Evaluation of a Mobile IPv6-Based Architecture Supporting User Mobility QoS and AAAC in Heterogeneous Networks

Victor Marques, Xavier Pérez Costa, Rui L. Aguiar, *Member, IEEE*, Marco Liebsch, and A. Manuel Oliveira Duarte

***Abstract*—This paper presents a Mobile IPv6-based overlay network architecture for heterogeneous environments, designed entirely based on IPv6, that aims to be implemented seamlessly irrespectively of the supporting network infrastructure. All transmission technologies are handled at the physical and data-link layers, imposing IPv6-based protocols for all higher layer communications and signaling. The architecture builds on Mobile IPv6 including improved fast handover, and integrates quality-of-service and authentication, authorization, accounting, and charging control per user. The most critical issues of the proposed architecture, mainly related to the handover process, were subject of a performance evaluation via ns-2 simulations. Finally, a field trial of the system was implemented, overlaying part of the GEANT infrastructure between Madrid and Stuttgart, which results are presented here.**

***Index Terms*—4G, authentication, authorization, accounting, and charging (AAAC), fast mobility, quality-of-service (QoS).**

I. Introduction

The EVOLUTION toward next-generation wireless networking brought life to the concept of technology and services integration. The success attained by the Internet, on one side, and by cell phone networks, on the other, led to visions of "next-generation networks," integrating the two different philosophies. On one hand, Internet protocol (IP) technology has been developed in order to support a new range and variety of services, previously only possible with circuit switching technologies, due to the very low cost of Internet access. On the other hand, cellular telephony technology also boomed during this Internet evolution, bringing commodities now very keen to humans: mobility and reachability. It was then natural to associate both concepts as a universal data-based mobile access network. Current General Packet Radio Service (GPRS) [1] and Universal Mobile Telecommunications System (UMTS) [2] technologies are just the starting points on these visions. Despite these ongoing developments, there are still many technical obstacles to overcome before the common citizen is able to achieve "multimedia communication anytime, anywhere, any style."

Although wireless transmission capabilities have steadily increased, the requirements of new applications have increased in a similar way. In this context, the fact that more resources are available does not necessarily mean that the user will have them when desired, a normal limitation with the relatively scarce radio spectrum. Specially in wireless networks, resources must be managed in an efficient way, and selected services [e.g., voice-over-IP (VoIP)] should receive better quality-of-service (QoS). This brings the need to integrate QoS support for differentiating between the traffic of different services (or users). Other control aspects, such as authentication, authorization, accounting, and charging (AAAC) systems (as traditionally considered within the Internet Engineering Task Force, but further enhanced with charging components) need also to be supported.

The work presented in this paper describes our view for a future Mobile IPv6-based overlay network architecture integrating the above mentioned requirements, and presents both experimental and simulation results on this overlay network. Although IPv4 could be considered as a possibility, the inherent address advantages and mobility efficiency of IPv6 makes this the technology of choice for future mobile networks. This architecture is positioned in a time in future beyond the traditional "All-IP" visions of industry groups (e.g., 3GPP, which still resort to specific transport protocols at the radio access level). We present an IPv6-centered heterogeneous architecture, i.e., IPv6 is the only protocol at the network layer, supporting mobility, and assuring QoS per user, and all management features required to deploy advanced services in a provider's framework (i.e., supporting AAAC functions). Security aspects are also included in this network, on the user access link. This Mobile IPv6 (MIPv6)-based architecture was implemented under the aegis of the IST Moby Dick[1] project. It supports seamless handover mechanisms across three different access networks [wideband code-division multiple access (W-CDMA), wireless local area network (LAN), and Ethernet], representative of different types of user access; mobility-enabled end-to-end QoS; IP-layer paging signaling; and AAA mechanisms enriched with auditing, metering, and charging (AAAC).

The rest of this paper is structured as follows. In Section II, we present the overall network architecture, showing access

Manuscript received December 1, 2004; revised June 21, 2005.

V. Marques is with the Portugal Telecom Inovação, Rua Eng. José Ferreira Pinto Basto, 3810-106 Aveiro, Portugal (e-mail: victor-m-marques@ptinovacao.pt).

X. P. Costa and M. Liebsch are with NEC Network Laboratories, NEC Europe, Ltd., D-69115 Heidelberg, Germany (e-mail: perez@netlab.nec.de; marco.liebsch@netlab.nec.de).

R. L. Aguiar and A. M. O. Duarte are with the Department of Electronica e Telecomunicacoes and the Institute of Telecommunications, Universidade de Aveiro, 3810 Aveiro, Portugal (e-mail: ruilaa@det.ua.pt; duarte@ua.pt).

Digital Object Identifier 10.1109/JSAC.2005.856825

[1]IST-2000-25394, Mobility and Differentiated Services in a Future IP Network: http://www.ist-mobydick.org

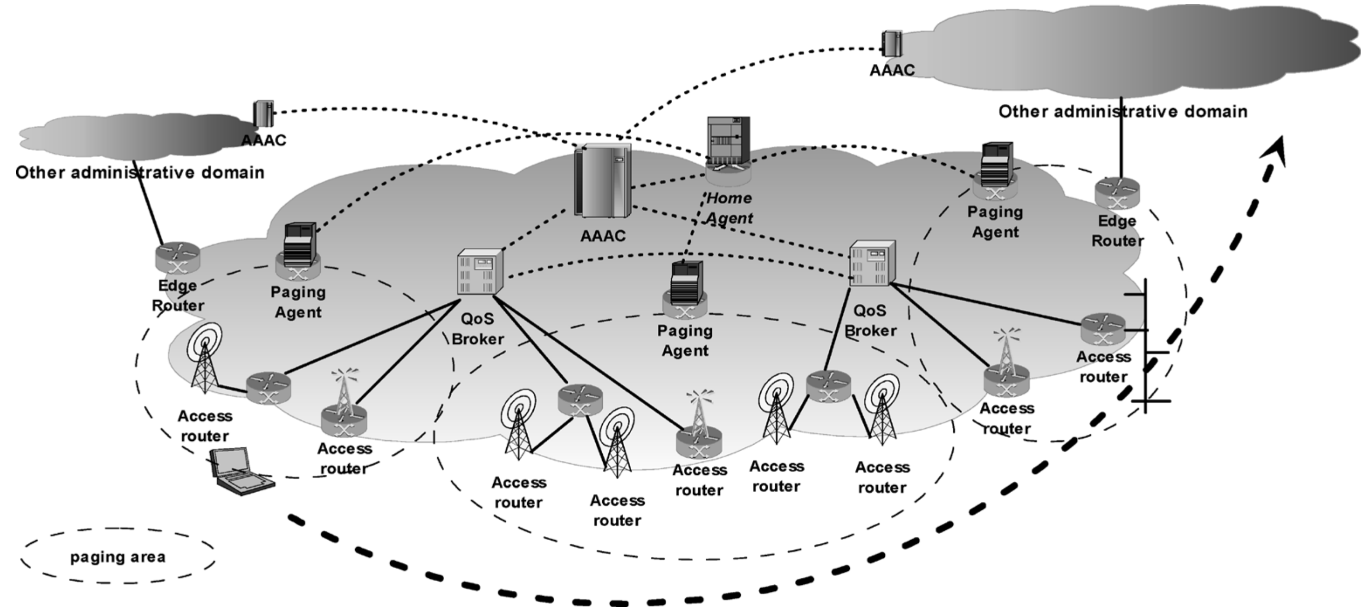

Fig. 1. General network architecture.

technologies, architecture elements, and their functions. In Section III, mobility related aspects are presented, covering the deployed fast handover (FHO) mechanism and IP paging in detail. Section IV presents the QoS framework, covering the main elements involved. In Section V, we present the AAAC support system. Section VI discusses the integration of these aspects (mobility, QoS, and AAAC) and provides an overall view of the system operation, describing basic signaling flows. Section VII discusses our simulation and experimental validation results and highlights the timing limitations of the architecture. Finally, a summary is presented in Section VIII together with our main conclusions.

II. Basic Network Architecture Elements

Our network architecture was designed to be agnostic to the type of layer 2 technologies used, and a real implementation was developed supporting three different types of access technologies (Ethernet, wireless LAN, and W-CDMA). The architecture is composed of mobile terminals (MTs, the user terminals), access routers, or gateways (ARs, the MT access point to the network), an AAAC system per administrative domain (supporting all AAAC functions), a home agent (HA) for IPv6 mobility, QoS brokers (QoSBs) (managing the access networks), and IP paging agents (PAs). These elements are depicted in Fig. 1. The system is controlled by a Network Management System (not depicted), which configures adequately all the network devices. All signaling between these entities is exchanged at the IP-layer level, in order to achieve a convergence layer independent of the L2 technology. We consider that in some situations, only one active connection to a wireless AR can exist, due to hardware and L2 technology limitations. The functions provided by these entities and their interworking build an architecture which supports an efficient seamless network access on this heterogeneous environment.

To provide seamless mobility, i.e., the user not aware that a handover is under progress, competitive to the existing in current cellular networks, Mobile IPv6 with FHO [3] is used with specific QoS enhancements and associated to context transfer techniques for access link security [4]. This IP-based approach has the advantage to increase the scope of mobility across cells based on different access technologies. Network access is provided by a (potentially wireless) AR, which controls a single (wireless) cell: on the wireless access, an IP subnet is directly mapped to a radio cell. Changing cells becomes then a process managed at the IP-level, as all other handovers in our network. For achieving this, W-CDMA support (used in TDD mode) had to be achieved by directly connecting base stations to the IP network, thus eliminating several of the elements defined in the 3GPP UMTS architecture (RNC, SGSN, GGSN) [5], simplifying its architecture and further achieving intertechnology handovers in that process.

QoS provision resorts to a differentiated services (DiffServ)-based model [6], because of its higher scalability and reduced signaling overhead. The association of a DiffServ framework with the use of localized QoSBs supports QoS on large scale, while simultaneously providing control and local optimization of the usage of access network resources. In our model, we assume that the core network is overprovisioned, and no explicit QoS control is required outside of the access links. In these links, the QoSBs instruct the access routers (ARs) regarding which flows to accept, performing connection access control based on measurements received from the ARs.

AAAC support is based on the AAA IETF Architecture [7], enriched with charging mechanisms to provide an overall architecture targeted to commercial use. Access link security is implemented based on IPSec, and adequate modifications to IPSec packages have been made to support mobility, allowing security related information to be exchanged between ARs during handover using context-transfer techniques. Furthermore, a novel IP paging concept was integrated into the overall architecture to decrease power consumption of the MT, and to further optimize the usage of the (scarce) radio access medium. All these mechanisms are explained in the next sections.

III. Mobility Support

Mobility support inside our architecture comprises both mobility (Mobile IPv6 and FHO) procedures able to operate across multiple networks and paging provision at the network level.

A. Fast Handover (FHO)

Mobile IPv6 [8] is used as the basic mobility management scheme for our network. While "standard" MIPv6 is deployed for global mobility management (i.e., inter-administrative domain handover) extensions for local mobility management (i.e., intra-administrative domain handover) are required in real-time environments (such as voice communications) for achieving seamless handover.

Early work on Mobile IPv6 almost exclusively dealt with IPv4 networks [9], [10]. Perkins and Wang [9] used *ns-1* to analyze the effects of route optimization and buffering (smooth handoff), Campbell *et al.* provided a comparison of IPv4 micromobility protocols in [10]. Because of the significant differences between MIPv6 and Mobile IPv4 (MIPv4) results obtained for MIPv4 do not take over for MIPv6. Regarding MIPv6 there are several proposals to enhance handover performance which can be divided in hierarchical [11] and nonhierarchical [3] approaches. A protocol overview of these proposals and improvement ideas are provided in [12]–[14]. Our analysis and simulation results [14], [15] concluded that a nonhierarchical FHO approach [3] was the most suitable mechanism for our network, as the best compromise between network complexity and seamlessness of the handover.

A handover is composed of two different types of connection changes: low-layer handover (which may pose rigid technology specific constrains) and IP-layer handover. The chosen IP FHO approach is independent of the access technology (W-CDMA, wireless LAN, or Ethernet) and follows a "make-before-break" philosophy: layer 3 handover will be prepared via the existing communication channel, before layer 2 handover is performed. In this way, the handover delay is reduced to the minimum amount of time necessary for the eventual reconfiguration or changing of the interface. During this preparation phase, the current AR is informed about the "intended handover" and for the duration of this handover process, bicasts all packets destined for the mobile node both to the current care-of-address (CoA) and to the new AR. This mechanism reduces packet loss during handover: explicit signaling controls the start of the bicasting in the access link, effectively providing seamless handover. Although extra bandwidth is required at the core network for the bicasting, this is not usually a problem in that part of the networks.

The detailed signal flow of this process, and the integration of QoS and AAAC aspects, is further discussed in Section VI.

B. Paging

As a requirement of the standard MIPv6 concept used in the network, an active MT acquires a different CoA in each cell,

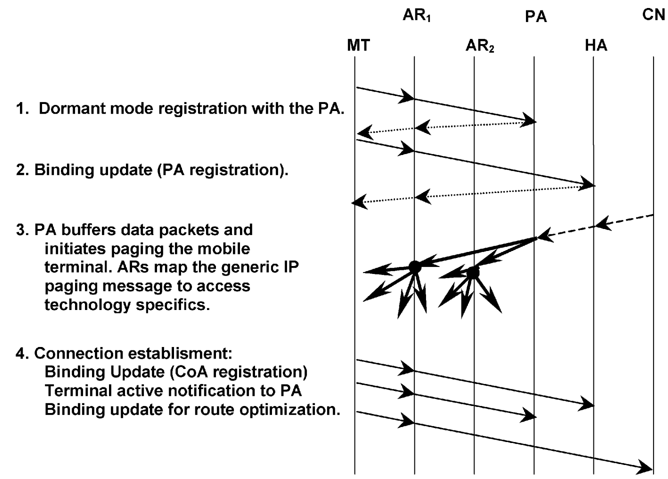

CN: Correspondent Node, MT: Mobile Terminal, HA: IPv6 Home Agent, PA: Paging Agent, AR1: Access Router 1, AR2: Access Router 2

Fig. 2. Illustration of the dormant mode registration with the PA, the paging process, and reestablishment of routing information after reactivation.

which identifies the IPv6 subnet in which the MT is currently located. The CoA is then to be registered with its MIPv6 HA. Our approach to reduce the amount of signaling between the MT and HA, while maintaining the MT reachability was to implement PAs, responsible for a paging area composed by several ARs and paging attendants in each AR.

An MT switching to dormant mode first discovers the responsible PA through the paging attendant function in its current AR, and notifies this PA of its current paging area. The paging area information is retrieved from router advertisement messages, as each AR advertises a specific paging identifier. After this, the MT registers the PA address with its HA by means of a standard MIPv6 binding update message carrying the "alternate CoA suboption" (see Fig. 2). This additional option allows a MT to register with its HA, as a binding, an IP address different of the MT CoA. After this, the MT can enter dormant mode and is allowed to roam within the registered paging area without the need to send location update information to the network. When roaming across paging areas (detected by the change in the paging identifier of the router advertisements), the dormant MT issues a paging area update message to the PA.

When a correspondent node (CN) addresses data packets to a dormant MT's home address, the HA intercepts these packets and forwards them to the registered PA by means of IP tunnelling, as standard Mobile IPv6. The PA terminates this tunnel, and buffers the initial data packets until the paging process has resolved the MT's current location, and it becomes again able to receive traffic packets. In this paging period, the PA sends an IP paging request message to (the paging attendant in) all ARs of the registered paging area (see Fig. 2). Individual paging attendants may then build link-level paging request messages, which are sent through the respective access technology (wired Ethernet, wireless LAN, and W-CDMA). On reception of one of the link-level paging messages, the MT reactivates itself and reestablishes detailed routing information with the network, notifying the PA and the HA of its current CoA (see Fig. 2). As a consequence, the data packets buffered at the PA are forwarded to the MT, and further data packets intercepted at the HA are now directly forwarded to the reactivated MT. Route optimization, through binding update information to the CN, is then also possible.

This solution, as shown in [16], saves frequent location updating (binding updates), decreases signaling costs, and saves scarce radio bandwidth, as well as increasing MT battery power usage efficiency.

IV. QoS Architectures for Mobile Service Provisioning

The QoS architecture has to support end-to-end QoS, easily manageable from an operator point of view. To achieve this objective, entities and methods had to be defined for the scalable allocation and control of the resources in the access networks, able to offer and guarantee end-to-end QoS and maintaining user connectivity and QoS, while the MT is moving.

A hard constraint to our architecture is the simultaneous support of mobility and QoS. Several IETF QoS frameworks were considered before designing the final QoS architecture, taking this integration in consideration. The pros and cons of Integrated Services (IntServ) [17] and DiffServ [6] models have been widely discussed in the literature and are well known; unfortunately, none of them has specific support for mobility. Not even the hybrid solutions we were aware of seemed to solve the integration of mobility and QoS. Thus, we developed an innovative usage of QoSBs, associated with FHO, incorporated in a traditional DiffServ approach, to be able to control and manage available resources in an efficient way, while enabling user mobility [4].

The architecture relies on the concept that the user will be granted with "services" previously contracted with the provider. The QoSBs are in charge of allocating resources in the access network, per user and per subscribed service, according to the contractual information to the user. As discussed in [18], these services will usually be fixed transport services (e.g., "best effort with 32 Kb/s," or "guaranteed 64 kb/s"), but mechanisms for potential dynamic service negotiation are also in place. QoSBs also manage flow aggregation of resources in the core network.

For providing user-oriented QoS services, the QoSB interacts with the AAAC system during the (required) user registration phase (see Section VI-A). When the user registers at the network, the AAAC system dumps to the QoSB all relevant user-specific information for QoS provisioning. Having this user-specific information, the QoSB can perform service admission control decisions on every service request done by the user's terminal. For that, the QoSB also interacts with the ARs in its QoS domain. These interactions are required both for AR's QoS configuration and for service authorization. To reduce signaling overhead, the system is designed in such a way that the user/terminal does not need to perform explicit resource reservation or release (although these messages can be supported). "Services" are requested by simple DSCP marking on outgoing packets, with each DSCP corresponding to a different *contracted* service. When the AR receives an incoming packet from a user with a certain DSCP value, it sends a configuration request to the QoSB. The QoSB, based on the information for this user, will then configure the AR with the appropriate QoS policies. Services are stopped implicitly, by an inactivity

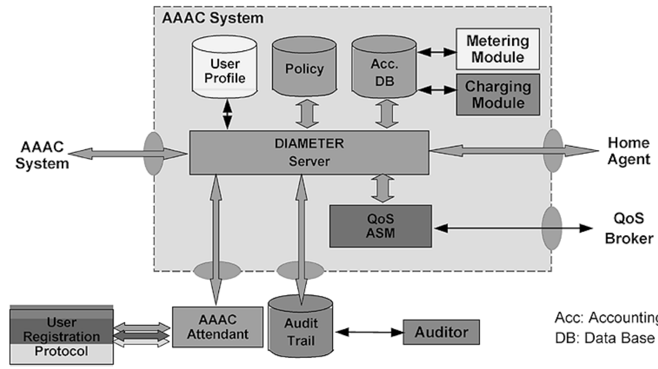

Fig. 3. Enhanced generic AAA architecture supporting QoS-enabled mobility management.

timeout. This concept is described in greater detail in [18], and detailed signal flows are presented in Section VI.

For each layer 2 technology supported, it is necessary to map between IP QoS parameters and those specific of that technology. The IP QoS control procedures, based on DiffServ codepoints, must be also directly mapped in the physical layer control, at the AR. Specific middleware exists both at the MT and the AR, which perform these tasks [4].

V. Mechanisms for Service Provisioning in a Heterogeneous Mobile Environment

Our AAAC architecture (Fig. 3) is a basic IETF-based AAA architecture (authentication, authorization, and accounting) [7] enriched with auditing, metering and charging, and optimized for IPv6. The fact that the AAAC services would be used in a QoS-enabled MIPv6 environment has been considered in the architecture design to provide features which enable new functionality, as well as performance optimization of the overall system.

This auditing function enables further functionality in the AAAC, with respect to the evaluations of audit trails generated by AAAC system and other entities. Within this context the policy repository is considered as being part of a policy-based AAAC system.

The AAAC system supports multiple interfaces. An AAAC attendant handles the interface with the MT: all communication is done through a user registration protocol (URP). An application specific module (ASM) communicates with the QoSB. The advantage of developing ASMs is the added flexibility: a variety of service equipment can be easily addressed with a uniform method from the AAAC system point of view. AAAC system to ASM communications use the AAAC protocol (DIAMETER), and the ASM to service equipment communications can be done by equipment-specific protocols (in our case, generally, either COPS [19] or URP). Note that in this architecture, a clear distinction is made between services offered to the user, such as QoS-enabled services, and services required for the operation of the AAAC system, such as charging. Therefore, the former ones are in general accessed and provided via an ASM and the extended AAAC protocol, while the latter ones are communicating directly with the AAAC system, using dedicated communication means if required.

AAAC systems communicate with each other via the DIAMETER base protocol suitably enhanced with adequate extensions.

A key element for QoS provisioning and charging is a meter that is able to account service usage. Within the IETF real-time flow measurement (RTFM) working group an IP-based metering framework was defined [20]. A publicly available network meter, NeTraMet, is commonly used for network monitoring and measurement. The basic component of this reference implementation is a traffic meter (NeTraMet agent), which captures IP header information according to a predefined rule and assigns this information to the different IP flows. The definition of these IP flows is quite flexible according to the needs of the network administrator. In our IPV6 network, the meter (enhanced to run on IPv6) is configured according with the description of the "services" subscribed by the user, and communicated to the QoSB.

VI. Key Scenarios/Mobility, QoS, and AAAC Integration

For the network, three distinct phases of operation and control can be clearly identified: 1) registration—in this architecture, a MT/user may only start using network resources after authentication and authorization, as in today's cellular networks; 2) authorization—the user has to be authorized to use specific services previous to its release by the network; and 3) handover—the user needs to have its existing resources reservations transferred from one AR to another, due to its movement.

A. Registration

The registration (see Fig. 4) process (AAA and mobility) is initiated after a CoA is acquired by the MT via stateless autoconfiguration, using layer-2 identifiers; the uniqueness of these addresses is checked at the registration stage, when duplicate address detection (DAD) is performed (through a combined process of ARP requests and special functionalities on the QoSB). Getting this uncertified CoA does not entitle the user to consume resources besides registration messages—but potentially allows emergency calls (ARs are configured to allow specific connections to emergency addresses, the equivalent to a 911 call). For accessing the network otherwise, the MT has to start the authentication process by sending the user authentication information (*message 1*, Fig. 4) to the AR. That request will be then forwarded to the AAAC system (*message 2*) responsible for that AR. Notice that the registration is done in function of the user, and not of the terminal (e.g., as happens in current cellular equipment).

In a roaming case, which is more complex, the *Domain A* AAAC will issue a registration request to the MT's home AAAC (*message w*). Here, the AAAC server A is playing the role of the foreign domain AAAC server, which needs to contact the MT home AAAC. This AAAC first checks if there is a formal contractual relationship between the administrative domain this request comes from and its own administrative domain (the equivalent to current roaming agreements). Then, in case of a positive result, the home AAAC performs authentication by verifying the provided credentials. Then, in the positive case, the home AAAC sends to the HA a request for that user (*message x*) to which the HA answers (*y*). Then, the home AAAC finally answers the *Domain A* AAAC (*z*).

One attribute of this positive acknowledgment is a user profile containing the required information to provide the services requested in the foreign domain. The user profile is a centrally managed profile containing all relevant user-specific information for service provisioning. One section of this profile, the network view of the user profile (NVUP), will be forwarded from the AAAC server to the QoSB (*message 3b*) (which also performs DAD at this stage), and a different set of the profile will be sent to the AAA Attendant located at the AR (*3a*). The NVUP contains all required information relevant to network service provisioning, while metering and security related information will be forwarded to the AAA attendant. The AAAC will also inform the MT of the successful registration, via the AR (*messages 3a and 4*). Afterwards, the AR will initiate accounting for that user, and informs the AAAC (*message 4a*). The authentication phase is thus completed, and the user can access the network.

B. Authorization/Session Setup

Fig. 4 also shows how each network service is authorized (*messages 5 to 11*). First, the MT sends a packet (*5*) with the DSCP code marked to request a specific subscribed service (e.g., a prioritary network access at 256 Kb/s). This trailer packet can be either a packet with real information or just a dummy packet, depending on the configuration on the MT. If the "requested service" does not match any policy already set in the AR, the AR issues a request (*6*) to the QoSB, through the QoS manager (*5a*). Upon analyzing the request, and based on the user NVUP and on the availability of resources, the QoSB eventually decides on an answer to the AR (*7*). The AR QoS manager will then either configure the AR with the appropriated policy for that user/MT service (*7a*), or informs the User/MT of a service denial (*7b*). After (*7a*), any other conformant packet sent from the MT that matches the configured policy rule will be able to cross the network (*8*). Packets with a different DSCP code will restart the authorization process once more. Upon reaching the end-domain where the other user or CN is, the marked packet starts another QoS authorization process (*8a*). The QoS manager on this AR will send a policy query to the QoSB (*9*), to which the QoSB answers (*10*). If there are resources, this answer is a positive answer and the QoS manager will configure the AR (*10a*) with the right policy, otherwise, the AR will send back a service denial message (*10b*). After (*10a*), the next packets matching the installed policy will be able to reach the other terminal (*11*).

With this approach both access networks provide contracted levels of QoS; core network is then monitored to check if the expected performance is being provided, such that the contracted end-to-end QoS is being fulfilled.

C. QoS-Enabled Handover

One of the most difficult problems when dealing with IP mobility is assuring a constant level of QoS. User mobility is performed in our network by means of FHO techniques in combination with message exchange to and between the QoSB during the handover, as shown in Fig. 5.

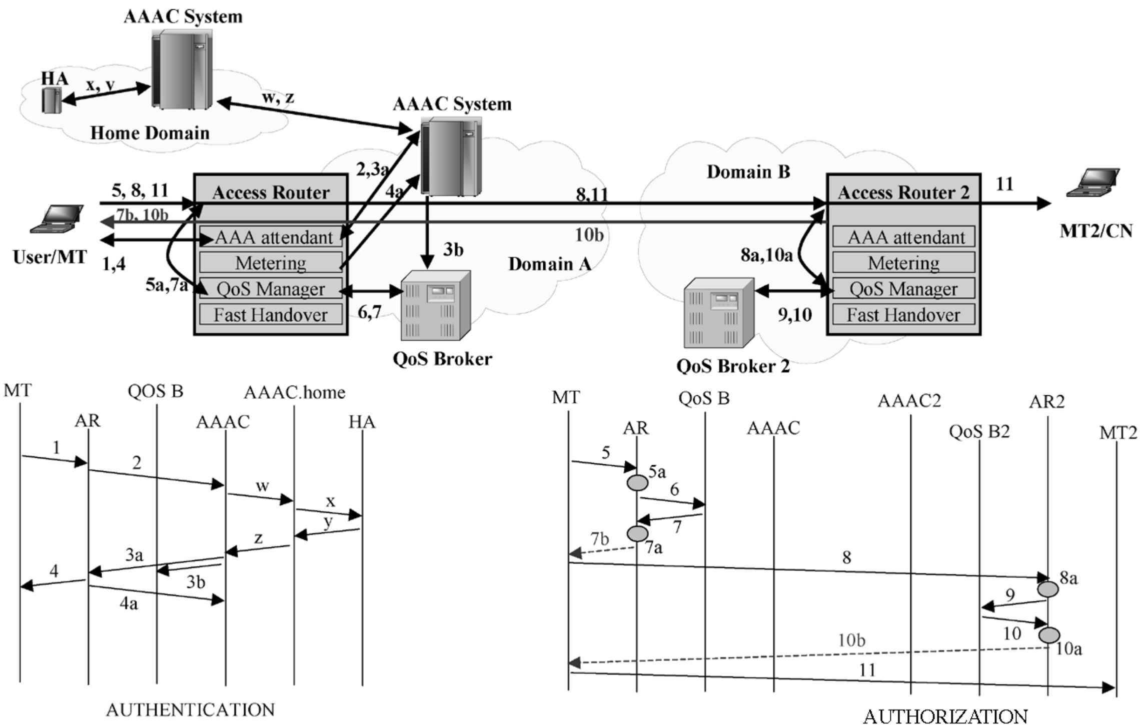

No.	Message/action	Content / Parameters	Remarks
1	AA Request	NA-I; credentials; CoA	NAI: Network Address Identifier
2	AA Request	NAI; credentials; CoA	AR as proxy of the MT
w	AA Request	NAI; credentials; CoA	Ask home AAAC for user information
x	AA Request	CoA	CoA to HA mapping request
y	AA Response	HA	CoA to HA mapping information
z	AA Response	**2x key** (MT, AR), Profile SubSet**,** Session	Info for AAAC foreign
3a	AA Response	**2x key** (MT, AR), Profile SubSet**,** Session	Info for AR and MT
3b	NVUP dump	MT, AR, Profile SubSet (includes timeout)	NVUP: Network View of the User Profile
4	AA Response	1x key(MT, AR), Profile SubSet (codes), Session Timeout	Dumps the proper DSCP codes to use in this network
4a	Accounting Start		
5	Service Request	DSCP, CoA, Destination Address	
5a	<event>		Request to QoS Manager
6	Authorization Request	DSCP, CoA, Destination Address	AR as proxy for MT
7	Authorization Confirmation/Denial	Policy for the requested Service/Denial of Service	AR configuration information or service denial
7a	<event>		AR configuration
7b	Service denial		If the user profile does not allow that service or if there are not enough resources available
8	Service Request	DSCP, CoA, Destination Address	
8a	<event>		Request to QoS Manager
9	Authorization Request	DSCP, CoA, Destination Address	
10	Authorization Confirmation/Denial	Policy for the requested Service/Denial of Service	AR configuration information or service denial
10a	<event>		AR configuration
10b	Service denial		If there are not enough resources available
11	User data		

Fig. 4. End-to-end QoS support, registration, and service authorization.

When a MT starts losing signal strength to the current AR ("old AR") (*message 1*, depicted in Fig. 5), it starts a handover procedure to a neighboring AR ("new AR") from which it receives a beacon signal with the network prefix advertisement (*2*). The MT builds its CoA and initiates the handover procedure, sending an IP-handover request to its new AR, but still through the old AR (*3*). The FHO module in the old AR will forward this request to its QoS manager (*message 3a*) and the FHO module on the new AR, known by the network prefix (*4a*). The QoS manager immediately forward this request to the QoSB (*4b*) ("old QoSB"). The old QoSB sends a handover request to the new QoSB (*5*), indicating the user's NVUP and the list of services currently being used by this user. Basically, this acts as a context transfer from the old QoSB to the new QoSB (in par-

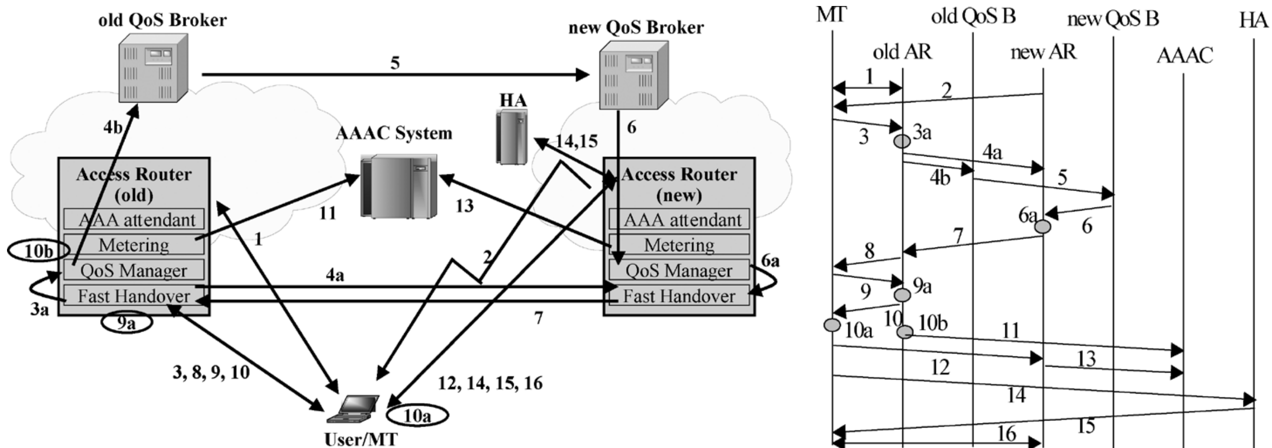

No.	Message/actions	Content / Parameters	Remarks
1	Current connection	User data	
2	Router Advertisement	Network prefix	New router on range
3	Router Solicitation for Proxy	NARaddr, new CoA	Request FHO
3a			QoS Manager is informed
4a	Handover Initiate	SubSubProfile, key, new CoA	CT info for AAA as byte-stream conveyed
4b	QoS message	NAR, oCoA	Indication of nAR ID, oCoA
5	QoS message	HoA, nCoA, DSCP in use	Carry NVUP (HoA, nCoA, DSCP in use)
6	QoS message	Configuration data, result info (command and DAD check result)	Carries configuration data for nAR or info on ResReserv failure, indication on DAD result
6a			Result of handover request sent to FH module
7	Handover Acknowledge	SPI	Confirm or reject handover
8	Proxy Router Advertisement	SPI	Proxy for the RG – confirm or reject HO
9	Handover Execute (FBU)		Start handover
9a	Start Bicasting (& Timer)		Start delay control for bicast.
10	Handover Execute ACK		FHO completed
10a	Leaving Old Link		Layer 2 handover
10b	Bicasting timer expired		Stop bicasting
11	Accounting Data (stop)		All accounting information is sent to AAAC
12	Neighbor Advertisement		Connect to new AR
13	Accounting Start		Start accounting process in new AR
14	Binding Update		Update information in the HA
15	Binding Acknowledge		Complete handover
16	User data		Data transfer

Fig. 5. End-to-end QoS support—handover with QoS.

allel, the new QoSB performs DAD). With this information, the new QoSB will verify the availability of resources, and sends a message to the QoS manager on the new AR (*6*) indicating whether the MT may or may not perform the handover. This mechanism allows the QoSB to abort the handover due to QoS constraints (e.g., missing bandwidth resources). If the handover is possible, the QoS manager, then sends this information to the FHO module (*6a*) and performs the configuration of the new AR to accommodate the MT when it moves. Meanwhile, the FHO module sends the handover reply back to the FHO module on the old AR (*message 7*). Upon receiving this message, the old AR sends it to the MT (*8*). Upon a positive answer, the MT sends a handover execute to the old AR (*9*), to which the old AR reacts initiating bicasting, setting up a timer (*9a*) and replying to the MT (*10*). The MT may now perform a successful layer 2 handover (*10a*), because all layer 3 resources have been reserved. Thus, it sends a neighbor advertisement to the new AR (*12*). The new AR starts an accounting process in the AAAC system for that user (*13*). To complete the handover process, the MT must send the binding update to its HA (*14*) that replies with a binding acknowledge. Meanwhile, the bicasting timer on the old AR expires (*10b*), which also makes it send all accounting information relative to that user to the AAAC system (*11*). The MT has now completed the handover to the new AR (*15*). If the handover is done inside the domain of only one QoSB, that is, if the QoSB controlling the both ARs is the same, then *message 5* will not exist. Everything else is exactly the same.

Notice that AAAC attendants are also informed of the handover, and the current user AAAC parameters (e.g., for metering) are exchanged directly via the handover initiate messages (*message 4a*). Security related information is similarly exchanged between the QoSB.

VII. Architecture/Protocol Validation

The described architecture and its different elements were implemented in a testbed and several validation tests were performed. This section presents some theoretical limits to the architecture performance, identifying performance bottlenecks, presents simulation results for the handover procedure and also measured values in a trial testbed [21]. This trial testbed was supported by Linux-based boxes, added with specific hardware for the W-CDMA interface, and closely resembles the network depicted in Fig. 1.

There are three important performance parameters for this architecture, in terms of user perception: the setup delay (or registration), the service initiation delay, and especially, the handover delay. This last parameter has to be kept at a minimum to guarantee a seamless handover.

A. Registration Delay

The registration delay is the delay since the user turns on its terminal until it becomes operational for service. There are two types of delays involved, the "low-layers" delay (layers 1 and 2), the delay since the terminal is enabled until is associates with one access technology; and the "high-layers" delay, the delay since the terminal sends *message 1* until *message 4* arrives (in Fig. 4)

$$Registration\_Delay = LowLayerDelays + HigherLayerDelays.$$

The first term is technology dependent, and will not depend on the specific network protocols or architecture used. The second term is dependent on the links speeds, devices performance (AR, AAAC) and, if the user is in roaming, the "electrical distance" between the foreign AAAC and the home AAAC

$$HigherLayerDelays = \sum LinkDelays + \sum ProcessingDelays + AAACsDistance.$$

The first term, the link delays between the ARs and the AAAC system, is usually negligible, as the management infrastructure has normally overprovisioned links with dedicated resources. Processing delays in the AAAC systems (and ARs), especially if they are overloaded (with other requests) or subdimensioned (databases too big for the processing power of the elements) can be dominant, but should be easily avoidable with proper computational and routing equipment in real-production networks.

Thus, the most restrictive factor seems to be the distance between the foreign AAAC and its home AAAC (in terms of propagation time), when the user is roaming. In this case, the overall registration delay will be mostly determined by this time.

This delay is not an essential performance parameter, as long as it is kept reasonably low (in human perception terms).

B. Session Setup Delay

The session setup delay is the amount of time since the user wants to start a communication until he is able to access the network. In Fig. 4, the session setup delay is the time it takes from *message 5* until *message 11* goes through. This delay is composed by the processing time in the ARs and QoSB, the distance between the ARs and link delays between: 1) the user terminal and the AR and 2) the ARs and the QoSB

$$SessionSetupDelays = \sum LinkDelays + \sum ProcessingDelays + ARsDistance.$$

The first term is usually very low, assuming normal network conditions. The processing delays are the delays in both QoSB and in both ARs. Once again, the previous comments on performance can be applied, especially, as the QoSB has an architecture which allows for easy distribution of load (that is: the ARs) between multiple machines.

The results obtained in the field trial (see Section VII-E1) showed that the processing delays are negligible. Therefore, the distance between the ingress AR and the egress AR may impose the dominant delay. However, this is a network limit rather than an architecture limitation.

C. Handover Delay

The handover delay is the critical parameter in terms of user sensitivity. The handover must be processed as fast as possible to be seamless to the user, since otherwise no multimedia communication can be supported. The handover delay is composed by several independent delays that must be minimized independently: transmission delays, computational delays and layer 2 handover delays. Some impose hard constraints, while others may be optimized or neglected

$$Handover\_Delay = \sum Transmission\_Delays + \sum Computation\_Delays + \sum Layer2Handover.$$

The transmission delays are composed by the sum of the delays between the MT and the AR, the AR and the QoSB, between both QoSB, and between both ARs. The global handover delay is the time spent between *message 3* (on Fig. 5) and *message 12*. However, the time that the user terminal is without connectivity and resources allocated is only the time spent in the layer 2 handover (which is the time elapsed between *event 10a* and *message 12*). The handover behavior was simulated (see next section) in order to access its "seamless" feasibility. After implementation, the measurements done in the trial environment (see Section VII-E2) proved our architecture concept, obtaining reduced global handover delays, and negligible packet loss (due to the bicasting process). Note that the global handover delay has impact on the relation between the coverage (radius) of each wireless cell and the speed the user can move itself, thus affecting cellular network planning.

D. QoS Enabled Handover Simulation

The most critical feature of this architecture, the handover process, was subject of a simulation study using the ns-2 simulator [22], to evaluate the performance of the handover signaling. For this, the ns-2 code was enhanced to implement the signaling presented in this paper.

Fig. 6 presents the simulation scenario used to evaluate the handover process. This base scenario consists on a MT (initially located at [10;185]) traveling between two ARs (located at [50;25] and [300;25]) on an area of 350×200 m. Each AR is connected to a core router. The QoSB, programmed to introduce 17 ms processing delay (larger than final values attained in real implementation), connected to the core routers, is the responsible for managing the resources of both ARs. The HA and a CN are also part of this scenario. The wireless connection between the base stations (ARs) and the mobile node is wireless

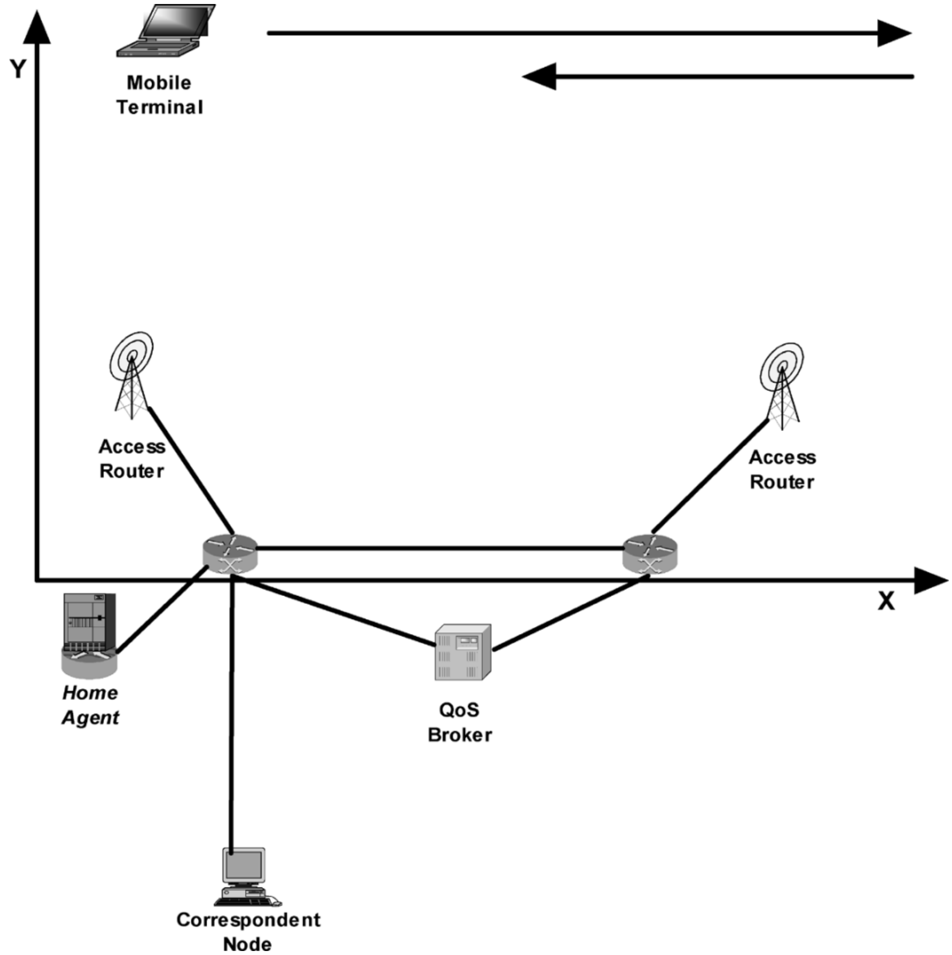

Fig. 6. QoS-enabled handover simulation scenario.

LAN. The physical links are 5 Mb/s links with 2 ms delay between nodes. Each simulation run had 50 s. Three seconds after simulation start, an application [transmission control protocol (TCP) or user datagram protocol (UDP)] is started. TCP is the most widely used transport protocol. We simulate endless FTP sources to understand the impact of this particular mobility architecture on the congestion control mechanism of TCP. The UDP flow simulates a video call with 100 bytes data packets sent each 0.005 s (190 Kb/s average bandwidth), in line with the expected next-generation personal services. The MT starts moving to [340;185] with a speed of 10 m/s, 5 s after simulation start. At $\mathrm{time} = 35$ s, the MT initiates the return to [10;185]. At $\mathrm{time} = 47$ s, the TCP/UDP applications are stopped.

The simulation study performed had two major goals:

1) to evaluate the signaling timings during the handover period;
2) to evaluate the handover performance while the mobile node is receiving data from TCP and UDP applications.

To better evaluate the handover signaling performance, some of the default scenario conditions were subject to variations, such as the increase of the link delay between the two ARs and between the CN and the ARs. Other variations performed were the increase of the MT speed to 80 m/s and the variation of the QoSB processing delay to 3 and to 37 ms.

TABLE I
HANDOVER SIGNALING SIMULATION RESULTS

	Message/Action	Simulation Time (s)	
		TCP	UDP
1	Router Advertisement	15.3006	15.5022
2	Router Solicitation for Proxy	15.303	15.5027
3	Handover Initiate	15.3284	15.5037
4	QoS message (handover request)	15.3284	15.5037
5	QoS message (handover decision)	15.35	15.5253
6	Handover Acknowledge	15.3543	15.5297
7	Proxy Router Advertisement	15.3606	15.536
8	Handover Execute (FBU)	15.4315	15.6092
9	Start Bicasting (& Timer)	15.4569	15.6305
10	Handover Execute ACK	15.4569	15.6305
11	Attaching to new Link	15.5477	15.6371
12	Neighbor Advertisement	15.5477	15.6371
13	Binding Update	15.5546	15.6473

1) Simulation Results: This section presents the simulation results both from the base scenario and from the variations.

Table I presents the time instants (from the beginning of the simulation) when the handover signaling messages are exchanged.

The table analysis allowed us to conclude the following.

1) The handover preparation phase (before the actual layer 2 handover) would take less than 150 ms (time difference between message 2 and 10).
2) The complete handover (including MT binding update) would take about 250 ms at most.

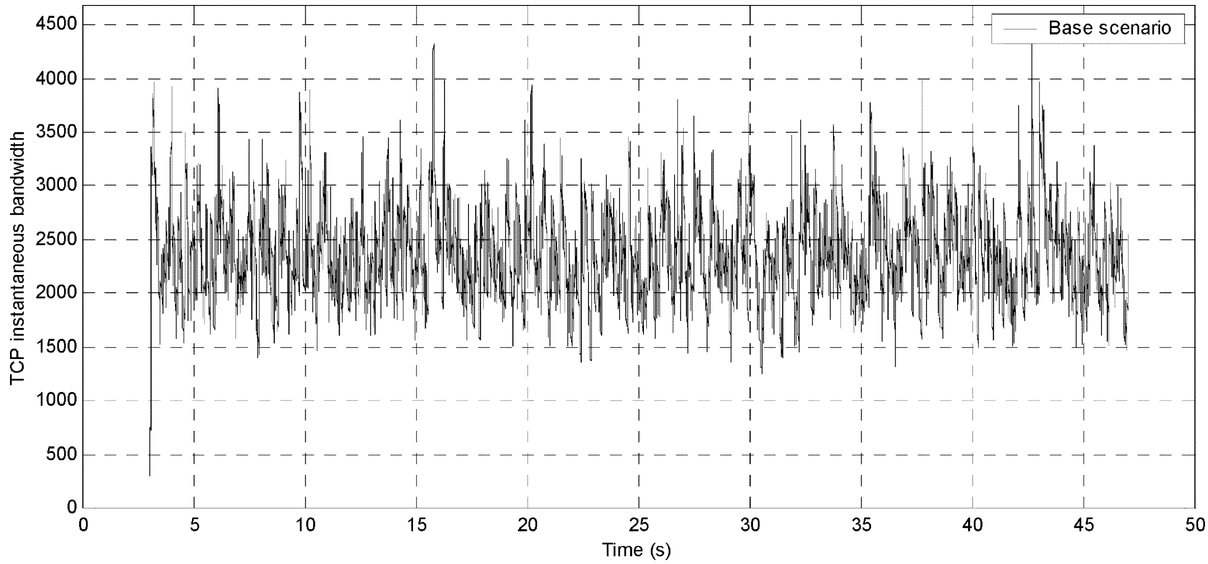

Fig. 7. Instantaneous TCP flow bandwidth.

3) When the wireless link is not congested, as it is the case of the UDP traffic (TCP by definition tries to use the full amount of available bandwidth), the handover timings were even lower (130 and 230 ms, respectively).

In this simulation scenario, the signaling messages and the traffic have exactly the same priority in every node (there is no QoS differentiation between them). This implies that, especially for the TCP case, the timings may be negatively influenced by the concurring traffic.

We may easily conclude that a MT travelling at 250 km/h needs less than 11 m to prepare the layer 2 handover, and all the handover process is concluded in less than 17 m. This is a reasonable value for cell overlapping and MT speed in next-generation networks.

However, to better understand how the applications are affected by the handover process, the next figures show the behavior of the TCP and UDP flows in terms of bandwidth (instantaneous and average), latency and jitter. For the UDP flow it is further possible to show the analysis of out of order packets at the destiny. It is also shown the impact of the variations performed in some of these flow characteristics. For the clarity of the figures, whenever it is not possible to differentiate the result of the base scenario from the result of the variations, these last are omitted (basically, meaning that the variation performed does not have any particular impact).

Fig. 7 presents the instantaneous TCP flow bandwidth between the CN and the MT. At time 15,4 and 42,9 (approximately) there were handovers. There are two small bandwidth peaks just after the handover. The bicasting process introduces a small delay to the packets that are delivered after travelling from old AR to the new AR. This delay causes these packets to arrive at the MT together with others that travel directly from the CN to the MT (after the binding update at the new network). These packets arriving at the same time cause an increase of acknowledges, and the consequence is a temporary instantaneous bandwidth increase after the handover. As an overview analysis to this result, it is possible to conclude that the handover process

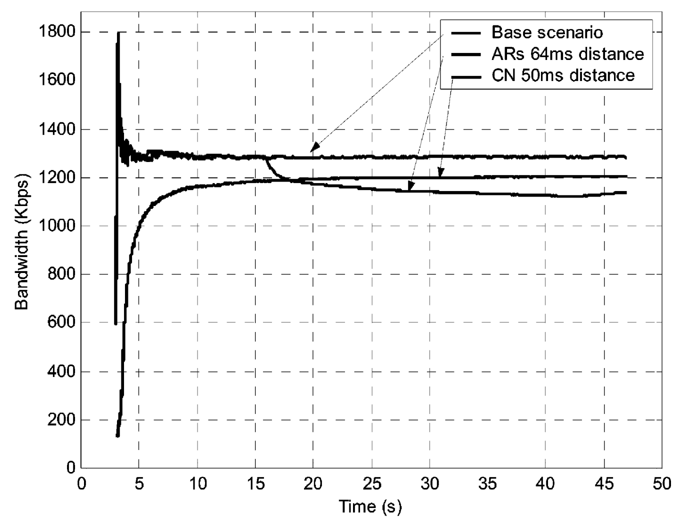

Fig. 8. Mean bandwidth of the TCP flow.

is seamless to the applications, since there were no packets lost (the TCP traffic did not decrease as it would in case of loss) and the bandwidth variation is small.

Fig. 8 shows the mean bandwidth used by the TCP flow for three different situations. The figure basically illustrates the TCP behavior. When the CN is very close to the MT, there is an initial peak in the bandwidth used. When the CN is farther away, we can observe the smooth increase in the mean bandwidth. The most relevant situation is when the MT moves to an AR that is 64 ms away from the old AR. In this case, we can observe a decrease in the mean bandwidth, but this solely happens due to the TCP design, that it is not optimized to deal with this type of reliable links. When the MT returns to the original AR (around 43 s), we can observe a slight increase in the mean bandwidth. For the other cases it is not possible to observe any relevant changes during the handovers.

With respect to the latency (Fig. 9) and jitter it is clear to identify the two handovers, since the bicasting process introduces some additional delay. Just after the handover, the bi-

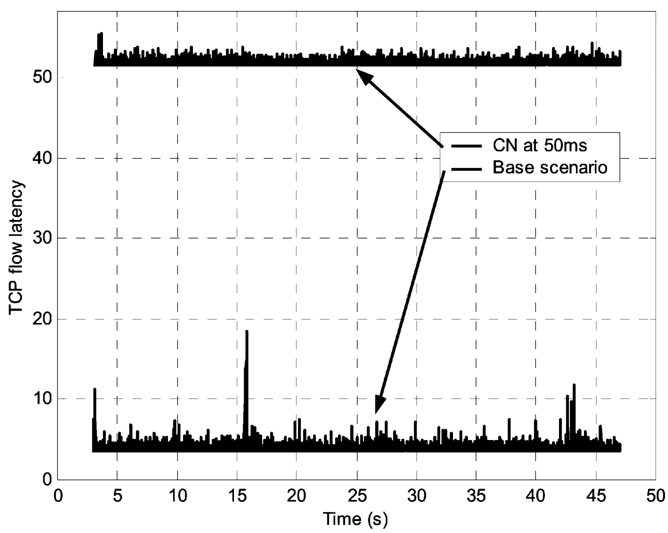

Fig. 9. TCP latency.

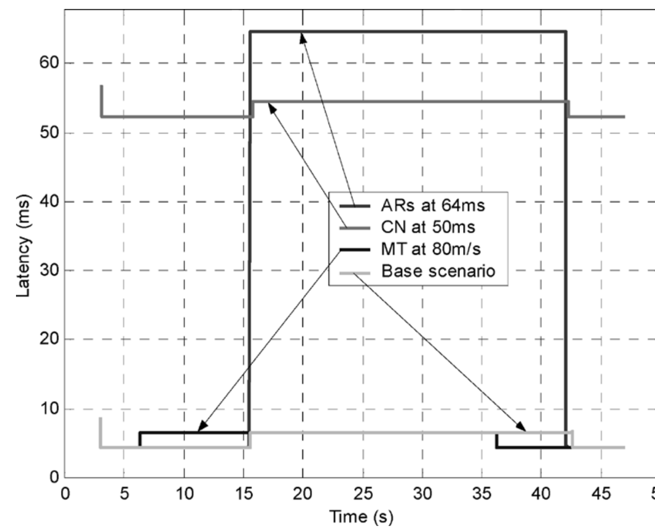

Fig. 11. UDP latency.

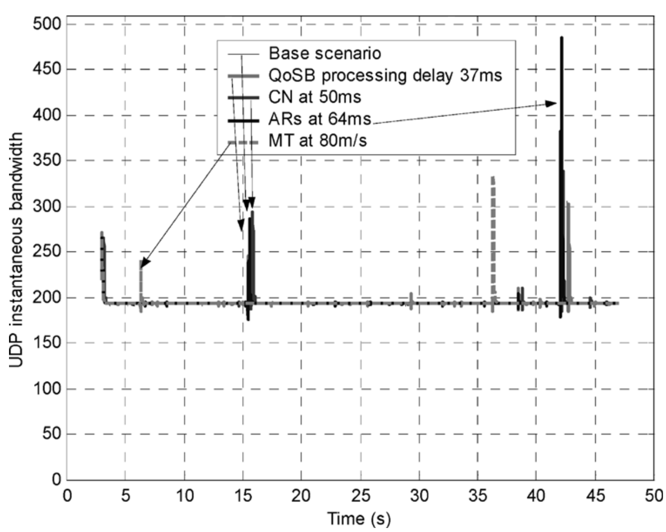

Fig. 10. Instantaneous UDP flow bandwidth.

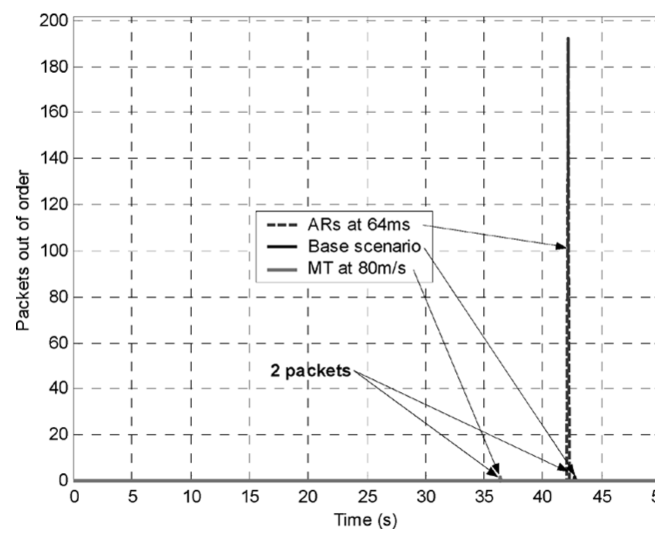

Fig. 12. Packets out of order.

casted packets arrive almost at the same time (sometimes after) as the packets being forwarded directly to the new AR. This justifies the negative jitter values.

Fig. 10 presents the instantaneous bandwidth of the UDP flow. Apart from the initial peak, due to the IPv6 mobility process (the first data packets flow through the home agent, arriving at the same time as the ones sent directly to the MT), it is possible to identify the two handovers at 15,5 and 42,6 s (except for the case when the MT travels at 80 m/s, around 6 and 36 s) as there is an instantaneous peak bandwidth in the graph. This is not due to an increase of the data at the sender side (it is a constant bit rate flow) but due to the bicasting process that causes the MT to receive some duplicate packets. The measurements were done at the MT side and repeated packets were accounted twice. By that reason, we observe those two peaks when the handover occurs.

In terms of latency (Fig. 11), we may observe that, besides the initial peak due to the IPv6 mobility (as explained before), there is an increase, equal to the difference of the electrical distance between the CN and the correspondent AR, when the mobile node travels to the new AR network. When the CN is farther away the latency is obviously higher. It is also possible to identify the earliest handovers when the MT travels faster (80 m/s). With respect to jitter, there exists a positive peak when the mobile node increases its distance to the CN and a negative peak when it comes closer, easily deductible by observing the latency graphs. In this last case, when the mobile node comes from a network farther away from the CN, there are some packets that, due to the bicasting will arrive after some others that were sent directly from the CN to the new AR. Naturally, the peak is higher when the electrical distance between the ARs is also higher. This jitter implies that some of the packets arrive to the MT out of order (Fig. 12). For some real time applications, packets out of order may be considered lost packets, which is a negative effect of the mobility. The amount of packets out of order is also higher when the ARs are more distant from each other.

The final comment is to the scenario where the QoSB processing delay was changed. The processing delay of the QoSB was varied from 3 to 37 ms, and the results, as expected, shown that the only effect is the proportional change in the time that the

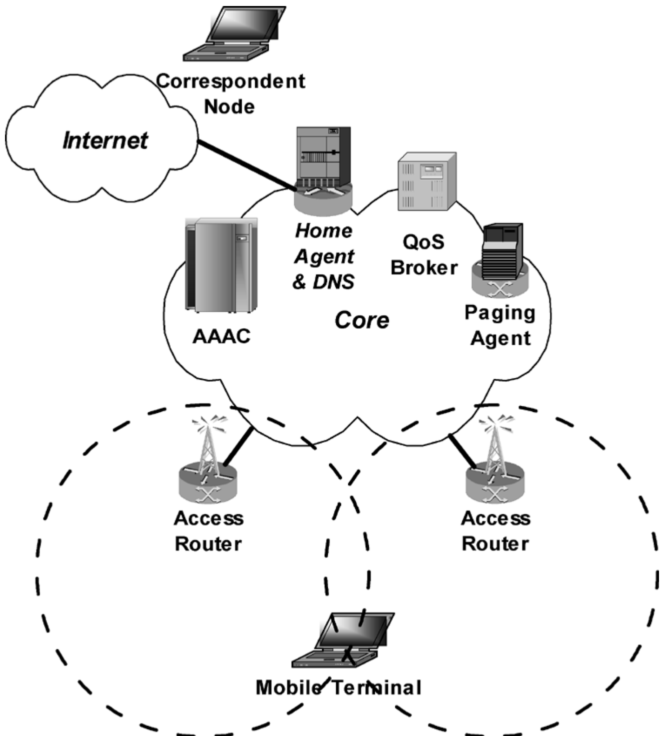

Fig. 13. Testbed scenario.

TABLE II
AUTHENTICATION TIMINGS

Message Nº	Registration Δt (s)
1	0
2	0.169
3	0.209
4	0.224
5	0.800

TABLE III
SERVICE AUTHORIZATION TIMINGS

Message Nº	Service Authorisation Δt (s)
5	0
6	0.754
7	0.767
8	3.494

TABLE IV
HANDOVER SIGNALING TIMINGS

Message Nº	Handover message Δt (s)		Message Nº	Handover message Δt (s)
2	0		7	1,087
3	1,006		8	1,114
4a	1,029		9	1,126
4b	1,065		10	1,144
6	1,086		12	1,257

handover preparation takes. From the values simulated (from 3 to 37 ms), which are reasonable and expected values of a real-QoSB implementation, we observed that there were no impacts on the architecture performance. These results, since they include processing delays associated with real implementation performance, mean that the presence of a QoSB in this architecture does not have a major impact on the performance, while the benefits of having it are very important from the network operator point of view.

E. Critical Scenarios Testbed Evaluation

The architecture presented here was also subject to a real testbed evaluation. The focus of this evaluation was on the signaling performance. As identified before, the three most critical situations are the registration, the service authorization and the most important, the handover.

Although our results were achieved in low load scenarios and accordingly slightly higher values may be obtained in real load scenarios, is should be noticed that most of the delay components are related with hardware and CPU: the field trial implementations were done using low-end PCs running the Linux operating system, that although good to achieve fast implementation results, are not optimal in terms of performance. Thus, performance delays do not relate with the proposed architecture and may be easily reduced with professional equipment.

The testbed scenario, very similar to the simulation scenario presented before, is depicted in Fig. 13.

1) Registration and Service Authorization: Table II presents the signaling timings obtained in the testbed for in typical authorization case. The messages numbers follow the conventions on Fig. 4.

This results show that the timings involved are similar to the ones we may experience in today's mobile networks (the time elapsed between the moment the user tries to join to a network and the instant it gets access).

Table III presents the signaling timings for the service authorization. Once more, the messages are numbered according with Fig. 4.

These results show that the service authorization takes about 0.8 s (message 5 to 7). Since this happens at service start, we may consider this value as negligible from the user experience point of view. Message 8 is the next application packet sent by the MT, which could be sent just after message 7. In the case of the applications used for this measurement, it takes some more time, but it is not architecture dependent.

2) Handover: Table IV presents the handover signaling timings measured on the testbed. Since only one QoSB was used, message 5 on Fig. 5 does not exist (it is QoSB internal).

The results obtained show that the handover preparation phase takes about 140 ms (from message 3 to 10). If we consider a MT speed of 250 km/h, this means that the handover is prepared in less than 10 m, which is a perfectly reasonable value for cell overlapping (usually, several dozens of meters). The complete handover takes about 250 ms (message 3 to 12), which is also a very comfortable value to be used in real networks.

The measurements performed during the fast handover procedure (with the make before break and bicasting mechanism) show that, using current UDP-based applications (typical for audio and video) the packet loss is null most of the times. At the most, one or two packets may be lost and that loss only happens due to the Linux driver's availability. Other measures were performed to evaluate the impact of the QoS components on the overall latency. This impact was the introduction of an increase of 8 ms in the handover time. Since during this time the MT is still attached to the old AR, and that the QoS component is essential to differentiate clients and services, we consider this value an excellent compromise between latency and added value.

F. Human-Based Tests

Given the delays measured in our trial network, we experimented multiple applications on this network, and evaluated user perception on network behavior. The applications run were video telephony (with netmeeting), video streaming (with VideoLan), data transfer (with mgen, ftp, and http) and multiuser gaming (Quake2). Although some of the applications are hard to test while moving (e.g., playing Quake), the overall user perception on the network supported our "seamless" objectives, and the users were not able to perceive any problem due to the behavior of the network. These good results were present in several audits and demonstrations.

VIII. Conclusion

We presented a Mobile IPv6-based overlay network architecture for heterogeneous networks including wired and wireless access technologies. This architecture is a first approach to 4G systems relying on the IPv6 protocol, and replaces most technology-dependent protocols by IP-oriented approaches.

This paper provides an evaluation of the architecture, showing how proposals under development in the IETF were reused and improved in order to achieve this goal, and analyzed their final performance in a testbed. The key problems analyzed in this paper are related with efficient handover solutions across different cells, including an integrated approach to the three traditionally often separated fields of QoS, AAAC, and mobility management. The architecture focuses on the use of Mobile IPv6 and its fast handover optimizations, a DiffServ-based QoS transport infrastructure managed by QoS broker, and a coherent AAA management structure, all embedded with integrated signaling. An innovative IP paging concept was also incorporated.

The simulations performed confirm the correct design of this architecture as a basis for 4G environments. It was interesting to verify that even when a mobile terminal is traveling at 80 m/s, there is no major impact on the architecture performance. The tests conducted in a controlled trial showed excellent results with typical multimedia applications with response times, similar to current cellular networks and well below human perception.

The architecture identifies similar elements across all access technologies, but maintains enough flexibility to support optimizations for the physical layers. This architecture is conceptually flexible and open, providing clear separation between the technology and administrative domains, while keeping the capability of providing specific services to specific users. This approach facilitates the deployment of multiple service provision models, as it decouples the notion of service (associated with the user contract) from the network management tasks. This open architecture represents a step toward IP-dominated cellular networks, and will allow easy and simple integration of other future access technologies (as, e.g., UWB [23]), since it uses only IP-based protocols for mobility, management, QoS provisioning, and AAAC tasks.

In summary, this network seems to provide a potential path to simple, flexible, architectures able to support IP-based multimedia service provision in future 4G networks.

Acknowledgment

The authors would like to acknowledge the European Commission support and the work developed by all partners in the Moby Dick consortium.

References

[1] "3GPP TS 23.060 v3.3.1 General Packet Radio Service (GPRS)," Service Description, Stage 2, 2000. Release 1999.
[2] H. Kaaranen *et al.*, *UMTS Networks: Architecture, Mobility, and Services*. New York: Wiley, 2001.
[3] R. Koodli, "Fast handovers for Mobile IPv6," IETF, RFC 4068, Jul. 2005.
[4] C. Beaujean, N. Chaher, V. Marques, R. L. Aguiar, C. García, J. I. Moreno, M. Wetterwald, and T. Ziegler, "Implementation and evaluation of an end-to-end IP QoS architecture for networks beyond 3rd generation," in *IST Mobile Summit 2003*, Aveiro, Portugal, Jun. 2003, pp. 221–226.
[5] "3GPP TS 25.331 V3.12.0 (2002–09), Third-Generation Partnership Project," Tech. Spec. Group Radio Access Network; Radio resource control (RRC); protocol specification. Release 1999.
[6] D. Black *et al.*, "An architecture for differentiated services," Internet Engineering Task Force (IETF), RFC 2475, 1998.
[7] C. de Laat, G. Gross, L. Gommans, J. Vollbrecht, and D. Spence, "Generic AAA architecture," IETF, Experimental RFC 2903, 2000.
[8] D. Johnson, C. Perkins, and J. Arkko, "Mobility support in IPv6," IETF, RFC 3775, 2004.
[9] C. E. Perkins and K-Y. Wang, "Optimized smooth handoffs in mobile IP," in *Proc. IEEE Symp. Comput. Commun.*, 1999, pp. 340–346.
[10] A. T. Campbell *et al.*, "Comparison of IP micro-mobility protocols," *IEEE Wireless Commun. Mag.*, vol. 9, pp. 72–82, Feb. 2002.
[11] H. Soliman, C. Castelluccia, K. El-Malki, and L. Bellier, *Hierarchical Mobile IPv6 Mobility Management*, Dec. 2004. Internet Draft, Work in Progress.
[12] N. Montavont and T. Noel, "Handover management for mobile nodes in IPv6 networks," *IEEE Wireless Commun. Mag.*, vol. 40, no. 8, pp. 38–43, Aug. 2002.
[13] R. Hsieh, Z.-G. Zhou, and A. Seneviratne, "S-MIP: A seamless handoff architecture for mobile IP," in *Proc. INFOCOM*, San Francisco, CA, 2003, pp. 1774–1784.
[14] X. Pérez-Costa, M. Torrent-Moreno, and H. Hartenstein, "A performance comparison of mobile IPv6, hierarchical mobile IPv6, fast handovers for mobile IPv6 and their combination," *ACM SIGMOBILE Mobile Comput. Commun. Rev. (MC2R)*, vol. 7, no. 4, pp. 5–19, Oct. 2003.
[15] X. Pérez-Costa, R. Schmitz, H. Hartenstein, and M. Liebsch, "A MIPv6, FMIPv6 and HMIPv6 handover latency study: Analytical approach," in *Proc. IST Mobile Wireless Telecommun. Summit*, Thessaloniki, Greece, Jun. 17–19, 2002, pp. 100–105.
[16] M. Liebsch and B. Lamparter, "A generic IP paging architecture and protocol," in *Proc. Eur. Wireless Conf.*, 2004, pp. 114–120.
[17] R. Braden, D. Clark, and S. Shenker, "Integrated services in the Internet architecture: An overview," IETF, RFC 1633, 1994.
[18] V. Marques *et al.*, "A simple QoS service provision framework for beyond 3rd generation scenarios," in *Proc. 10th Int. Conf. Telecommun.*, Papeete, French Polynesia, Feb. 2003, pp. 1475–1481.
[19] D. Durham, Ed., "The COPS (common open policy service) protocol," IETF, RFC 2748, 2000.
[20] N. Brownlee, C. Mills, and G. Ruth, "Traffic flow measurement: Architecture," IETF, RFC 2722, 1999.
[21] V. Marques *et al.*, "Second test and evaluation report," IST-2000–25 394 Project Moby Dick, Deliverable 503, 2003.
[22] Network Simulator (ns), Version 2 [Online]. Available: http://www.isi.edu/nsnam/ns
[23] M. Z. Win and R. A. Scholtz, "Impulse radio: How it works," *IEEE Commun. Lett.*, vol. 2, no. 2, pp. 36–38, Feb. 1998.

Victor Marques received the Licenciatura and M.Sc. degrees in electronics and telecommunications and the Ph.D. degree in electrical engineering from the Universidade de Aveiro, Aveiro, Portugal, in 1994, 1997, and 2005, respectively.

He has worked as a Researcher at the Instituto de Telecomunicações, Aveiro, Portugal, and on April 2001, he started working with Portugal Telecom Inovação. He is coauthor of several national and international publications, both in conferences and magazines. He has been involved in several international research projects, namely, RACE, ACTS, IST, and EURESCOM. His main interests are related to the provisioning of quality-of-service (QoS) in multiservice, heterogeneous networks.

Xavier Pérez Costa received the Telecommunications Engineering and Telematic Engineering Ph.D. degrees from the Polytechnic University of Catalonia (UPC), Catalonia, Barcelona, in 2000 and 2005, respectively.

He did his diploma thesis in 2000 at NEC Laboratories, Heidelberg, Germany, in the area of quality-of-service (QoS) for wireless LANs and was hired after that to focus on IP-based mobility management issues in the framework of the European project Moby Dick. Since then he works at NEC and is currently a Project Manager leading a team which evaluates and configures wireless LAN QoS and power saving mechanisms for NEC's 3G/WLAN mobile terminals. He is the author of several technical papers and patent applications in this area and participates in the 802.11 standardization effort, as well as in the WiFi certification program. His general research interests include, among others, wireless technologies, mobile networking, and theoretical computer science.

Rui L. Aguiar (S'88–M'95) received the Licenciatura degree in electronics and telecommunications engineering, the M.Sc. degree in microelectronics for telecommunications, and the Ph.D. degree in electrical engineering from the University of Aveiro, Aveiro, Portugal, in 1990, 1995, and 2001, respectively.

He is currently an Auxiliary Professor in the Department of Electronics and Telecommunications, University of Aveiro and a Researcher at the Institute of Telecommunications. He is currently leading research activities in heterogeneous networking at the Institute of Telecommunications. His current research interests are centered on the implementation of advanced wireless networks, systems, and circuits, with special emphasis on high bit rates and quality-of-service (QoS) aspects, areas in which he has more than 100 published papers.

Marco Liebsch received the Diploma degree in communication electronic engineering from the University of Applied Sciences, Mannheim, Germany, in 1997 and the M.Sc. degree from the University of Central England Birmingham, Birmingham, U.K., in February 1999.

Until August 1997, he extended the work done on his diploma thesis within the company Freudenberg Weinheim, Germany, and worked on the development of electronic boards for measurement of ferromagnetic quality parameters and for feedforward control. His M.Sc. thesis on "The New Generation Internet Protocol (IPv6) Within an ATM/ADSL-Based Access Network Architecture" was carried out at NEC Network Laboratories, NEC Europe, Ltd., Heidelberg, Germany. Then, he joined NEC Network Laboratories as a Research Staff Member, first working on IP telephony protocols within the EU Ten-Telecom TIPHON project (TTTnet) 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.

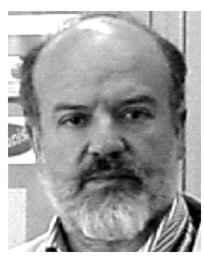

A. Manuel de Oliveira Duarte is a Full Professor with the Department of Electronics and Telecommunications, University of Aveiro, Aveiro, Portugal, which he joined in 1978. In 1988, created the Broadband Systems Group of the University of Aveiro which, over the years, diversified into a series of interrelated research teams, acting in the areas of broadband technologies, optical wireless networks, teletraffic, organization and structure of telecommunications networks and services, and economic and social aspects of telecommunications. More recently, he is in charge of an educational and vocational training program (Programa Aveiro-Norte, www.aveiro-norte.ua.pt) at the University of Aveiro focusing in the areas of industrial design, production technologies, and management. Information and communication technologies are widely used in the context of this program as enablers of a learning environment available to a wide variety of users (undergraduates and postgraduates, lifelong students/trainees, teachers/trainers, etc.).