Scalable QoS-Aware Mobility for Future Mobile Operators

Rui L. Aguiar and Susana Sargento, Instituto de Telecomunicações

Albert Banchs, Carlos J. Bernardos, María Calderón, and Ignacio Soto, Universidad Carlos III de Madrid

Marco Liebsch and Telemaco Melia, NEC Network Laboratories

Piotr Pacyna, AGH University of Science and Technology

ABSTRACT

Telecom operators and Internet service providers are heading for a new shift in communications paradigms. The forthcoming convergence of cellular and wireless data networks is often manifested in an "all IP approach" in which all communications are based on an end-to-end IP protocol framework. The approach to network design becomes user and service-centered, so that continuous reachability of mobile users and sustained communication capabilities are default requirements for a prospective architecture. In this article, we describe a network architecture which is able to provide seamless communication mobility, triggered either by the user or by the network, across multiple technologies. The architecture allows for media independent handovers and supports optimized mobility and resource management functions. The main focus of the article is on major technical highlights of mobility and quality-of-service (QoS) management subsystems for converged networks.
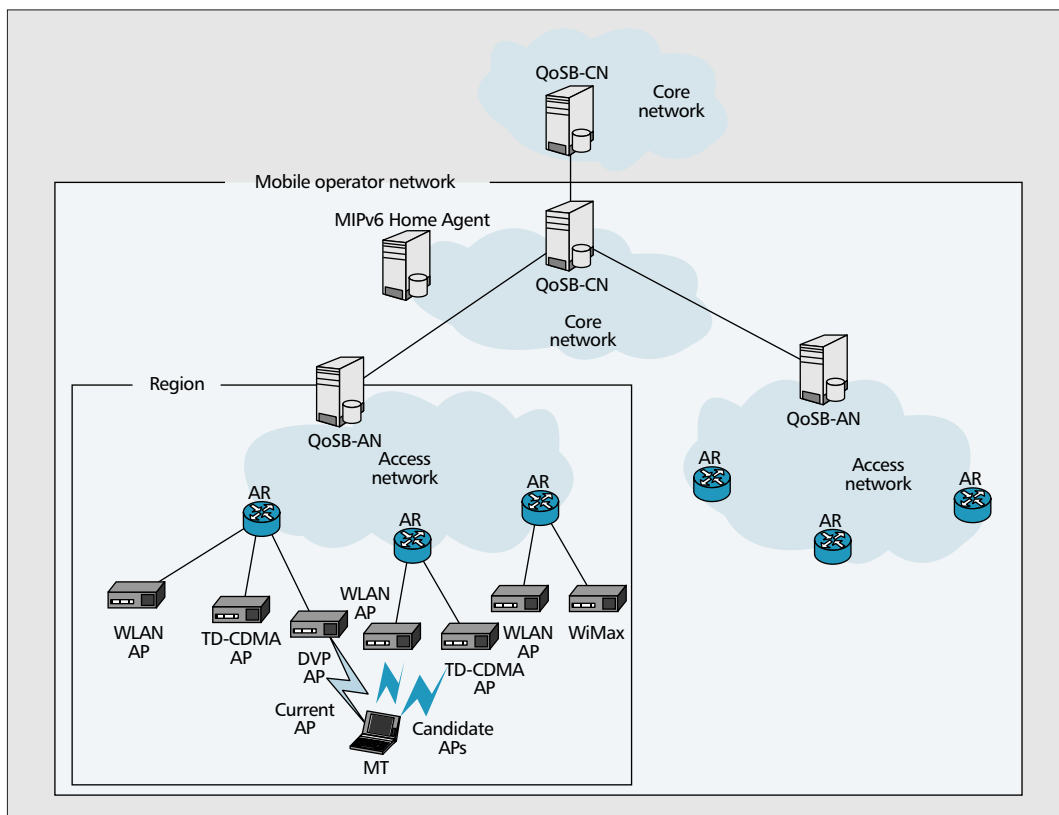
INTRODUCTION

Several indicators point towards the coexistence of heterogeneous networks in the future. These relate to multiple types of access technologies spanning the whole range from home networks and wireless LANs to campuswide wireless access, and from 2G/3G cellular networks to satellite networks. Both operators and manufacturers have taken up the development and introduction of dual-mode and multimode handsets to permit connectivity across 3G and WLAN-based networks. Currently discussed standards, such as Unlicensed Mobile Access and Voice Call Continuity, strive to address this integration, but still fall short of providing universal solutions that support seamless integration of these networks. Users on the move will experience service discontinuities, so such standardized solutions can only be regarded as a first step.

Tomorrow's customers will expect the network, and in particular its technological structure, to "disappear" and be of no concern. Along these lines, previous works [1, 2] have proposed solutions that support seamless mobility based on the Internet Protocol version 6 (IPv6). While these works have shown that the basic concepts are viable, the Daidalos project [3] has moved to a comprehensive approach to provide seamless end-to-end services accessible anytime and anywhere across heterogeneous technologies.

Daidalos architecture proposes an enhanced IPv6 mobility platform that provides mobility and QoS, as key drivers of future all-IP based 4G networks. Fast intra- and intertechnology handovers are a solution to the requirement of seamlessness. While IETF protocols have been proved to work properly in a "standalone" manner, their applicability in integrated environments requires further enhancements. For next-generation integrated systems, additional requirements are the optimization of resource usage, scalability for an increasing number of customers, and increased network flexibility. Previous studies and prototypes [4–6] have already demonstrated the feasibility of parts of these integrated IP-based solutions.

In this article we show how the Daidalos architecture has improved, enhanced, and optimized existing micro-mobility schemes to work in an integrated way with the macro-mobility management scheme of the de facto standard Mobile IP version 6 [7]. The main highlights of the architecture include support for both network-initiated and mobile-initiated handovers in heterogeneous environments integrated with QoS resource management. The next section presents an overview of our scalable architecture. We then describe how it operates during the handover. We describe complexity issues related with our implementation, and then scalability is addressed. The article's conclusions are presented in the final section.

■ **Figure 1.** *Hierarchical structure of mobility and QoS architecture.*

DAIDALOS MOBILITY ARCHITECTURE OVERVIEW

The Daidalos mobility architecture aims to provide an efficient and scalable integration of multiple network technologies, with sustained QoS support. The (simplified) general view of the architecture is illustrated in Fig. 1. As it can be observed from the figure, the architectural design follows a hierarchical structure: the network of each mobile operator consists of a core network (two such networks, from different operators, are represented in the figure) and a set of access networks. The access networks contain multiple access routers (ARs), with multiple radio access points (APs) each. The architecture supports multiple access technologies, including WLAN, WiMAX, TD-CDMA, and DVB.

Each access network is called a region. Resources in each region are independently managed by an access network QoS broker (QoSB-AN), thus providing a first scalability step. Resources in the core are managed by the core network QoS broker (QoSB-CN), which communicates for end-to-end QoS with the QoSB-ANs of the mobile operator's network as well as with the QoSB-CNs of the other operator's networks. The architecture is based on widely accepted standards for mobility and QoS. Mobility is implemented by means of the MIPv6 protocol [7], with fast handover extensions [8], and QoS is based on the DiffServ architecture [9]. However, additional mechanisms that integrate and complement MIPv6 and DiffServ are needed in order to achieve the objective of providing QoS to mobile users while optimizing the overall performance. Such mechanisms have been designed in the architecture.

Handover decisions in our architecture are sustained both by measurements on signal quality as well as QoS measures (such as load and resource availability). Handovers can be started either by the terminal or by the network. We refer to the former as a mobile-initiated handover (MIHO) and to the latter as a network-initiated handover (NIHO). Handover execution is improved with functions for maintaining quality during handovers, along with tight coupling with QoS functions.

In the following we describe the enhanced functionalities of the architecture in further detail. The modules needed to instantiate the functionalities are illustrated in Fig. 2, organized according to their physical location:

Enhanced mobile-initiated handover decisions. Handover decisions in the case of a MIHO are enhanced with the objective of ensuring that, from all the possible AP candidates, the "best" one is chosen. The module responsible for the handover decision at the MT is intelligent interface selection (IIS). This module relies on the mobile terminal controller (MTC) to obtain the information it uses to take a decision. This includes signal quality measurements, obtained from the mobility abstraction layer (MAL), as well as QoS measures, such as the load of the APs, retrieved from the candidate APs. The latter information is obtained from QoS abstraction layer (QAL) in the neighboring ARs, and conveyed by means of the Candidate Access Router Discovery (CARD) protocol [10] to the
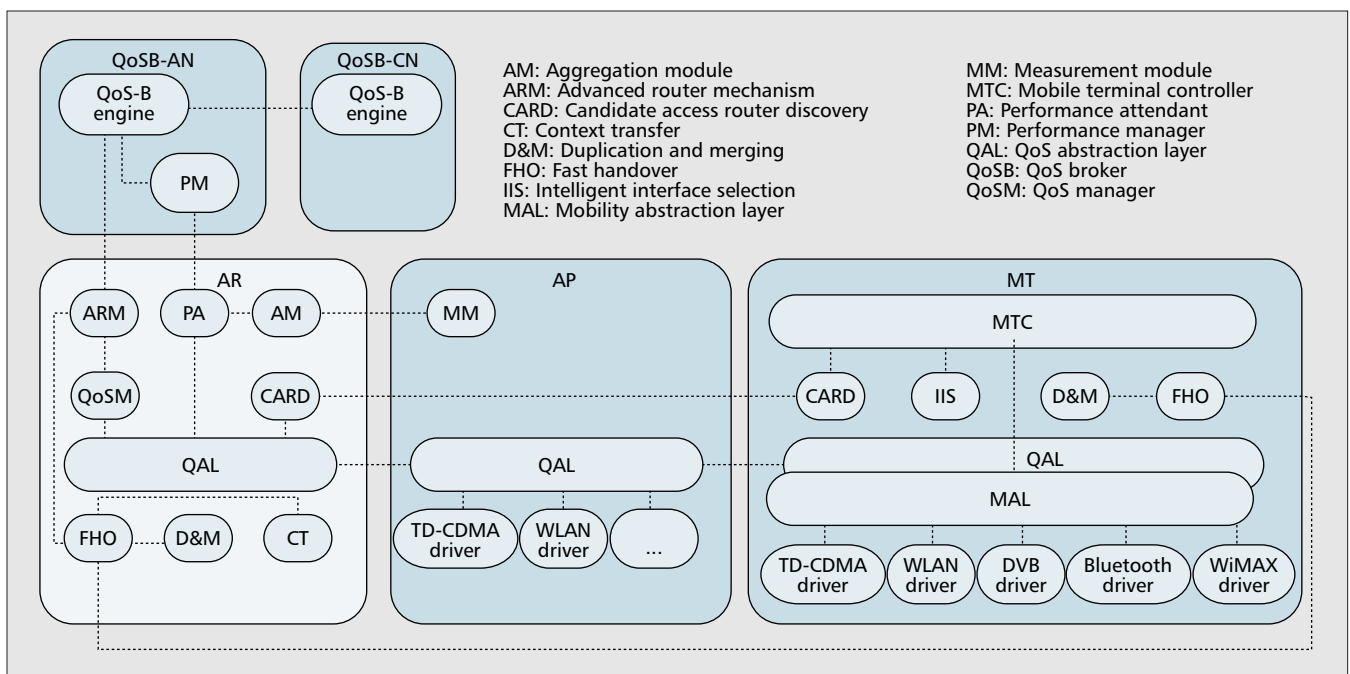
■ Figure 2. *Architecture modules.*

MT. With this information, the target AP for the handover is chosen so that both signal strength and QoS requirements are met in the new AP, thus guaranteeing appropriate operation and service quality after the handover.

Network-initiated handover functionality. The enhanced MIHO functionality ensures that handover decisions are taken optimally according to local information, but does not guarantee that the overall distribution of resources will be optimal from an operator perspective — which is essential for a realistic network. In order to achieve this, NIHO support is required in order to allow the optimization of the overall capacity by properly balancing the load among the various APs of a region [11]. For this purpose, the performance management (PM) module at the QoSB-AN collects information about the load of the different APs and the radio link quality between the MT and its candidate APs, and based on this information (eventually) reorganizes the wireless connections. Information on the load of the APs is obtained by the performance attendant (PA) modules at the APs, from their interface with the QAL, and delivered to the PM. Signal strength measurements are taken by the measurement modules, filtered out and aggregated by the aggregation module (AM), provided to the PA at the AR, and from there conveyed to the PM [12]. Based on all these data, the PM then reorganizes the connections of all MTs for achieving optimized global performance. This reorganization takes into consideration QoS beyond the wireless access, by means of the interaction between the PM and the QoSB engine at the QoSB-AN. The NIHO execution is then triggered by the communication between the QoSB and the fast handover (FHO) execution module at the AR, through the advanced router mechanism (ARM).
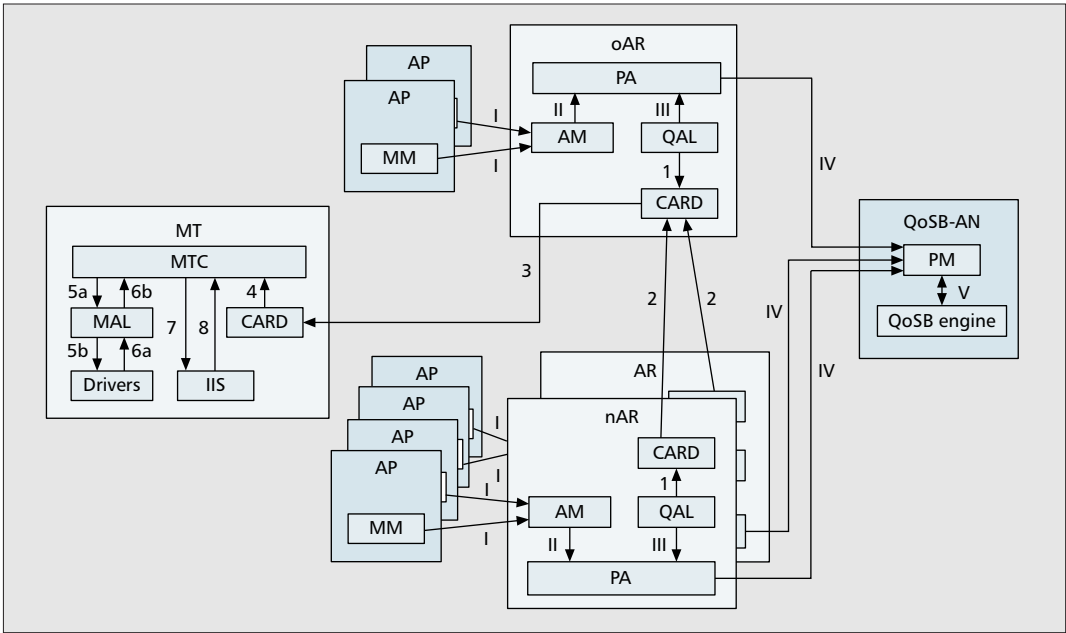
Seamless handover execution. In the execu-tion of a handover involving the old AR (oAR) and the new AR (nAR), it is required that conti-nuity of communication be maintained. To per-form a low-latency lossless handover, the fast handovers for Mobile IPv6 protocol [8] is enhanced with duplication and merging (D&M) functions. These functions improve performance by duplicating the packets addressed to the MT at the old AR to avoid packet loss. In order to set up the MT's context in the nAR, the context transfer (CT) function is used to transfer the mobility-related state (including security infor-mation).

Quality of service. QoS is based on the Diff-Serv architecture. Admission control and resource reservations are handled by the QoSBs, which act jointly to perform QoS reservations over an end-to-end path. QoS reservations at the routers are performed through the interaction between the QoSB Engine at the QoSB-AN and the ARM module at the AR, which performs the reservation via the QoS manager (QoSM). Simi-larly, reservations in the wireless access part are performed through the interaction between the QoSB engine at the QoSB-AN and the ARM module at the corresponding AR. The latter communicates with the QoSM, which communi-cates with the QAL at the AR. QoS reservations in the wireless access are then performed by the QAL modules at the AP and MT.

Multiple technology support. The support of multiple technologies in the architecture is pro-vided by means of a modular design based on the use of abstraction layers (ALs): the mobility abstraction layer (MAL) and the QoS abstrac-tion layer (QAL). These ALs interface with drivers of the different technologies and offer a unique interface to the upper-layer modules of the architecture, while hiding the specifics of the underlying technologies. The QAL offers a tech-nology-independent interface for QoS functions

■ **Figure 3.** *NIHO and MIHO decision operation.*

such as the setup of a QoS connection or the measurement of available resources in an AP. Similarly, the MAL offers a technology-independent interface for mobility-related functions such as the execution of a handover or measurement of signal strength received at MT.

In the following sections, the above functions and their interactions are described in detail. First, we describe the operations related to the decision of performing a handover for the mobile and network initiated cases (MIHO and NIHO). Then, we address the process of handover execution, which is almost identical for both cases.

HANDOVER DECISION FUNCTION

In existing IP-based architectures, handovers are typically initiated by terminals upon detecting that the quality of the signal received from the AP degrades below a certain threshold. In our architecture, handover decisions take into account both the signal strength and QoS measures, and can be triggered either by the terminal (MIHO) or the network (NIHO). In this section we describe the decision process in the handover operation with MIHO and NIHO. The protocol operation for both cases is illustrated in Fig. 3.

MOBILE-INITIATED HANDOVER

The main decision to be made with regard to a MIHO is the choice of the new AP to handover. In the preparatory phase of this process, the MT discovers the available candidate APs. This is done by means of the CARD protocol [10]. With this protocol, neighboring ARs exchange QoS information about attached APs and provide it to the MT. Specifically, QoS-related information is provided to CARD by the QAL module (message 1 in Fig. 3). This information is sent from candidate ARs (including the future nAR) to the current AR (oAR), and from there it is conveyed to the MT (messages 2 and 3, respective-

ly). Once the MT has obtained information about available candidate APs and their QoS, it proceeds to measure their signal strength. The decision of which are the APs whose signal strength has to be measured is made by the MTC based on the CARD information (message 4) and is provided to the drivers of the respective technologies via the MAL (messages 5a and 5b). The measured signal qualities are reported back to the MTC (messages 6a and 6b) and this information, together with the QoS-related information obtained previously, is then provided to the IIS (message 7). Based on these data, the IIS decides which is the most appropriate AP for handover. This decision also takes into account the user preferences, which are provided to IIS by the user via a graphical interface and include preferred technologies and providers. Note, however, that strictly the IIS decision may result in a failed handover, since it does not account for the availability of resources in the end-to-end path. In order to avoid a failed handover, the IIS does not choose a single AP candidate, but a prioritized list of up to three possible candidates (all of them with enough resources in the wireless access) ordered according to signal strength and preferences. This list is provided to the MTC (message 8) and at this point the handover execution process starts (see the following section).

NETWORK-INITIATED HANDOVER

A NIHO may be triggered either by the degradation of the quality of the signal received from the MT by the AP[1] or by the QAL upon detecting that the load of an AP exceeds a predefined threshold. NIHO decisions are taken with the goal of optimizing global performance in a region controlled by a QoSB-AN. These decisions are taken by the PM module in conjunction with the QoSB engine module. To this end, available QoS at APs as well as data related to signal strength are sent to the PM module,
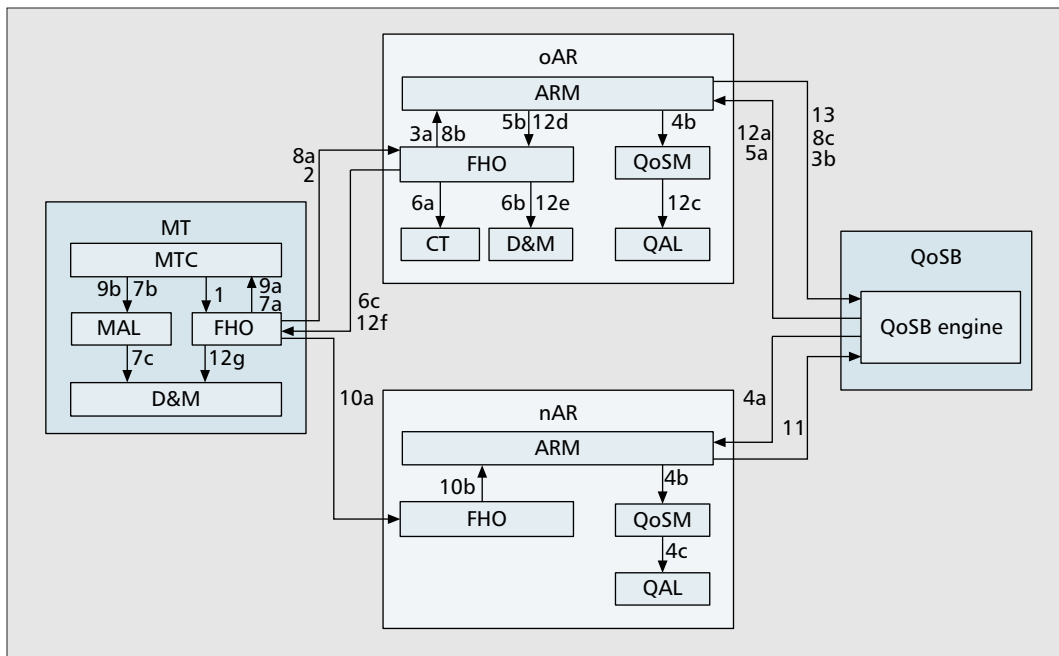
■ **Figure 4.** *FHO execution operation.*

whereas the QoSB engine accounts for end-to-end QoS information. The process of transferring information to the PM is described in the following.

Signal strength measurements are taken by MM modules at the APs and from there they are transferred to the AM modules at the ARs (message I in Fig. 3), which aggregate all the information received and provide it to the PA (message II). By measuring the strength of the signal received from a given MT at all the APs of a region, it is possible to estimate all the candidate APs that provide good signal quality to this MT [11]. In addition to signal strength data, the PA also collects QoS-related information from the QAL modules (message III). The data of all the PAs of a region are sent to their corresponding PM module, which is located at the QoSB-AN controlling the region (message IV).

With all the above information, the PM is aware (through the QoS-related data) of the load of the various APs of the region, and is also aware (from the measurements taken) of the possible candidate APs that each MT may be handed over to while preserving a good signal quality. Based on these data, the PM can then decide the AP to which each MT should be attached to such that i) load is optimally distributed among all the APs of a region, and ii) the signal strength of all connections is good. These decisions are checked against the QoSB engine in order to make sure that end-to-end QoS requirements are available for the connections (message V). At this point the handover execution process starts (see next section).

One of the crucial aspects of the operation described above for NIHO regards the triggers for this process. If NIHOs are triggered too often, performance may be severely degraded due to too many (consecutive and simultaneous) handovers. On the other hand, if infrequently triggered, the network performance will not be optimal and the NIHO mechanism will be under-exploited. To find a good trade-off between these two extremes, in our architecture, NIHO triggers are event-based with event filtering and smoothing functions on the measurement and aggregation modules.

Having two different triggers to initiate the handover, one in the network and one in the user terminal, it is essential to avoid the bouncing effect in which MT handovers from an AP to another because of a network trigger, and immediately handovers back because of the MT trigger. The proposed architecture prevents this effect, as a network trigger will always handover a terminal to an AP where signal level and QoS is good enough so that the MT will not have the need to trigger a handover again.

HANDOVER EXECUTION FUNCTION

In this section we address and describe the fast handover execution operation. This is a time-critical procedure, as sessions need to be transferred from the old to the new network without noticeable disruptions. Figure 4 presents an overview of our fast handover operation. The protocol operations affect the MT, the old and new ARs (oAR and nAR) and the QoSB-AN (hereafter simply QoSB). In the following, we first explain the handover execution for the MIHO case, and afterwards we describe the differences between this case and the NIHO one.

MOBILE-INITIATED HANDOVER

The handover in a MIHO is performed as follows. The MTC, upon receiving information from the IIS to trigger the handover procedure, activates the FHO module in the MT (message 1) to initiate a handover. It then sends a message to its current AR containing a list, ordered by preference, with up to three possible candidate APs and the respective ARs (message 2).
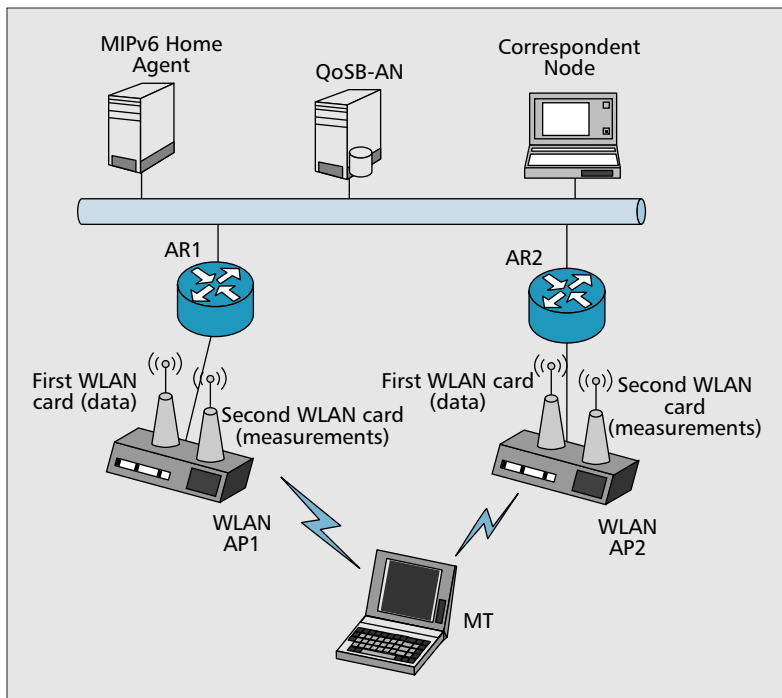
■ Figure 5. *The testbed.*

The AR, through the ARM module, forwards the request for approval to the QoSB (messages 3a and 3b). Upon receiving this request, the QoSB verifies the availability of the required end-to-end QoS and answers (if possible) with the first occurrence of the list that can be allocated in the selected AR. Then, it informs the nAR of the QoS requirements (message 4a), which triggers the QoS reservations (messages 4b and 4c), and sends back to the oAR its handover decision (messages 5a and 5b). Upon receiving the handover decision, the oAR triggers the CT, instructs the D&M to start duplication, and informs the MT that it can now move to the chosen AP (messages 6a, 6b, and 6c). When duplication is activated, the oAR forwards any data directed to the MT to the nAR for its delivery to the MT, in order to minimize data loss during the handover process. The context information to be transferred includes security-related information (note that QoS-related context does not need to be transferred by CT, since it was previously transferred directly through the QoSB).

As soon as the MT receives message 6c, it starts the merging process (messages 7a, 7b, and 7c), whose function is to filter out the duplicated packets received from the oAR and the nAR and deliver only one copy of each packet to the applications. Before performing the handover, the MT informs the oAR of its decision (message 8a), which in its turn informs the QoSB (messages 8b and 8c). Then, the FHO module in the MT notifies the MTC of the decision to perform disconnection from the previous interface (message 9a), and the MTC notifies the MAL about the target AP the terminal should attach to (message 9b). This is followed by disconnection from the current link and attachment to the new one. Upon connection to the nAR interface, the MT informs the nAR, which updates the

IPv6 neighbor cache and forwards packets to the terminal (messages 10a and 10b). The nAR informs the QoSB that the MT is attached to the new link (message 11), and this indication is then forwarded to the oAR (message 12a) in order to delete the reservations (messages 12b and 12c), and stop D&M at the oAR (messages 12d and 12e) and at the MT (messages 12f and 12g). After this process, the oAR informs the QoSB (message 13) that the reservation release actions have been successfully performed.

NETWORK-INITIATED HANDOVER

When the network initiates a NIHO, the procedure is similar to the one described above. However, now the handover is not triggered by the MT (message 2) and message 6c contains a flag, indicating that this is a NIHO. The remaining process is the same.

IMPLEMENTATION EXPERIENCES

In order to validate the proposed architecture in terms of protocol design and basic functionality, we implemented the proposed architecture and evaluated its performance in the testbed illustrated in Fig. 5. Our implementation was done using the Linux distribution Mandrake 10.0 with kernel 2.6.8.1 and MIPL basic Mobile IPv6 support. The description of the testbed and results hereafter are based on the wireless LAN (WLAN) technology, although implementation is also available for TD-CDMA with some hardware limitations.

In order to provide the mobility features required for MIHO and NIHO, we introduced modifications to the behavior of the Host AP WLAN driver in MTs and APs. The first modification we implemented was to disable automatic handovers and use a function to force handover execution when required. This allows handover decisions to be taken by the network (NIHO), or by the terminal accounting for additional factors other than signal strength (MIHO). For MIHO, we also implemented a function for scanning a selected subset of channels, which allows the signal strength measurements to be limited to only the channels of the neighboring APs (provided by CARD), in contrast to the default automatic scanning, which covers all channels and therefore results in larger latencies.

In NIHOs, APs are required to measure the signal strengths of the MTs connected (using different channels) to other APs. In order to make these measurements, we installed a second WLAN card at the APs whose function was to scan all channels periodically and perform passive measurements on the signal strength detected from the MTs. QoS functions, required both for MIHOs and NIHOs, were developed based on the admission control algorithm of [13].

With the above testbed, we performed experiments for MIHO and NIHO and obtained some preliminary results. By physically moving the MT from AP1 towards AP2, MIHOs were triggered. We measured the time that the MT takes to select the new AP and perform the signaling that precedes handover. The measured times were about 1 s. We argue that these times are low enough for realistic scenarios, as they allow

speeds of above 100 Km/h with overlapping coverage areas of 30 m, without interrupting the communication. Handover times were also measured. The results, of about 50 ms, were also satisfactory; according to the measures of [4], this disruption is perfectly acceptable, for example, for VoIP communication.

NIHOs were forced by issuing a new QoS request such that one AP (AP1) with two MTs became heavily loaded and one of the MTs was moved to AP2 in order to unload AP1. We measured the time elapsed between the QoS request and the beginning of the handover execution (this is the extra time that NIHO needs to collect the measurements information). The times we measured were below 2 s; however, note that in a running commercial system, measurements will probably be regularly collected and already present in the QoSB, and thus these times can be substantially reduced. We also measured the handover execution time. The results obtained were similar to those in the MIHO case, as both cases involve almost the same functions (the only difference is that NIHO handovers have the advantage of not needing the scanning step in the search for handover candidate APs).

SCALABILITY AND PERFORMANCE

The experiments presented in the previous sections validate our architecture for a small scenario. In this section we analyze the applicability in a large scenario by addressing performance and scalability issues. The focus of our analysis is on the scalability of the MIHO and NIHO functions of the architecture.

Scalability of the MIHO function is jeopardized by the signaling load between CARD modules. However, since CARD signaling is only executed between neighboring ARs, this represents a natural limit for the signaling, which provides scalability guarantees. Indeed, through proper planning we can impose a limitation on the number of neighboring ARs and thus limit the signaling load. In addition, the CARD signaling load can be further reduced by performing filtering on the information exchanged and setting large intervals to regulate the periodic exchange combined with the use of event-driven notifications.

It is important to note that with the proposed architecture, MIHO decisions are performed taking into account QoS and signal strength considerations. However, to cope with this flexibility, the complexity of mobility signaling is slightly increased as compared to traditional fast handover approaches. The advantage is that this makes it unlikely that the new AP does not meet the MT requirements. In contrast, in traditional IP mobility architectures that do not consider QoS in handover decisions, it is more likely that handover decisions do not meet QoS requirements and therefore more than one handover is necessary to find a suitable AP. We argue that by reducing the number of failed handovers, our architecture is more efficient and this contributes to scalability.

Scalability in the NIHO function is challenged by the design of the PM module. Indeed, this module stores information for each AP-MT
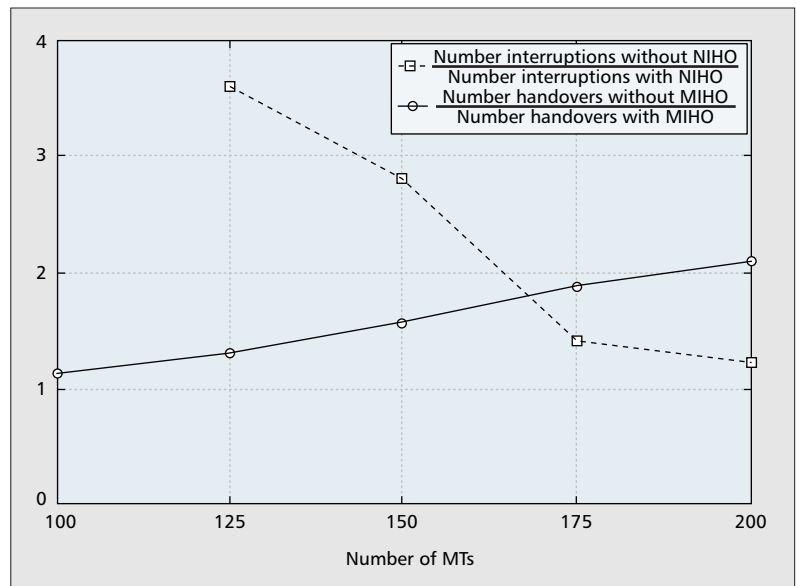

■ **Figure 6.** *Performance improvement achieved with the MIHO and NIHO functions.*

pair in a region and performs handover decisions based on these data. Therefore, scalability of the PM critically depends on the size of the region; if too large, the number of AP-MT pairs grows drastically and PM performance degrades. On the other hand, if we define too small regions, it may be difficult to optimize performance because NIHOs do not currently work between regions. Therefore, by finding the right trade-off for the size of a region and planning the network accordingly, scalability can be guaranteed while providing near-optimal performance. As NIHO yields a more efficient use of the network resources, we argue that (similarly to the MIHO case) this mechanism contributes to the scalability of the architecture.

In order to gain insight into the performance improvement achieved with the MIHO and NIHO functions, we performed the following simulation using OPNET. We had a number of MTs moving around at a speed uniformly distributed between 0 and 10 m/s in an area covered by 25 APs separated 250 m from each other. Their QoS requirements were such that at most 10 MTs could be admitted at one AP. Given this scenario, we measured the number of handovers needed with MIHO as compared to the traditional approach of handovers driven by signal strength. We then added the NIHO function and studied the improvement in terms of the number of communications interrupted due to moving into an area with all APs fully loaded. The results as a function of the number of MTs are given in Fig. 6. These results show that a very substantial improvement in performance is achieved when the MIHO and NIHO techniques are used.

From the analysis conducted in this section we conclude that, with proper planning and configuration, the proposed architecture scales well and can be used in arbitrarily large networks. Furthermore, as a result of incorporating MIHO and NIHO, the architecture is more efficient, which improves scalability.

CONCLUSIONS

The IP-based architecture presented here integrates multiple technologies in a seamless environment, is very flexible in terms of the handover possibilities, and is fully integrated with QoS support. The overall design integrates and enhances some work in progress and trends inside IETF, 3GPP, and IEEE.

The IP-based architecture presented in this article integrates multiple technologies in a seamless environment, is very flexible in terms of the handover possibilities (MIHO and NIHO, intra- and intertechnology), and is fully integrated with QoS support. The overall design integrates and enhances some work in progress and trends inside IETF, 3GPP, and IEEE. The architecture is highly scalable both at mobility and QoS-support levels. The usage of monitoring and resource management functions, and their integration in the mobility process, is also an added advantage of our work. The changes developed are conceptually simple, and can be deployed with low-cost equipment. Furthermore, the architecture has added advantages not described here, such as integrated security and AAA functionalities [3].

ACKNOWLEDGMENTS

REFERENCES

[1] J. Hillebrand *et al.*, "Quality-of-Service Signaling for Next-Generation IP-Based Mobile Networks," *IEEE Commun. Mag.*, June 2004, pp. 72–79.
[2] J. Jähnert *et al.*, "The 'Pure-IP' Moby Dick 4G Architecture," *Elsevier Comp. Commun.*, vol. 28, Issue 9, 2 June 2005.
[3] R. L. Aguiar *et al.*, "Designing Networks for the Delivery of Advanced Flexible Personal Services: the Daidalos approach," *Proc. IST Mobile & Wireless Telecommun. Summit*, Lyon, France, 2004.
[4] C. J. Bernardos *et al.*, "Experimental Evaluation of a Handover Optimization Solution for Multimedia Applications in a Mobile IPv6 Network," *European Trans. Telecommun.*, vol. 16, no. 4, July/Aug. 2005.
[5] V. Marques *et al.*, "Evaluation of a Mobile IPv6-Based Architecture Supporting User Mobility, Quality of Service and AAAC in Heterogeneous Networks," *IEEE JSAC*, Special Issue on Wireless Overlay Networks Based on Mobile IPv6, vol. 23, no. 11, Nov. 2005.
[6] X. Perez-Costa *et al.*, "A Performance Comparison of Mobile IPv6, Hierarchical Mobile IPv6, Fast Handovers for Mobile IPv6 and Their Combination," *ACM SIGMOBILE Mobile Comp. and Commun. Rev.*, vol. 7, no. 4, Oct. 2003.
[7] D. Johnson *et al.*, "Mobility Support in IPv6," RFC 3775, June 2004.
[8] R. Koodli, Ed., "Fast Handovers for Mobile IPv6," RFC 4068, July 2005.
[9] S. Blake *et al.*, "An Architecture for Differentiated Services," RFC 2475, Dec. 1998.
[10] M. Liebsch *et al.*, "Candidate Access Router Discovery (CARD)," RFC 4066, July 2005.
[11] T. Melia *et al.*, "Network-Controlled Fast Handovers," Internet draft, work in progress.
[12] J. Danda *et al.*, "Measurement Module for Mobility Management in WLAN Networks," *2nd Int'l. Wksp. on Multimedia Interactive Protocols and Systems*, Grenoble, France, Nov. 2004.
[13] A. Banchs *et al.*, "End-to-end Delay Analysis and Admission Control in 802.11 DCF," Elsevier Comp. Commun., to appear.

BIOGRAPHIES

RUI L. AGUIAR (ruilaa@det.ua.pt) received a Ph.D. degree in electrical engineering in 2001 from the University of Aveiro, Portugal. He is currently an assistant professor at the University of Aveiro and is also leading a team at the Institute of Telecommunications, Aveiro, on next-generation network architectures and protocols. His current research interests are centered on the implementation of advanced wireless networks, systems, and circuits, with special emphasis on QoS and mobility aspects.

ALBERT BANCHS (banchs@it.uc3m.es) received his M.Sc. and Ph.D. degrees in telecommunications from the Technical University of Catalonia, Spain, in 1997 and 2002, respectively. He worked for ICSI in 1997, for Telefonica I+D in 1998, and for NEC Network Laboratories from 1998 to 2003. Since 2003 he is with University Carlos III of Madrid. His research interests include QoS and resource allocation in wired and wireless networks.

CARLOS J. BERNARDOS (cjbc@it.uc3m.es) received a telecommunication engineering degree in 2003 from University Carlos III of Madrid. He is currently a Ph.D. candidate in the Telematics Department of the same university, where he has been working as a research and teaching assistant in telematics engineering since 2003. His current work focuses on IP-based mobile communication protocols.

MARÍA CALDERÓN (maria@it.uc3m.es) is an associate professor at the Telematics Engineering Department of University Carlos III of Madrid. She received a computer science engineering degree in 1991 and a Ph.D. degree in computer science in 1996, both from the Technical University of Madrid.

MARCO LIEBSCH (marco.liebsch@netlab.nec.de) received a diploma in communication electronic engineering in 1997 and an M.Sc. degree from the University of Central England, Birmingham, England, in 1999. Then he joined NEC Network Laboratories as a research staff member, first working on IP telephony protocols and on MGCP-based VOP Gateway control. In September 2000 he joined the Laboratories' Mobile Internet Group and is now working on IP-based mobile communication protocols.

TELEMACO MELIA (Telemaco.Melia@netlab.nec.de) received his informatics engineering degree in 2002 from the Polytechnic of Turin, Italy. In June 2002 he joined the Mobile Internet Group of NEC Europe Ltd., Heidelberg, Germany, as a software engineer. His main research areas are IP mobility support across heterogeneous networks and resource optimization control. He is the author of several publications in international conferences and is currently employed as a research staff member.

PIOTR PACYNA (pacyna@kt.agh.edu.pl) received an M.Sc. degree in computer science in 1995 and a Ph.D. in telecommunications in 2005 from the AGH University of Technology in Krakow, Poland. He has been working as a lecturer in the Department of Telecommunications of the AGH University of Science and Technology. His research concentrates on next-generation IP networks, mobility, and security, and he has authored several research papers in these areas.

SUSANA SARGENTO (ssargento@det.ua.pt) was graduated in electronics and telecommunications engineering from the University of Aveiro in 1997, and concluded her Ph.D. in 2003. Since February 2004 she has been an assistant professor at the University of Aveiro. Her main research interests are in the areas of next-generation heterogeneous networks, infrastructure, and ad hoc networks, specifically in QoS, mobility, routing, multicast, and charging issues.

IGNACIO SOTO (isoto@it.uc3m.es) received a telecommunication engineering degree in 1993, and a Ph.D. in telecommunications in 2000, both from the University of Vigo, Spain. He was a research and teaching assistant in telematics engineering at the University of Valladolid from 1993 to 1999. In 1999 he joined University Carlos III of Madrid, where he has been an associate professor since 2001. His research activities focus on mobility support in packet networks and heterogeneous wireless access networks.