution. Problems with handoff control are somewhat more complicated in hybrid environments. Since the control process has to monitor multiple interfaces, and also use them simultaneously, it has to decide the best instant of when it can safely execute the handoff. Further, the decision process is compounded since handoff execution is simply not restricted to simple policies, but to a complex set of rules. For example, a handoff agent in a mobile client can also decide based on the link load, QoS support, security, user preferences, or rules derived from some other physical context variables. We believe a scheme which takes all this into account will obviously optimise user experience in heterogeneous spaces.
- **Handoff decision rules.** Most of the current horizontal handoff schemes operate based on certain link-layer information, mainly signal strength. However, this is not by itself sufficient to assist the handoff process in hybrid environments. Although signal quality has to be interpreted in an overlay, future multimode nodes will receive signal strength information from multiple interfaces, and so handoffs need not to be based on this exclusively. The decision can depend on some other information, for example, a handoff agent in a highly mobile multimode device can implement a rule-base, which underpins its decision based on, for instance, the overall link characteristics, security, link cost, and robustness.
- **Context-Awareness and Mobility.** Operating conditions change as mobile clients move. This is particularly true for highly mobile hosts that not only keep changing their location, but also their direction and velocity. The Sentient Car [16] (figure

3) provides an excellent example. Context conditions in mobile devices can have a direct bearing on the flows' end-to-end performance, the handoff schemes implemented and to the overall user connectivity. A mobile device context involves aspects such as physical context variables (e.g. device location, movement direction, velocity etc.), application characteristics and, of course, user-based preferences. Hence, any sophisticated handoff mechanism meant for hybrid environments must contemplate using *context-awareness* in future implementations. For example, based on the exact position (e.g. available from a GPS system) and velocity information available to a mobile host (e.g. speed sensors), a co-located proxy in the infrastructure can assist mobility by tracking and accurately predicting when a handoff should occur. Once it is equipped with this information, it can communicate this to the correspondent nodes, which in turn can assist flow adaptation even before a handoff occurs.

5 Related Work

Related research on vertical handovers conducted some years ago, as part of the BARWAN project [4], builds on a 4-network wireless overlay network. Handoffs make use of a MIPv4 solution, using a multicast address in the mobile host (the care-of-address) to receive advertisements from potential access points in the overlay. However, scalability is a fundamental constraint for this solution; managing multicast address is a complex task, and in an environment with hundreds of mobile hosts this can become a huge limitation. Furthermore, handoff latency is bounded by the time needed for the mobile host to discover that it has moved in or out of a network in a wireless overlay. In contrast, our client-based handoff scheme improves latency by controlling and forcing the handoff, using a *router advertisement cache*.

In the MosquitoNet project [14], two flow-based handoff mechanisms are built typically as extensions to Mobile IPv4. The first mechanism supports multiple packet delivery methods and adaptively selects the most appropiate one to use, typically relying on the characteristics of each flow. The other, however, enables mobile hosts to make use of multiple network interfaces simultaneously, and controls the selection of the most appropriate network interfaces on ongoing flows. Both approaches are based on traffic flow details that construct the selection policies. However, an interesting extension to these approaches can make use of other information such as device context, layer-2 triggers, and the TCP connection information, rather than simply relying on data flow characteristics.

Mobility is currently a hot topic in the IETF. They have three main protocols to manage mobility: MIPv6 [13], Hierarchical MIPv6 [23], and Fast Handover Protocol [15]. The primary driver of these proposals, as well as ours, is to minimise the effects of handoffs in mobile networking.

MIPv6 [12] can be particularly disadvantageous during high mobility scenarios. An approach here can make use of a scheme similar to those used in micro-mobility protocols that reduce the handoff latency and improve performance under high mobility scenario. Micro-mobility protocols are broadly aimed at improving transparent mobility at the subnet level of a network domain. A. T. Campbell et al. [8] gives a nice survey of micro-mobility protocols. *Seamoby* [9], resolves complex interaction of parameters and protocols needed for seamless handoffs. The two main issues being currently dealt in Seamoby are the dormant mode host alerting problem (i.e. paging) and context transfers between nodes in an IP access network (i.e. handoff).

6 Conclusions

We reviewed networking needs for future mobile environments. Especially, the need for a seamless handoff mechanism to efficiently roam in heterogeneous environments was highlighted. To that end, we consider our testbed implementation as an important step, to develop further mobility related research, and understand its effects on existing applications and protocol stacks. We envisage Mobile IPv6 as our base for the design and implementation of novel handoff algorithms over *all-IP* – wired and wireless networks. A Mobile IPv6 client-based handoff mechanism has been implemented, which improves performance for horizontal handoffs in IEEE 802.11b based subnets. Currently, similar schemes are being developed to suitably adapt to vertical handoffs.

Some other critical aspects should also be accounted for while designing next-generation inter-system and inter-network handoff algorithms. Issues such as multiple network interface management, link adaptation, physical context-effects, and impact of vertical handoffs on transport layer and applications need to be addressed. Our work-in-progress points towards these issues.

Finally, we believe multimode mobile terminals augmented by novel handoff mechanisms can improve performance by offering seamless roaming in hybrid environments. Designers, however, need to be very careful in implementing such handoff schemes, as these mechanisms might involve certain judicious trade-offs. More importantly, mechanisms here need to exhibit low latency, incur little or no data loss, scale typically to large networks and user base, adapt to different environments, and act as a convergence point between heterogeneous spaces and technologies. Our on-going work explores at least one such handoff scheme.

Acknowledgements

We wish to thank the system administrators of the Computer Laboratory and the Department of Engineering for making the testbed possible, most notably Michael Gray, Paul Whitehouse, Caroline Blackmun, Piete Brooks, Martyn Johnson and Brian Jones.

Further gratitude goes to Andy Hopper, Frank Stajano, Glenford Mapp, Ian Pratt, Jon Crowcroft and Calicrates Policroniades for providing the support and necessary advice for this project. The project is supported by Vodafone Group R&D, Sun Microsystems Inc., BenchMark Capital, and the Engineering & Physical Science Research Council (EPSRC) of the UK, and ATT Labs Cambridge.

Pablo Vidales has a scholarship from the Mexican government through the National Council of Science and Technology (CONACYT). The Sentient Car project was funded by the Cambridge-MIT Institute (CMI). We would like to thank all of these organisations for making this work possible.

References

[1] BT exact Technologies, “White paper: Technology Timeline”, Reference:42270, Issue 1, November, 2001.

[2] M. Weiser, “The Computer of the 21st Century”,*Scientific American*, September 1991.

[3] T. Narten, E. Nordmark and W. Simpson, “Neighbor Discovery for IP Version 6 (IPv6)”, *Request for Comments*, RFC 2461, 1998.

[4] E. A. Brewer, et al., “A network Architecture for Heterogeneous Mobile Computing”, *IEEE Personal Communications Magazine*, Oct 1998.

[5] Rajiv Chakravorty, Pablo Vidales, Leo Patanapongpibul, Kavitha Subramanian, Ian Pratt, Jon Crowcroft, “On Internetwork Handover Performance using Mobile IPv6”, University of Cambridge, Technical Report. `http://www.cl.cam.ac.uk/users/rc277/overlay.html`

[6] L. Patanapongpibul, “A Client-based Handoff Mechanism for Mobile IPv6 Wireless Networks”, *8th IEEE Symposium on Computers and Communications (ISCC)*, IEEE Computer Society Press, June 2003.

[7] “An Introduction to the Vodafone GPRS Environment and Supported Services”, Issue 1.1/1200, December 2000, Vodafone Ltd., 2000.

[8] A. T. Campbell and J. Gomez-Castellanos, “IP Micro-Mobility Protocols”, *ACM Mobile Computing & Communications Review (MC2R)*, Vol. 4, No. 4, pp 45-54, October, 2001

[9] J. Kempf, “Problem Description: Reasons For Doing Context Transfers Between Nodes in an IP Access Network”, *IETF Request for comments*, RFC 3374, September 2002.

[10] J. Mitola, “The software radio architecture”, *IEEE Communications Magazine*, May 1995

[11] C. E. Perkins and D. B. Johnson, “Mobility Support in IPv6”, In *Proceedings of ACM MOBICOM*, 1996.

[12] C. E. Perkins, “Mobile IP”, *IEEE Communications Magazine*, May 1997.

[13] D. B. Johnson, C. E. Perkins and J. Arkko, “Mobility Support in IPv6” (draft-ieft-mobileip-ipv6-22.txt), May 2003 `http://www.ieft.org`

[14] X. Zhao, C. Castellucia and M. Baker, “Flexible Network Support for Mobile Hosts”, In *Proceedings of ACM/IEEE MobiCom 98*, Dallas, TX

[15] R. Koodli, “Fast Handovers for Mobile IPv6”, (draft-ietf-mobileip-fast-mipv6-05.txt). `http://www.ietf.org`

[16] P. Vidales and F. Stajano. “The Sentient Car: Context-Aware Automotive Telematics” in *Proceedings of the First European Workshop on Location Based Services (LBS-2002)*, London, 2002.

[17] “Mobile IP for Linux (MIPL) Implementation by HUT Telecommunications and Multimedia Lab”, `http://www.mipl.mediapoli.com`.

[18] “6BONE: Testbed for the deployment of IPv6”. `http://www.6bone.net`

[19] Cambridge Open Mobile Systems (COMS) Project Initiative. `http://www.cl.cam.ac.uk/coms/`

[20] IEEE 802.20 Initiative. `http://grouper.ieee.org/groups/802/20/`

[21] R. Chakravorty, A. Clark and I. Pratt, “GPRSWeb: Optimizing the Web for GPRS Links”, In Proceedings of *ACM/USENIX First International Conference on Mobile Systems, Applications and Services (ACM/USENIX MOBISYS 2003)*, May 2003, San Francisco, USA (to appear). `http://www.cl.cam.ac.uk/users/rc277/gprs.html`

[22] R. Chakravorty, S. Katti, J. Crowcroft and I. Pratt, “Flow Aggregation for Enhanced TCP over Wide-Area Wireless”, In *Proceedings of the IEEE INFOCOM 2003*, March-April, San Fransisco, USA (to appear). `http://www.cl.cam.ac.uk/users/rc277/gprs.html`

[23] H. Soliman, et. al., “Hierachical MIPv6 Mobility Management”, IETF draft, draft-ietf-mobileip-hmipv6-05.txt, July 2001.