Ingeniero de Telecomunicación por la Universitat Politécnica de Catalunya desde enero de 2003. Durante sus estudios trabajó a media jornada en el departamento de redes de la consultoría Nextret en Barcelona (Nov-99, Mar-01) y realizó una *internship* en los laboratorios de British Telecom en Inglaterra, donde colaboró con el "Personalisation Server Team" en el diseño e implementación de un sistema de conexión entre agentes móviles con el protocolo SIP (Abr-01, Nov-01). Desde enero hasta octubre de 2002 formó parte del "Mobility Group" de los laboratorios de redes de NEC en Alemania donde, participando en el proyecto europeo de la IST "MobyDick", llevó a cabo un estudio para la mejora del protocolo de movilidad MIPv6 como su proyecto fin de carrera, el cual fue galardonado con el premio ALCATEL al "Mejor Proyecto Fin de Carrera en Movilidad y Multimedia en las Telecomunicaciones" por el COIT. En enero de 2003 se une al equipo de investigación de los laboratorios de telemática de DaimlerChrysler, en Palo Alto, California, donde colabora en el desarrollo de la tecnología DSRC de comunicación entre vehículos para la mejora de la seguridad vial. En enero de 2004 vuelve a Alemania ingresando en el "Institut für Telematik" de la Universidad de Karlsruhe, para continuar con la investigación y desarrollo de nuevos sistemas de comunicación entre vehículos en cooperación con las principales empresas automovilísticas del país.

Nombre:
Marc Torrent Moreno

Nombre del Tutor del Proyecto:
Xavier Pérez Costa

Entidad Investigadora en la que se realizó el Proyecto:
NEC Network Labs Europe

Fecha de lectura:
31 de enero de 2003

Calificación:
Matrícula de Honor

Publicaciones y méritos:

Este trabajo es parte del proyecto de la IST *Moby Dick*, "Mobility and Differentiated Services in a Future IP Network," IST-2000-25394, 2000.

Los resultados obtenidos dieron lugar a tres publicaciones en conferencias internacionales y en el *journal* 'ACM Mobile Computing and Communications Review':

- "A Performance Study of Hierarchical Mobile IPv6 from a System Perspective", X. Pérez-Costa and M. Torrent-Moreno. IEEE International Conference on Communications (ICC), May 2003.
- "A Simulation Study on the Performance of Hierarchical Mobile IPv6", X. Pérez-Costa, M. Torrent-Moreno and H. Hartenstein. International Teletraffic Congress (ITC). August 2003.
- "A Performance Study of Fast Handovers for Mobile IPv6", M. Torrent-Moreno, X. Pérez-Costa and Sebastià Sallent Ribes. IEEE Conference on Local Computer Networks (LCN). October 2003.
- "A Performance Comparison of Mobile IPv6, Hierarchical Mobile IPv6, Fast Handovers for Mobile IPv6 and their Combination", X. Pérez-Costa, M.Torrent-Moreno and H. Hartenstein. ACM Mobile Computing and Communications Review, Nov. 2003

La publicación "A Performance Study of Fast Handovers for Mobile IPv6" aparece como artículo de investigación recomendado en la asignatura "UNSA: DEA Réseaux et Systèmes

Distribués" del año académico 2003-2004 en la universidad de Ingeniería EURECOM, Sophia Antipolis. (http://www.eurecom.fr/~filali/teaching/dea-rsd/)

A Performance Study of Mobile IPv6, Hierarchical MIPv6, Fast Handovers for MIPv6 and their Combination

I.Introduction and Motivation

The fast Internet evolution together with the enormous growth in the number of users of wireless technologies has resulted in a strong convergence trend towards the usage of IP as the common network protocol for both, fixed and mobile networks. Future *All-IP* networks will allow users to maintain service continuity while moving within a multi-access enviroment.
The IETF working group in Mobile IP is proposing Mobile IPv4 (MIPv4) [3] and Mobile IPv6 (MIPv6) [4] as the main protocols for supporting *IP mobility*. Various enhancements to the MIPv6 base protocol have been already proposed since it is believed that in certain cases Mobile IP could result in a poor performance. For environments where the mobile nodes could change its point-of-attachment frequently and the standard Mobile IP protocol could result in a high signaling load as well as high handoff latency and packet losses, micro-mobility protocols, as they are commonly referred to, have been proposed [5]. Hierarchical Mobile IPv6 [20] is the current IETF IPv6 micro-mobility proposal. Additionally, for applications that could suffer from long interruption times due to handoffs, Fast Handovers for Mobile IPv6 has been designed.
From a network operator's point of view as well as from a technology provider view though, one should consider which is the level of improvement that can be expected from implementing these proposals on top of MIPv6. Note that their performance can be affected by various situations or channel conditions, degrading it in different manner. Moreover, it is necessary be aware of the tradeoffs to be faced between achieved improvement and the implementation cost.
This research work investigates the impact of various parameters on the overall performance as experienced by a mobile node of a Mobile IPv6-based wireless access network and compares the performance obtained by the proposed enhancements, i.e., Hierarchical Mobile IPv6 (HMIPv6), Fast Handovers for Mobile IPv6 (FMIPv6), and our proposed combination of both (H+F MIPv6), with the performance of the MIPv6 base protocol.
We are primarily interested in quantifying the degradation of quality of service a mobile user perceives during a handoff when receiving a data stream (e.g., video or voice over IP) and the signaling load costs associated with Mobile IPv6 and its enhancements. More specifically, we are interested in performance metrics like handoff latency, packet loss rate, obtained bandwidth per station and signaling load. Moreover, the impact of different traffic sources is studied: CBR, video, VoIP and TCP transfers. The scenario chosen for this study resembles a 'building block' of a potential wireless LAN 'hot spot' deployment, as one of the possible wireless access networks in an *All-IP* network. It comprises four access routers and up to 50 mobile nodes that communicate in accordance with the IEEE 802.11 wireless LAN standard. We study the performance metrics as observed by one single mobile node that either moves deterministically or randomly while the other mobile nodes move randomly all the time providing realistic 'interference' with respect to the observed mobile node. We consider the impact of different parameters like number of mobile nodes, handoff rate of the observed MN, number of correspondent nodes, wired link delay, and specific protocol options over the various performance metrics. Due to the complexity and broadness of the required study, simulation was chosen as the most suitable analysis method.
As simulation tool we used the network simulator *ns-2*, which we improved and extended not only with the three new proposals, HMIPv6, FMIPv6 and H+F MIPv6, but also with some MIPv6 functionalities that were not properly implemented or just missing, e.g. Neighbor Discovery. I invested about 4 months into the implementation of the new modules at NEC Network Labs Europe.
In contrast to the related literature, we perform a detailed study of Mobile IPv6, Hierarchical Mobile IPv6, Fast Handovers for Mobile IPv6 and their combination focusing not only on handoff latency but on a complete picture of the overall performance taking into account a variety of performance metrics as well as impacting factors. Moreover, while previous analysis usually studied a single mobile node without the interference of others, our work considers a more realistic scenario with up to 50 mobile nodes and random movement patterns. Some simulation results give insights not easily gained without performing simulations. Our goal is not

to determine which protocol performs 'best' but to assess the performance that can be expected for each protocol, broaden our knowledge of the reasons that influence the difference in the performance and help in the design decision of which is the best suited protocol for a specific scenario. This study provides a deep understanding of the overall performance of the various protocols and supports the design process of a Mobile IPv6-based network when a decision of whether it is appropriate to implement any of the proposed Mobile IPv6 enhancements has to be made.
The rest of this summary is organized as follows. Section II describes the basics of MIPv6, HMIPv6, FMIPv6 and our combined H+F MIPv6 approach. In Section III we describe the simulation setup. Performance aspects subject of interest are given in Section IV. Simulation results are provided and discussed in Section V. Finally, Section VI summarizes the results and concludes the study.

II.Mobile IPv6

Mobile IP supports mobility of IP hosts by allowing them to make use of (at least) two IP addresses: a home address that represents the fixed address of the node and a care-of address (CoA) that changes with the IP subnet the mobile node is currently attached to. Clearly, an entity is needed that maps a home address to the corresponding currently valid CoA.
In Mobile IPv4 [8] these mappings are exclusively handled by 'home agents' (HA). A correspondent node (CN) that wants to send packets to a mobile node (MN) will send the packets to the MN's home address. In the MN's home network these packets will be 'intercepted' by the home agent and tunneled, e.g. by IP-in-IP encapsulation [31], either directly to the MN or to a foreign agent to which theMN has a direct link. In MIPv6, home agents no longer exclusively deal with the address mapping, but each CN can have its own 'binding cache' where home address plus care-of address pairs are stored. This enables 'route optimization' compared to the triangle routing via the HA in MIPv4: a CN is able to send packets directly to a MN when the CN has a recent entry for the MN in its corresponding binding cache. When a CN sends a packet directly to a MN, it does not encapsulate the packet as the HA does when receiving a packet from the CN to be forwarded, but makes use of the IPv6 Routing Header Option. When the CN does not have a binding cache entry for the MN, it sends the packet to the MN's home address. The MN's home agent will then forward the packet. The MN, when receiving an encapsulated packet, will inform the corresponding CN about the current CoA.
In order to keep the home address to CoA mappings up-to-date, a mobile node has to signal corresponding changes to its home agent and/or correspondent nodes when performing a handoff to another IP subnet. Since in MIPv6 both, HA and CN, maintain binding caches, a common message format called 'binding updates' is used to inform HA and CNs about changes in the point of attachment. Additionally, since the BUs have associated a certain lifetime, even if theMN does not change its location a BU to its HA and CNs is necessary before the lifetime expires to keep alive the entry in the binding caches. In the rest of the document those BUs will be referred as periodic BUs. Binding updates (BU) can be acknowledged by BU Acks (BAck).
In contrast to MIPv4, where signaling is done using UDP, Mobile IPv6 signaling is done in extension headers that can also be piggybacked on 'regular' packets. To acquire a CoA in Mobile IPv6, a mobile node can build on IPv6 stateless and stateful auto-configuration methods. The stateless autoconfiguration mechanism is not available in IPv4. In our work, we assume stateless auto-configuration for all tests since with this mechanism it is not necessary to contact any entity to obtain a new CoA, reducing the handoff process duration.

II.A. Neighbor Discovery

Neighbor Discovery [14] is one of the main differences when comparing IPv4 and IPv6. It is used by nodes to resolve link-layer addresses and keep track of the reachability of their neighbors. Hosts use it as well to locate routers on their link. The main difference is the IPv6 way of learning MAC addresses and the Neighbor Cache, previously ARP Cache, which can be set in five different states: Incomplete, Reachable, Stale, Delay and Probe.
A MN, when performing a handover, has to learn the Access Router's (AR) MAC address before being able to inform about the new point of attachment via the BUs. In IPv4 a MN runs the ARP process and has to wait until its completion, delaying thus the BUs transmission. On the other

hand, the IPv6 Neighbor Discovery protocol optimizes this process obtaining the AR's MAC address from the Router Advertisement. This results in the MN being able to send the BU without any delay after a handover and running the neighbor unreachability detection process in parallel. However, in IPv4, after the ARP process is completed, MAC addresses on both sides are obtained. This is not the case for IPv6 where the AR that has a packet to transmit to the MN must run the address resolution process to obtain the MN's MAC address. In fact, in the IPv6 case, when a MN learns a node's MAC address in a different way than the usual Request-Reply exchange or when it wants to send a packet after some time without using the entry, the neighbor unreachability detection has to be launched to resolve the MAC address, but this is a one way process (only one address is resolved). Note that in both cases, addresses will be resolved in parallel while sending packets, no delay is added. Additionally, some channel utilization can be saved if confirmation of reachability is received from upper layers.

II.B. Fast Handovers for Mobile IPv6

To reduce the service degradation that a mobile node could suffer due to a change in its point of attachment *Fast Handovers for Mobile IPv6* has been proposed [6]. During the IETF discussions regarding this proposal two different mechanisms have been described: anticipated and tunnel-based handover. Tunnel-based handover relies on link layer triggers to potentially obtain better results than Anticipated Handover, introducing though a link layer dependence that could make the solution unfeasible for some link layer technologies. In principle, a link layer independent solution would be a more desirable solution. Therefore, we have focused on the performance study of the *Anticipated Handover* proposal, which is solely based on network layer information.

Anticipated Handover proposes a 'make-beforebreak' approach. When a MN has information about the next point of attachment to which the MN will move, e.g., via reception of a Router Advertisement from a new AR (nAR), it sends a Router Solicitation for Proxy (RtSolPr) to the old AR (oAR) with an identifier of the point of attachment to which it wants to move. Once the oAR receives information that a MN wants to move to a nAR, it constructs a nCoA based on the MN's interface ID and the nAR's subnet prefix. It then sends a Proxy Router Advertisement (PrRtAdv) to the MN containing the proposed nCoA and the nAR's IP address and link layer Address. At the same time, the oAR sends a Handover Initiate (HI) message to the nAR, indicating the MN's oCoA and the proposed nCoA.

Upon receipt of the HI message, the nAR first establishes whether there is already an active Neighbor Cache entry for the proposed nCoA. If the nCoA is accepted by the nAR, the nAR adds it to the Neighbor Cache for a short time period so it can defend it. The nAR then responds with a Handover Acknowledge (HAck), indicating that the proposed nCoA is valid. Upon receipt of the HAck the oAR is prepared to forward packets for the MN to the nCoA. As soon as the MN received confirmation of a pending network layer handover through the PrRtAdv and has a nCoA, it sends a Fast Binding Update (F-BU) to oAR, as the last message before the link layer handover is executed.

On receipt and validation of the F-BU, the oAR responds with a Fast Binding Acknowledgment (FBAck), destined to the nCoA. The oAR waits for a FBU from the MN before actually forwarding packets. On receipt of the F-BU, the oAR forms a temporary tunnel for the lifetime specified in the F-BAck, and the F-BAck is sent through the tunnel to the MN on the new link. When the MN arrives to the nAR and its link layer connection is ready for network layer traffic, it sends a Fast Neighbor Advertisement (F-NA) to initiate the flow of packets that may be waiting for it. The nAR will deliver packets to the MN as soon as it receives an indication that the MN is already attached to it, usually receiving a F-NA from the mobile node. The oAR is responsible for forwarding any packets that arrive for the MN under its oCoA after the MN has moved. Once the fast handoff process is completed, the MNwill follow the MIPv6 normal procedure of informing the HA and correspondent nodes about its new location.

II.C Hierarchical Mobile IPv6

Hierarchical Mobile IPv6 (HMIPv6) is a localized mobility management proposal that aims to reduce the signaling load due to user mobility. The mobility management inside the local domain is handled by a Mobility Anchor Point (MAP). Mobility between separate MAP domains is handled by MIPv6.

The MAP basically acts as a local Home Agent. When a mobile node enters into a new MAP domain it registers with it obtaining a regional care-of address (RCoA). The RCoA is the address that the mobile node will use to inform its Home Agent and correspondent nodes about its current location. Then, the packets will be sent to and intercepted by the MAP, acting as a proxy, and routed inside the domain to the on-link care-of address (LCoA). When a mobile node then performs a handoff between two access points within the same MAP domain only the MAP has to be informed. Note, however that this does not imply any change to the periodic BUs a MN has to sent to HA, CNs and now additionally to the MAP.

HMIPv6 presents the following advantages: it includes a mechanism to reduce the signaling load in case of handoffs within the same domain and may improve handoff performance reducing handoff latency and packet losses since intra-domain handoffs are performed locally. However, since the periodic BUs are not reduced but the ones due to handoffs, the gain depends on the mobility of the mobile nodes.

II.D. Hirearchical Mobile IPv6 plus Fast Handovers for Mobile IPv6

In this section we describe our proposed combination of FMIPv6 and HMIPv6 which was designed to add up the advantages of both and provide additional improvements. In [20] a sketch on how to combine FMIPv6 and HMIPv6 is provided. However, some issues are left open as for example when should the mobile node decide to perform the handoff. The main ideas of our approach and its differences with respect to a simple aggregation of the proposals described in the above sections are as follows.

Our approach is based on two main observations that show that a simple aggregation of HMIPv6 and FMIPv6 would be inefficient. First, consider a MAP placed in an aggregation router above the ARs involved in a handover. The usual fast handover process of forwarding packets from the oAR to the nAR would be inefficient in terms of handover latency since packets would traverse the MAP-oAR link twice and could arrive disordered. On the other hand, if the entity responsible of establishing the redirection prior to the handoff would be the MAP, then this inefficiency would be removed. Therefore, in our approach, as suggested in [20], the entity performing the functionality of the Fast Handover process is the MAP instead of the mobility agent in the old access router.

Second, note that with FMIPv6 the traffic is redirected when the oAR receives the F-BU but in our case if the mobile node would perform the handoff right after sending the F-BU to the MAP, all the packets forwarded to the oCoA, during the period that the F-BU requires to arrive to the MAP, would be lost. Additionally, if the MN would perform the handoff right after sending the F-BU, it would not immediately receive any redirected packet for the same reason, increasing the handoff latency and packet losses. As a solution, we propose to wait as long as possible (until connectivity is lost) for the F-BAck at the old link to start the handover. In this case we assure that when we receive the F-BAck there are no packets lost sent to our oCoA and the ones redirected to our nCoA are buffered, i.e., no packet losses. Additionally, assuming that the packets experience a similar delay in the path between the MAP and the ARs involved in the handoff, the reception of the F-BAck would act as a kind of synchronization packet telling us that new packets are already waiting or about to arrive to the new AR and therefore, the handover latency due to the wired part would be almost removed.

Our approach requires, as in the case of FMIPv6, that the MN has some time from the moment it realizes that a handover should be performed until it is necessary to perform it because of losing connectivity to the current AR. In the cases where this is not possible we apply the same recovery mechanisms as FMIPv6.

Addendum: During the presentation of the thesis a new internet-draft appeared proposing a combination of HMIPv6 and FMIPv6 basically explaining in detail what was indicated in [20] but without the proposed optimization of waiting at the old access router for the F-BAck.

III. Simulation Setup

The studied scenario was designed in order to be large enough to provide realistic results but to be small enough to be handled efficiently within *ns-2*. The chosen scenario, depicted in Figure 1, is composed by the Home Agent and the Correspondent Nodes that are connected via the 'Internet' (modeled by adjusting the link delay *ld*) to a central router (CR). Four access routers

(AR) –each one representing a different IP subnet– are connected via two intermediate routers (IR) to the central router. When Hierarchical MIPv6 is considered, the functionality of the Mobility Anchor Point is placed on the central router and the CR, IRs, and ARs form the micro-mobility domain. At the start of the simulation the mobile nodes are uniformly distributed over the system area.
The access routers have been positioned in a way to provide total coverage to an area of approximately 700x700 square meters considering a transmission range of 250 meters, see Figure 2. The mobile nodes move randomly within the coverage area following the random waypoint mobility model (RWP). This model has been previously used mainly for ad-hoc simulations but it is well suited as well also for our purposes, more details are given in Section V. As wireless medium the 2Mbps Wireless LAN 802.11 DCF [30] provided by *ns-2* [21] is used.
Within the micro-mobility domain each wired connection is modeled as a 5Mbps duplex link with 2ms delay. The 'Internet' connecting the central router and the HA or CNs is modeled also as a 5Mbps duplex link with a default link delay (ld) of 10ms. In the simulations, the *ld* value has been varied to model various 'distances' between the MNs and the HA and CNs.

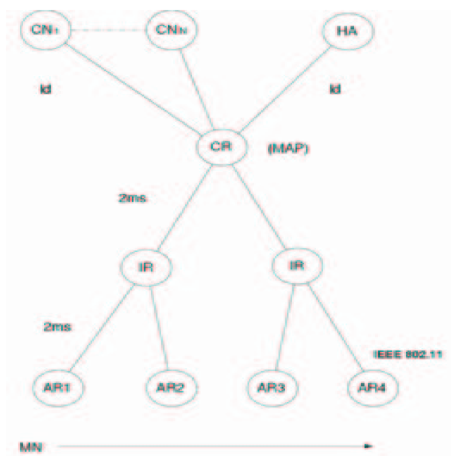

Figure 1: Simulation scenario

Figure 2: Access routers location

While moving within the overlapping area, the mobile nodes are able to send/receive data *only* via the access router that corresponds to their current care of address. Technologies like 802.11 allow the mobile nodes gathering information about the neighboring access routers, but do not allow to receive IP flows at different frequency bands simultaneously from two access routers, except for particular cases like having an additional wireless interface.
In order to simulate a realistic case where aMN will receive packets from the shared AR queue and where a MN will also compete with other MNs and with an AR to access the channel, half of the MNs receive data from the CNs and the other half send data to the CNs. The CNs sending to the MNs introduce delay in the AR queue and the MNs sending to the CNs introduce delay in the wireless link. The study though focuses on the MNs receiving data from the CNs which are the most affected ones by the handoffs since the purpose is to analyze the degradation of the user experienced quality of service due to mobility.
It is important to note the following fact that results from using a shared access: whenever we work close to the maximum throughput of the channel, the MNs that will first experience a reduction in their throughput will be the ones *receiving* from the CNs. The reason is that these stations receive all the packets from the same station, i.e., the AR, sharing the access queue to the wireless channel, while the other MNs *sending* to the CNs do not share their access queue.
In our simulations we study the performance metrics as observed by one single mobile node but affected by other moving mobile nodes. In most of the simulations the observed mobile node follows a deterministic path while all other mobile nodes move randomly. This case allows for full control of the mobility – and handoff rate – of the observed node while the interference of other nodes is still realistic due to their random movements. As a second case we allow the observed mobile node to move randomly, too. By doing this, mobility is less 'controllable' but

random movement effects – like going back and forth between two ARs – can be analyzed. Thus, with both deterministic and random movements of the observed node studied separately, impact of the different parameters over the various protocols can be studied in a clear as well as realistic way.

The first type of sources used in our simulations will be UDP CBR sources. These sources provide constant traffic where no acknowledgments are required. This kind of traffic is usually generated by real-time applications and due to its deterministic characteristics, without recovery mechanisms, eases the protocols study and comparison. Unless otherwise noted, UDP CBR sources are used.

One of the applications expected to be used with MIPv6 is VoIP. We have implemented a VoIP model based on the one provided in [27]. The model assumes silence suppression and models each voice source as an on-off Markov process. The alternating active *on* and silence *off* periods are exponentially distributed with average durations of 1.004s and 1.587s. As recommended by the ITU-T specification for conversational speech [28], an average talk spurt of 38.57% and an average silence period of 61.47% is considered. A rate of 88 kbps[1] in *on* periods and 0 kbps in *off* periods is assumed for a voice source that generates CBR traffic.

As streaming application for real-time video traffic we have used a real H.263 [25] video encoding provided by [26] (film: "Star Trek: First Contact") for a target bit rate of 64 kbps. The obtained frame sizes (in bytes) of the individual encoded video frames are used as input for the *ns-2* real-time video traffic application. Since these traces include only the raw packetized video, additional streaming protocol overhead has been added. As in the case of VoIP sources we consider a 12 byte RTP header plus 8 byte UDP header and plus 40 byte IPv6 header as the streaming protocol overhead.

TCP is the most widely used transport protocol. We simulate endless FTP sources to understand the impact of IP mobility on the congestion control mechanism of TCP.

The simulation code used for the experiments was designed on top of INRIA/Motorola MIPv6 [24] code for *ns-2* [21] implementation. Although *ns-2* is a very powerful network simulator, with plenty of technologies already implemented, the existent code was only able to perform one single mobile node MIPv6 simulations with limited functionalities. We have fixed it in order to simulate complete MIPv6 scenarios and extended it with four main modules: Neighbor Discovery, Hierarchical Mobile IPv6, Fast Handovers for Mobile IPv6 and their combination. The whole functionality described in Section II has been implemented.

IV. Performance metrics

The parameters to be studied are as follows:

Handoff Latency: Handoff latency is defined for a receiving MN as the time that elapses between the last packet received via the old route and the arrival of the first packet along the new route after a handoff.

Packet Loss: Packet loss is defined for a receiving MN as the number of packets lost during the handoff. While one usually assumes that packet losses are directly proportional to latency it will be shown that this is not true in some cases. We have studied separately the packet losses due to the address resolution process, the packet losses in the old access router and the packet losses in the Home Agent.

Signaling Load: The signaling load is defined for MIPv6 and HMIPv6 as the number of BUs and BAcks received during the simulation. Additionally, in the FMIPv6 and H+F MIPv6 case the BUs, BAcks, PrRtAdv, PrRtSol, F-NA, F-BU, F-BAck, HI and HAck signaling messages are also considered.

Bandwidth per Station: We study the probability to obtain the required bandwidth and the corresponding expected variance for CBR and TCP sources for an increasing number of competing stations.

Note that the whole set of performance metrics have been obtained for each scenario but only the most relevant results have been included.

---

[1] Assume 8KHz 8 bits/sample PCM codec was used with 20ms frame per packet. With 12 byte RTP header, 8 byte UDP header and 40 byte IPv6 header, the size of each voice packet is 220 bytes. The bandwidth required will be $(220x8)/20x10^{-3}$ =88kbps

V. Performance evaluation & discussion

With our *ns-2* simulations we study the parameters explained in Section IV for the scenario described in Section III. Unless stated otherwise, we analyze the degradation of the performance metrics from the point of view of a single mobile node that follows a *deterministic* path while all other mobile nodes in the system follow the random waypoint mobility (RWP) model with a maximum speed of 5m/s. The RWP model is well suited to represent movements of mobile users in campus or 'hot spot' scenario at moderate complexity.

To obtain accurate results we have chosen a UDP probing traffic from the CN to our specific mobile node of 250 bytes transmitted at intervals of 10 ms. The other mobile nodes create background traffic sending or receiving data at a rate of 32 kbps.

All simulations have a duration of 125 seconds with a 5 seconds warm-up phase. Each point in the following graphs represent the average of at least 100 simulations. The sample size necessary to achieve a confidence interval of 99% with respect to the average value has been selected as indicated in [23]. This required in some cases to perform up to 1000 simulation runs, e.g., in the 50 mobile nodes or random movement case.

We assume a system where mobile nodes use the IPv6 stateless address auto-configuration feature [18] performing Duplicate Address Detection (DAD) in parallel to avoid the introduction of an additional delay to the handoff process. Note that the delay introduced by DAD would be too time consuming resulting in a noticeable disruption of the service.

V.A. *Impact of number of stations*

We present here the results of the impact of the number of competing stations on the following parameters: handoff latency, packet loss, obtained bandwidth and the fast handoff process probability of success.

The studied MN performs 4 handoffs during a simulation run moving at 10 m/s from center to center of the ARs' coverage areas until it reaches again the starting point. The values represented in the graphs correspond to the analyzed MN.

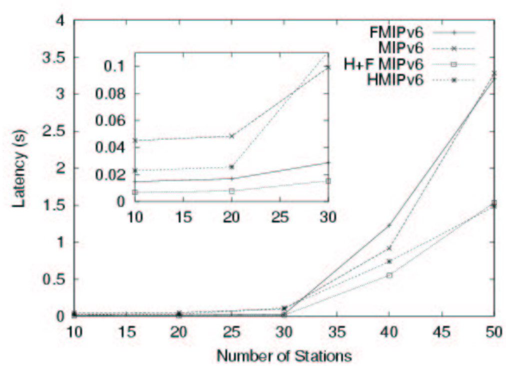

Figure 3: Impact of number of stations on handoff latency

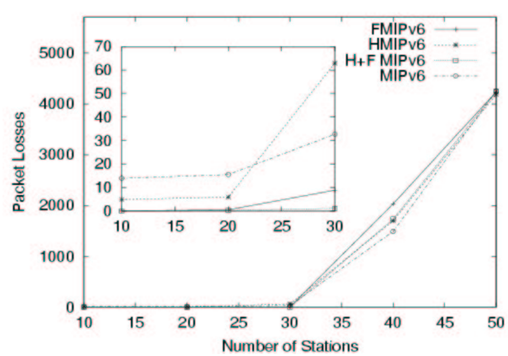

Figure 4: Impact of number of stations on packet losses

Figures 3 and 4 show the increase in handoff latency and packet losses due to an increase in the number of MNs sharing the wireless channel. We can observe that up to 20 MNs the results are as expected considering that for a small number of mobile nodes, e.g. 20 or below, the dominating factor of the handoff latency is the wired delay. HMIPv6 latency outperforms standard MIPv6 one since the wired 'distance' in order to update the entity that forwards packets to the mobile node is always shorter. FMIPv6 outperforms standard HMIPv6, since the MN prepares the handoff in advance and thus, after a handoff, does not have to wait for the oAR to be updated to start receiving packets again. With FMIPv6 packets are redirected by the oAR to the nAR through the wired link and therefore only this delay is noticed. H+F MIPv6 performs better than all the other solutions since, as explained in Section II.D, when the MN receives the F-BAck from the MAP indicating that the handoff should be performed, the re-directed packets are already waiting in the new AR.

An exceptional case can be observed for 30 MNs where MIPv6 shows a slight better performance than HMIPv6. Due to the encapsulation that HMIPv6 always does from the MAP to the current point of attachment we have a higher load on the channel, i.e., 40 additional bytes per packet, and thus HMIPv6 reaches earlier saturation conditions, increasing the wireless delay that now dominates over the wired one. This difference can not be noticed in the H+F MIPv6 case because although we have the same encapsulation problem, the higher load in the channel does not have a direct impact on the handoff performance due to the fast handover mechanism that prepares the handover in advance and re-tries up to three times. However, when the wireless delay becomes very high due to saturation in the channel, e.g., 40-50 stations case, we have again a better performance of HMIPv6 in comparison with MIPv6 due to two reasons. First, in the HMIPv6 case the BU to the MAP is sent right after attaching to the new link while MIPv6 sends a BU to the HA before the one to the CN, i.e., introducing an additional wireless delay. This difference could be removed sending the BU first to the CN and then to the HA. Second, while the BAcks to HA and MAP are mandatory, the BAck to the CN is optional. In our implementation BAcks to CN BUs are not sent to avoid additional overhead and because in case of the BU being lost, the MN will re-send it again when receiving a data packet from the HA instead of directly from the CN. Under high saturation channel conditions the probability of a BU to be lost is higher, therefore, when using standard MIPv6 if a BU to the CN is lost, it is not retransmitted, increasing significantly the latency value. On the other hand, when the BAck from the MAP is not received, the BU will be retransmitted.

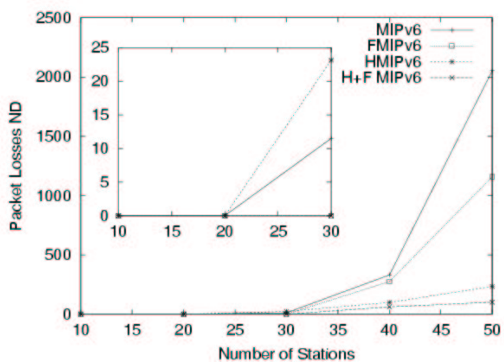

Figure 5: Impact of number of stations on packet losses in the Neighbor Discovery resolution queue

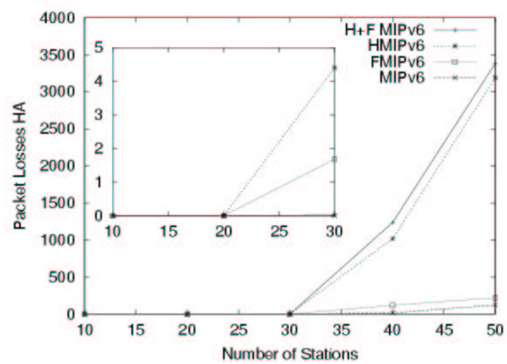

Figure 6:Impact of number of stations on packet losses at the HA

Although the Fast Handover protocol is designed to minimize packet losses and latency during a handoff, we can observe a worse performance with respect to MIPv6 when saturation arises. To understand this behavior a few factors must be considered. In the scenarios with 40 or more MNs the load in the wireless channel is high, resulting in a channel with a long access time and high collision rate. If we take a look at the packets lost at the neighbor discovery resolution queue[2] (ND), Fig.5, we can see that they are higher when FMIPv6 is not used (they double with 50 MN). Those packets, that are dropped in the ND entry queue, are not sent through the wireless channel, which results in a lower channel saturation and, what is more important, a shorter access delay. In the FMIPv6 scenario though, the nAR learns the link layer address of the MN before having to send a packet to it (via the reception of the PrRtSol by the oAR which triggers the HI-HAck handshake) even if the FMIPv6 process has not been successfully performed. Therefore, the AR will send packets through the wireless medium without waiting for the address to be confirmed, once the F-NA has been received, introducing a higher load on the channel.

H+F MIPv6 and HMIPv6 present, under saturation conditions, similar packet losses results since the process to update the MAP and afterwards HA and CN about the new point-of-attachment is the same for both approaches (see Figure 6). They show the worst performance in packets lost at the HA, which is actually a good measure of whether the route updating mechanisms are working properly. Packet are lost at the HA only when the BU lifetime of both, CN and HA, has expired. As it can be seen in the figure, the higher load for 30 or more MNs produces a higher rate of

[2] During the address resolution process only a small amount of packets are buffered for the same destination address, e.g., three in our implementation [15]

packet losses at the HA. Which actually are most of the packet losses experienced in the H+F MIPv6 and HMIPv6 case. The reason is that the MN has to wait for the MAP's BAck to send the BUs to HA and CN, what can take a long time when the wireless channel is highly congested, resulting in the expiration of the BU lifetime (10s in our experiments) of the HA and CN. H+F MIPv6 obtains a slight higher HA packet loss rate due to its additional signaling load (see Section V.B). These higher packet losses in the HA are compensated by lower packet losses due to Neighbor Discovery. Note that if the first signaling message of the fast handover procedure (PrRtSol) arrives at its destination, triggering the HI, the nAR will already have the link layer address before having to forward data packets to the MN, which explains the slight difference between H+F MIPv6 and HMIPv6. Another remarkable aspect of the ND packet losses graph is the big difference in saturation conditions between the protocols that use a hierarchical approach and the others. The ND procedure is triggered by the first packet received in the nAR, the BAck from the MAP. Using a hierarchical approach and under saturation conditions the BAck is not always immediately followed by data packets (because the HA and CN have not been updated on time and packets are being dropped in the HA) providing some additional time to the nAR to resolve the link layer address.

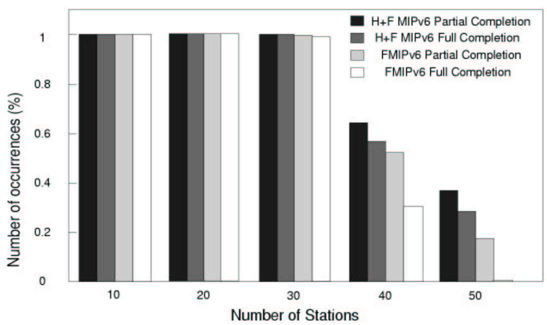

Figure 7: Fast Handover process success histogram

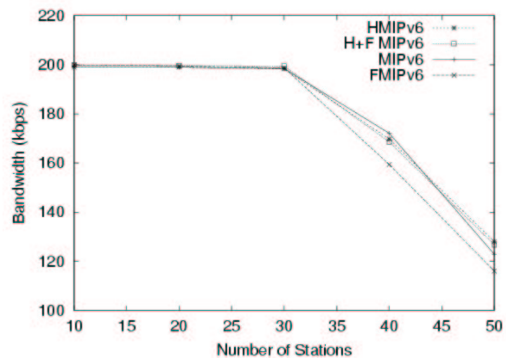

Figure 8: Impact of number of stations on bandwidth obtained by observed MN

Although all the differences (either in ND or HA Packet Losses) described for the congestion case, we can observe that once the saturation level has been reached by all the protocols, if we increase the number of MNs the packet losses tend to converge, since for all cases the wireless channel presents a high collision rate and long channel access time reducing thus, the impact of the differences between the approaches. Figure 7 perfectly shows the saturation of the channel depending on the number of MNs. Up to 30 MNs the wireless channel conditions allows for a proper completion of the fast handoff process. However, for a higher number of MNs the probability of the process success decreases dramatically. We have differentiated between two cases: full completion of the fast handoff process and partial completion, i.e., the redirection of the traffic from the oAR to the nAR has been established. We believe that the latter case is a significant value since it means that the F-BAck packet has been lost but not the previous FMIPv6 corresponding messages, resulting in a smoother handoff compared to MIPv6. H+F MIPv6 presents a better performance than FMIPv6 since most of the packet are lost in the HA reducing the load introduced in the wireless channel compared to FMIPv6.

Figure 8 corresponds to the bandwidth obtained by our specific mobile node. As we can see, the bandwidth correlates almost perfectly the results shown for packet losses. The slight difference between both graphics (in the 40 and 50 MN case) is a consequence of the higher wireless load of the different enhancements. A higher number of data packets sent through the wireless channel and signaling load yields a longer channel access delay and higher collision rate, resulting in a higher number of packets waiting to be sent in the current MN's AR interface queue at the end of the simulation and, therefore, lower bandwidth achieved. As commented above, H+F MIPv6 and HMIPv6 experience a lower load on the wireless channel since most of the packets are lost in the HA.

For the following studies we have focused on the case of 20 MNs since this represents the case with a highest number of MNs in the network where the cannel can still be accessed without experiencing a high degradation in the quality of service due to competing nodes.

V.B. *Impact of handoff rate*

In Section V.A we have shown some of the performance improvements obtained introducing the MIPv6 enhancements. However, as explained in Section II several additional signaling messages have been introduced to achieve those results. A trade-off between additional signaling load and performance improvement has to be considered. In Figure 9 we study the differences in signaling load between MIPv6 and the proposed enhancements for a handoff rate range varying from 0 to 10 handovers per minute for a simulation of 125 seconds.

H+F MIPv6 presents the higher signaling load within the local domain, as expected, since it introduces the HMIPv6 signaling load plus the FMIPv6 signaling load. The next highest signaling load within the local domain belongs to FMIPv6 since, in the event of a handoff, a higher number of signaling messages are required. One of the purposes of HMIPv6 is to keep constant the signaling load outside of the local domain. Figure 9 shows that this goal is achieved by HMIPv6 and H+F MIPv6. In the scenarios where a MAP is placed on the CR and when roaming within the local domain, HA and CNs do not realize any change in the point of attachment and receive only periodic BUs, therefore the signaling load is constant outside the local domain. However, with standard MIPv6 and FMIPv6, when a MN performs a handoff, it must inmediately inform its HA and CNs, and thus, although the periodic BUs are re-scheduled, the total signaling load is increased within and outside the local domain. Note though, that the introduction of a MAP in the system results in a quantitative increase of the signaling load in the local domain, i.e., additional MAP's BU-BAck plus the encapsulation for the BAcks originated by the HA.

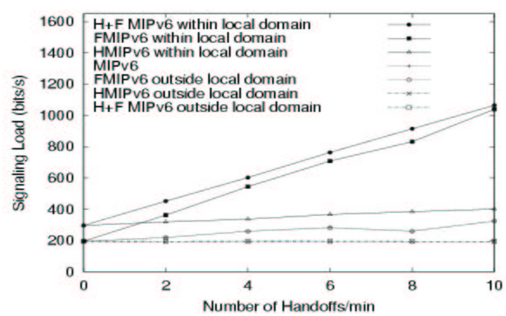

Figure 9: Impact of handoff rate on signaling load

As we can observe, MIPv6 and FMIPv6 introduce the same signaling load outside the local domain since all the additional FMIPv6 signaling is sent only within the local domain. The same case applies to H+F MIPv6 and HMIPv6 that only differ in the signaling behavior within the local domain obtaining thus, the same results outside the local domain.

The signaling load corresponding to standard MIPv6 presents, a priori, a strange behavior having a local minimum for the case of 8 handoffs/min. However, if we recall that for each handoff the MN reschedules the periodic BUs to be sent we realize that if the timer of the periodic BUs is below the time between two consecutive handoffs we will observe the periodic BUs and afterwards the ones due to a handoff. On the other hand, if the time between two consecutive handoffs is below the timer of the periodic BUs, they will be always re-scheduled without being sent during the whole simulation. Thus, in the case of 8 handoffs/min, considering a timer of 10 seconds for the periodic BUs, they are always re-scheduled due to a handoff and never sent, resulting in a reduction of signaling load compared to the previous case.

V.C. *Impact of number of correspondent nodes*

One of the advantages of HMIPv6 is that when performing a local handoff the only entity that has to be informed via a BU is the MAP, which reduces the signaling load. This becomes specially important when the number of correspondent nodes increases, i.e., while the number of BUs to be sent increase linearly with MIPv6 and FMIPv6 remain constant for HMIPv6 and H+F MIPv6. However, HMIPv6 and H+F MIPv6 do not reduce the number of periodic BUs to be sent but increase it by the additional one sent to the MAP. Based on the above comments, a trade-off has to be considered between the number of handoffs performed within periodic BU periods and the number of correspondent nodes. This trade-off was already addressed in [11].

Figure 10 shows the impact of increasing the number of correspondent nodes over the signaling load for the different protocols in the case of a mobile node performing 4 handoffs[3] in 120 seconds. FMIPv6 and H+F MIPv6 perform exactly as MIPv6 and HMIPv6, respectively, concerning to the signaling load sent outside of the local domain since there are no differences in the protocol behavior for the signaling messages sent outside of it. As we can observe, the usage of HMIPv6 or H+F MIPv6 reduces the signaling load outside of the local domain compared to MIPv6 and the difference tends to increase according to larger number of correspondent nodes. The difference though, is not very big since in our scenario the number of handoffs per periodic BU periods is small resulting in a small differentiation of HMIPv6. For a scenario with higher mobility or with larger BU periods the HMIPv6 signaling load reduction would be larger including also the local domain signaling.

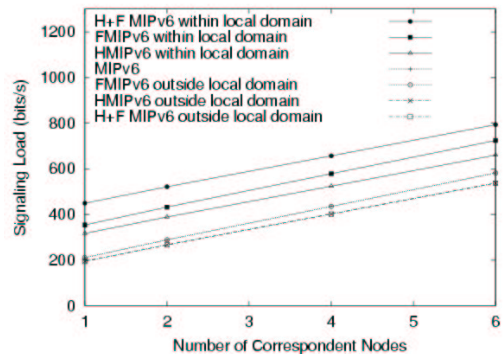

Figure 10: Impact of number of CNs on signaling load

V.D. *Impact of wired link delay*

We have measured the differences in handoff latency and packet losses between MIPv6 and its enhancements when the wired link delay *ld* from the CR to the HA and CN is increased. The different *ld* values model different 'distances' to the HA and CNs.

MIPv6's enhancements reduce the time that elapses between a MN change of point of attachment and the traffic redirection to its nCoA by introducing a new forwarding entity within the local domain, either oAR or MAP, responsible to re-direct the traffic. Thus, the delay experienced by the re-directed traffic does not depend on 'how far' is the MN from its HA and CNs outside of the local domain. On the other hand, with MIPv6, the BUs sent after performing the handover, have to reach the HA and CNs (outside of the local domain) in order to send the traffic to the proper CoA resulting on a direct dependence with the *ld* value.

As we can see in Figure 11 the results are as expected: while an increase in the wired link delay implies an increase in the handoff latency for MIPv6, it does not affect the other proposals' handoff latency.

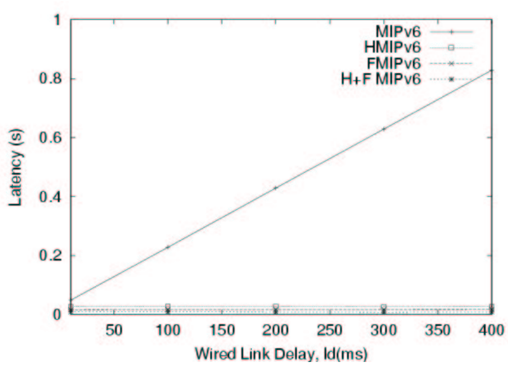

Figure11:Impact of wired link delay on handoff latency

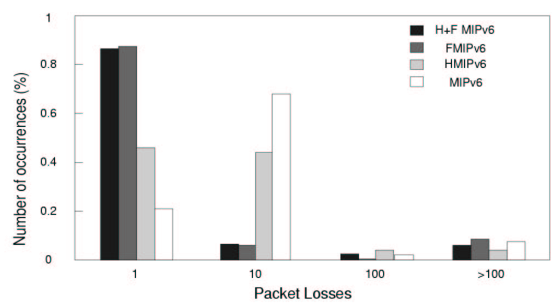

Figure 12:Packet losses histogram considering random movement

V.E. *Impact of random movement*

Mobile users are unaware of overlapping areas where handoff decisions are taken. This section studies whether the differences on the performance metrics observed in previous sections for a mobile node following a deterministic path still hold considering random movement. Note that unexpected movements can have a quite negative effect on the packet losses experienced due to

[3] In [26] a twelve-week trace of a building-wide local-area wireless network was studied. The results presented there that 2 handoffs per minute is a high handoff rate for pedestrian mobile users

back and forth movements around the overlapping areas. This effect could potentially prevail over the protocol enhancements.

Figure 12 shows the histogram of packet losses experienced by the studied mobile node moving randomly in the case of 20 mobile nodes for the four different protocols. The packet losses occurrences have been grouped in lower or equal than 1, 10, 100 and over 100. As we can observe from the figure, the results are consistent with the ones presented in Section V.A. FMIPv6 and H+F MIPv6 show the better packet loss performance keeping for most of the cases values below or equal to 1. HMIPv6 outperforms MIPv6 but without reaching the level of the protocols that include the fast handover approach.

V.F. *Impact of traffic sources*

Until this section we have studied the impact of different parameters over a target station receiving a high constant traffic load (*probe*) in order to obtain results with a significant precision and without the interference of source burstiness (VoIP, Video) or recovery mechanisms (TCP). In this section we repeat the experiment of Section V.A but considering more realistic traffic sources and a simulation scenario where all the MNs send or receive the same type of traffic at the same rate. By doing this, we analyze whether the different performance improvements observed in previous sections are affected by the traffic source type, i.e., whether a user would realize a service improvement or the improvements are 'masked' by the traffic sources characteristics. Specifically, three different types of traffic are studied: VoIP, video and TCP transfers.

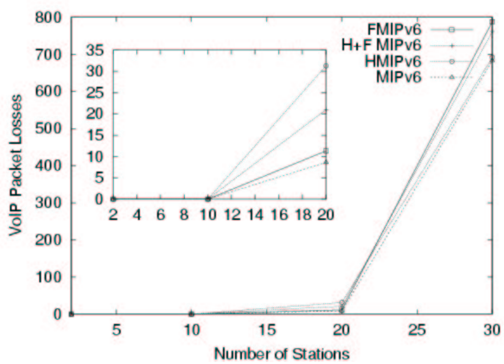

Figure 13: Impact of number of sources on VoIP packet losses by a receiving user

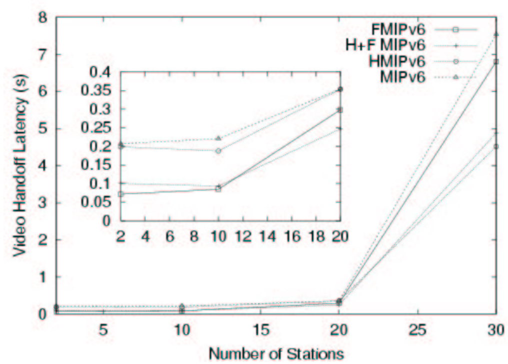

Figure 14: Impact of number of sources on Video handoff latency by a receiving user

As explained in Section III, our VoIP source produces bursty traffic following an on-off Markov process that results in a high variance between packet arrivals. Figure 13 shows the impact of the number of stations over the packet loss rate of VoIP traffic until the congestion level is reached. Since VoIP sources produce a relatively low traffic load (24kbps per source) no packet loss is observed for any of the protocols until the 20 MNs case. In this case, surprisingly, MIPv6 is the protocol that performs best in packet losses terms and HMIPv6 worst. The additional load introduced by the different enhancements in the wireless channel is the reason for this behavior. HMIPv6 is the worst one due to the encapsulation of all packets directed to the MNs performed by the MAP, FMIPv6 performs better since is 'better equipped' to avoid packet losses and H+F MIPv6 is in the middle since is the one producing more overhead but equipped as well with a mechanism to reduce packet losses. When the wireless channel is congested, i.e., 30 MNs case, we observe the same behavior as the one

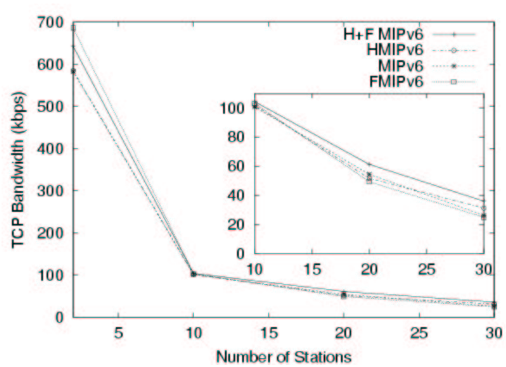

Figure 15: Impact of number of sources on TCP bandwidth obtained by a receiving user

already described in Section V.A. We can conclude that, for a scenario with low rate traffic sources sending small packets (compared to the additional encapsulation header) and in no congestion conditions, the overhead introduced by the different enhancements would result in a worse performance in handoff latency an packet losses terms compared to the baseline Mobile IPv6.

The H.263 video source produces packets of different length at a variable bit rate for a target rate of 64 kbps. We show the impact of the number of stations over the handoff latency. As we can observe in Figure 14, the results are similar to the ones already described in Section V.A, i.e., H+F MIPv6 and FMIPv6 are the ones that perform best in handoff latency terms and MIPv6 is the worst. In this case, in contrast to the VoIP one, the implementation of the Mobile IPv6 enhancements results, as expected, in a better user experienced service since the additional signaling load is less relevant compared to the data traffic load.

Finally, we study whether a regular user downloading a file using TCP would notice any difference in the received service by using one of the different proposals. For a user performing a download, handoff latency or packet loss rate are not relevant performance metrics but the experienced bandwidth during the TCP transfer is of major interest.

Figure 15 shows the differences on the available bandwidth for TCP users depending on the MIPv6 protocol enhancement used. In the figure we can observe the TCP sources adjustment of the sending rate to the available channel capacity when the number of mobile users increases. For a number of mobile nodes below 10, a lower packet loss rate obtained via the enhancements results in users achievement of larger bandwidth. H+F MIPv6 presents better packet losses results than FMIPv6; however, with the latter proposal a larger bandwidth value is obtained. In Section V.A we have shown that there is not a direct relationship between packet losses experienced and obtained bandwidth. The reason is that in the H+F MIPv6 case the MAP encapsulates all the data packets addressed to the mobile nodes, and this overhead reduces the available bandwidth in the wireless channel. The same explanation applies to HMIPv6, where the lower packet loss rate does not result in a significant higher bandwidth compared to MIPv6 because of the packet encapsulation within the local domain.

When the number of mobile nodes increases, the probability of experiencing a collision while trying to access the channel increases, too. This, in turn, triggers the TCP congestion avoidance mechanism more often reducing the packet losses experienced by the MNs and thus, decreasing the bandwidth differences between the proposals. These differences would otherwise be much bigger, as it has been shown is Section V.A, when the users try to get a larger bandwidth than the one actually available in the channel.

As a conclusion, TCP users would also benefit from the implementation of one of the MIPv6 protocol enhancements even though the improvement would be lower than for other types of traffic, e.g., CBR.

VI. Conclusion

Mobile IPv6 represents a key element of future *All- IP* wireless networks to allow a user to freely roam between different wireless systems. In this paper we have provided quantitative results on Mobile IPv6 performance as experienced by a mobile node and on the level of improvement that can be achieved by using the proposed Mobile IPv6 enhancements. The results were achieved through a thorough study via simulation that required to implement Neighbor Discovery, HMIPv6, FMIPv6 and our combination of HMIPv6 and FMIPv6 for *ns-2*.

We performed a ‘stress test’ of the protocols where we studied how handoff latency, packet loss rate, obtained bandwidth and fast handoff process success probability are affected by the number of mobile nodes, i.e., by competition for the wireless medium, or by protocol interactions, e.g., with the Neighbor Discovery process of IPv6. The behavior of the protocols for a general case considering random movements and more realistic traffic sources, i.e., VoIP, video and TCP, were also studied. Finally, the signaling load costs associated to the different proposals compared to the performance improvements obtained were analyzed, considering a broad range of handoff rates and number of correspondent nodes. These factors were shown to have a significant influence over the performance metrics and we indicated the points to be taken into account in a real implementation.

Specifically, we have shown that while some simulation results corroborate the intention of the protocols specifications, other results give insights not easily gained without performing simulations. Some of the key results are that *i)* random movements of the observed mobile node do affect the experienced performance but the improvements with respect to the perceived quality of service when using one of the various protocol enhancements is still clearly noticeable, *ii)* in scenarios where the users produce a low rate with small packets, e.g, VoIP sources, the additional overhead introduced by the proposed enhancements can result in a worse performance than the baseline Mobile IPv6 one, and *iii)* Mobile IPv6 can eventually outperform its proposed enhancements in packet losses terms in saturation conditions due to the higher number of packets discarded directly that lower the load in the wireless channel.
Through this analysis a deep insight on the different overall performance of the various protocols and their causes was acquired. Therefore, the results of this study are twofold. First, we provided quantitative results for the different IETF proposals of the overall performance for a realistic 'hot spot' scenario. Second, we provided the reasoning behind the impact of the different parameters over the performance of the various protocols in saturation and no saturation conditions. This reasoning can applied when other scenarios are analyzed.