Seamless vertical handover of VoIP calls based on SIP Session Border Controllers

Stefano Salsano(1) , Luca Veltri(2), Gianluca Martiniello(1), Andrea Polidoro(1)
(1) University of Rome "Tor Vergata" Dpt. of Electronic Engineering, Rome, Italy
(2) University of Parma, Dpt. of Information Engineering, Parma, Italy

***Abstract*-High data rate wireless networks enables the provisioning of real time services to mobile devices over IP. At the same time, the different nature of the adopted radio technologies (802.11/WiFi, Bluetooth, 2.5G/3G networks) has introduced the problems of supporting seamless communication across these heterogeneous wireless accesses, which is a goal for 4G network. This paper describes a Session Initiation Protocol (SIP) based solution for mobility management aiming to provide seamless voice service in this heterogeneous scenario. The novelty of the solution is that it relies on the so called "Session Border Controllers" which are currently used in commercial SIP telephony solutions to deal with NAT traversal. A prototype of the proposed solution has been implemented in a test bed where the 802.11 and 3G (UMTS) technologies have been used respectively as typical WLAN and cellurar access networks. Measurements results are reported which analyze the performance of the solution in a real world environment, using commercial WiFi and 3G services.**

I. Introduction

Over the past few years network and service providers have shown a special interest in the opportunities of offering voice over IP (VoIP) services in the wireless environment. In fact the pervasive penetration of wireless technologies featuring high data rates has enabled the spreading of Internet access among roaming users and has created a strong demand for access to a number of IP multimedia services like VoIP.

Till now *wireless VoIP*, the name given to this new trend in VoIP services, has been limited to the Local Area Network (LAN) level where voice communication are provided to a company over an enterprise wireless LAN (WLAN) or to people accessing public WLAN hot spots. Wireless VoIP has been also applied to the Personal Area Network (PAN) level, with the aim of offering in a domestic environment voice communication in a cordless-like fashion (e.g. using the Bluetooth technology). Now, thanks to high data rate 2.5G/3G cellular technologies, wireless VoIP can also be attainable over the Wide Area Network (WAN). Besides, as wireless data networks based on different access technologies (802.11, Bluetooth, 2.5G/3G networks, etc.) become more and more deployed, mobile terminals as cellular phones, PDAs and laptops will be equipped with multiple wireless networking interfaces. As a consequence we aim at an application scenarios where a multi-mode device will place mobile voice calls over wireless PAN/LAN networks (at home, office or public hot spots) as well as over public cellular networks.

This work was partly performed in the context of a cooperation with Bull Italia

The ability to use multiple networks in parallel gives the user a possibility to choose for example the most economical solution at the time, and for the operator or service provider the possibility to use the most suitable connection for each application. The issue is how to support seamless mobility to multi-mode terminals: that is to place and receive calls over the most suitable wireless interface and to maintain VoIP sessions alive while handing off between two heterogeneous wireless networks. As most of the wireless access networks are currently using private IP addresses and connected through NAT (Network Address Translation) elements, the support of the "NAT traversal" combined with mobility and handover functionality is a very important requirement.

In this paper we describe a solution for this issue. The solution takes care of the "vertical mobility" of a user among different access networks/technologies, considering that for each different attached network the terminal will receive and use a different IP address. On the other hand, the movement of the user among base stations of the same technology/network (e.g. among different access points in the same WiFi campus, or among different 3G cells in the cellular network of an operator) is handled by specific mechanisms of the given access network and no IP re-configuration is required.

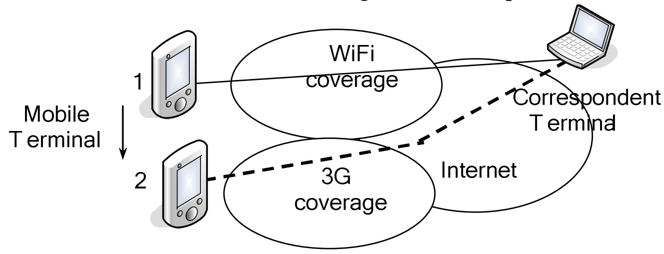

Fig. 1. VoIP handoff scenario in a Wi-Fi/3G interworking scenario

In our solution, we consider a Mobile Terminal with multiple network interfaces that can be active at the same time making it possible to realize a "soft handover". As an example Fig. 1 illustrates the scenario of a handover between a WiFi and a 3G network (this scenario has been also implemented in a testbed that will be described). Our solution is based on the Session Initiation Protocol (SIP) which will be used for both "traditional" VoIP signaling and for supporting terminal mobility. Moreover our solution integrates mechanisms to enable Mobile Terminals to make and receive VoIP calls regardless if they are located inside a public or private IP network.

The paper is organized as follows. Section II provides an overview of the problems to be faced and surveys related works. Sections III and IV presents our solution. Section V describes our testbed and the achieved measurement results.

II. Problem Overview and Related Work

The purpose of this section is to introduce some issues that will be addressed in the paper, as well as to present related work in the area. For space constraint we will not be able to introduce the SIP protocol [1] and its architecture, therefore we will assume that the reader is familiar with SIP elements (e.g. User Agents, Proxy Servers, etc.) and to the SIP signaling procedures.

A. Mobility support mechanisms for VoIP

There are two basic mobility management approaches to support mobility in VoIP services [2]. The first one exploits Mobile IP (MIP) [3] and operates in the network layer. Mobile IP is not directly related to VoIP applications, it provides the mobile terminal with a single IP address that is used for all applications and that can "follow" the terminal in its wandering. The other scheme seeks to support mobile VoIP services in the application layer by exploiting the features of SIP. This second mechanism is often referred to as "application layer mobility" and it provides mobility for the active sessions. Its goal is to keep a session (e.g. a phone call) active while the terminal is changing the IP access network and is consequently changing the used IP address (maybe also the used network interface). In both the approaches (MIP and SIP) there is the need to minimize the service disruption during a handoff procedure and this issue has been addressed in several works (see for example [4] and [5]). In our proposal, we focus on an "application layer" mobility management scheme, based on SIP.

B. NAT traversal issues

The "NAT traversal" is one of the most critical issues to ubiquitously deploy VoIP services in the real world. While a large part of the target terminals are using private IP addresses and are "hidden" behind a NAT, signaling protocols for VoIP, like SIP, do not "natively" support full NAT traversal. NAT traversal mechanisms are needed to allow terminals in "private" networks to be reached (i.e. "invited" to a phone call), and to allow the media streams to cross the border between the private and the public network, where the NAT box is placed. IPv6 promises to overcome the NAT issues but wide spread diffusion of IPv6 is not foreseen in the short/medium term. Several mechanisms have been added to SIP to deal with NATs [6], [7]. Proper combination of these mechanisms allows a SIP terminal in a private network to communicate with a terminal using a public IP address crossing a "typical" NAT. Unfortunately the general scenario where both terminals can be in private networks needs the intervention of a new element in the media path, which is not considered in the "canonical" SIP architecture. There are different ways to achieve this goal, some of these are under consideration for standardization, (e.g. STUN, TURN, ICE). Another solution is the insertion of an element typically referred to as "Session Border Controller" (SBC) [8]. This element is now an important component of several VoIP solutions in real world. The SBC acts as intermediate node in all signaling and media sessions of a user that wants to access the public network through a private access network. From the media point of view the SBC acts as "B2BUA" (Back to Back User Agent) while from the signaling point of view it can also act as a Sip Proxy. In practice the SBC "represents" the user in the public network by providing him/her with a public routable IP address through which the user can be reached even behind a NAT. The SBC can be owned by an enterprise to allow its hosts on a private network to make and receive VoIP calls or by a public VoIP provider to offer VoIP services to enterprises. In our solution we foresee to extend the functionality of the SBC to take care also of the handoff between heterogeneous wireless access networks.

C. WLAN/3G interworking

Several works addressed the issue of 3G/WLAN interworking/integration, both in the research community (see for example [9] and [10]) and in standardization fora like 3GPP or 3GPP2. The underlying idea is to include WLAN access in the set of services provided by 3G operators to their subscribers. Most of the work done focused on the authentication and security aspects, while the issues related to handoff are still to be investigated. On the other hand, in our work we consider complementary aspects. We do not consider the authentication issues and we assume that the terminal is able to authenticate (separately or in an integrated manner) to the different access networks; instead, we focus on the problem of seamless mobility amongst heterogeneous networks and vertical handoff management. A similar approach can be found in [11], which considers the classical SIP based mobility where the handoff is handled by the correspondent terminal. As explained in the next section, our solution envisages a different approach for SIP based mobility management.

III. Proposed Solution

The fundamental concepts of the proposed solution can be illustrated with the help of Fig. 2. Mobile terminals have access to different networks (in the figure, WLAN and 3G), which can overlap their coverage areas. The Mobile Terminal has separate interfaces, each one dynamically receives its (private or public) IP address from the corresponding wireless network. The Mobile Terminal contains a SIP client (User Agent) and a Mobility Management Client (MMC). The Session Border Controller contains a Mobility Management

Server (MMS) which is the main entity controlling the user mobility. Thanks to the interaction between the MMC in the mobile terminal and the MMS in the SBC the device can move between IP subnets, being reachable for incoming requests and maintaining VoIP sessions across subnet changes. The “CT” node shown in the picture is the Correspondent Terminal that communicates with the Mobile Terminal. A SIP Registrar is also included in the picture with a dashed line, as in our solution the SIP Registrar is not directly involved in the mobility management procedures for the handover.

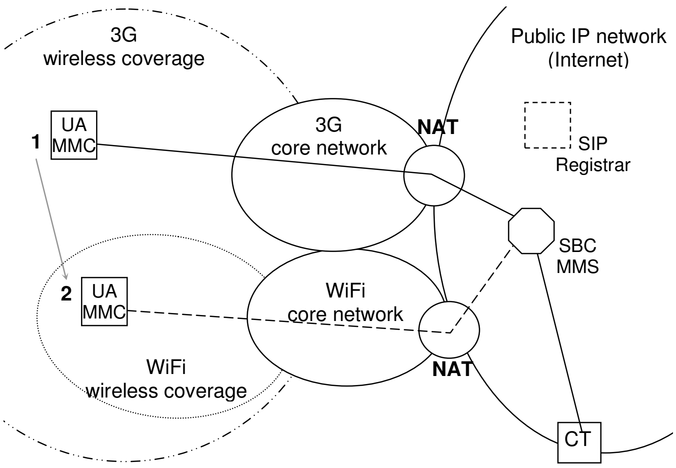

Fig. 2. Architecture

The MMS is an “anchor point” for the media flows which are transmitted over the wireless access networks directed to (and coming from) the MT. When the MMC in the MT detects that a handover is needed, it will request the handover to the MMS (via a SIP message) over the “target” network. The MMS in the SBC then will update its media proxy and will start transmitting and receiving the media over the target network (details are provided in the next section). Note that the entire handover procedure is handled by the MT and the SBC, letting the CT completely unaware of what is occurring. Instead, in the traditional approach the CT is directly involved by the MT with a new “Re-INVITE” transaction. In such case, the CT is in charge of performing the handoff by establishing a new media flow using the IP address of the terminal in the “target” network. Executing the handover with the CT could lead to high delays, for this reason various solutions have been proposed trying to overcome this problem by introducing intermediate entities as temporary anchor points (see [5]). In our solution the anchor point is not temporary, but permanent. Obviously, having a permanent media relay in the path is not an efficient and elegant solution, but we are not adding this new relay with our solution. We are assuming that the Session Border Controller will already be in the path. Hence, the proposed mobility anchor point is integrated with SIP processing and media relay capabilities provided by SBCs. It is important to note that a SBC works by processing all SIP signaling coming from and directed to the UA and corresponding media flows, modifying opportunely IP addresses and port numbers for NAT traversal. Extending its capabilities to support the handoff is a straightforward step.

In order to configure the IP addresses on the MT interfaces, existing mechanisms are used (e.g. DHCP on the WLAN and PPP on the 3G interface). Then, by interacting with the underlying operating system, the MMC always maintains an updated status of such networking configuration. When multiple interfaces are active, the MMC needs to select the preferred interface for sending/receiving the media flows (while the terminal is involved in a call) or for exchanging SIP signaling (both during calls and in idle state). The choice of the selected interface performed by MMC may depend on cost aspects and/or on QoS issues (signal strength, perceived packet loss and/or delay). The MMC will also control the sending and receiving of media packets, dealing with changes in IP addresses and ports during the handover procedures.

Two architectural solutions can be envisaged in the Mobile Terminal to support the proposed solution.. In the first case the User Agent is aware of the handoff and integrates the MMC (or closely interacts with it) to handle the Mobility Management signaling and the sending/receiving of packets over the different interfaces. In the second case the MMC acts as a so called “B2BUA” (Back to Back User Agent) running on the same MT and masquerading all mobility and NAT traversal functionality by relaying both signaling and media flows. In this case the UA sees the MMC as default outbound proxy and has no knowledge of the handovers. Note that the two solutions only differs in the internal implementation in the Mobile Terminal, there is no difference in the external behavior of the procedures.

The former solution has the advantage that it requires less processing resources to the MT with respect to the second approach. On the contrary, the latter solution has the advantage that existing SIP UAs can be easily supported/reused without any changes. Therefore we choose to implement the MMC as a separate element and from now on we will use this assumption.

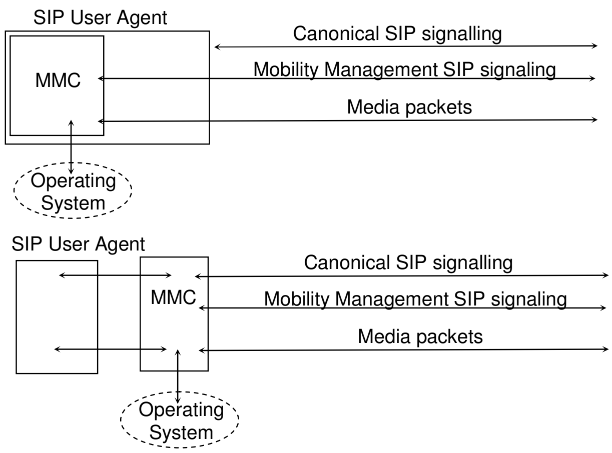

Fig. 3. Two solutions for the MMC in the mobile terminal

IV. Specification of the Procedures

As described in the previous section, the mobility management involves four main functional entities. On the MT sides there are the SIP UA and the MMC, while on the network side there are the SBC with the MMS and a SIP Registrar.

The Session Border Controller (SBC) enhanced with a Mobility Management Server (MMS) is needed to manage MT handoffs between different access networks providing service continuity and NAT traversal. The SBC is able to process both SIP protocol header fields and Session Description Protocol (SDP) [12] bodies in order to force itself as relay for the media packets.

In the SIP architecture, the SIP Registrar records the current location(s) of the user. In our proposal the mobility of the Mobile Terminal between different access networks is controlled by the SBC. One could conclude that there is no need for an external SIP Registrar and that the SBC could play the role of the user's "home" Registrar. We argue that a better solution is to keep the two entities separated, at least from a logical point of view. The SIP Registrar can handle more complex situations where the user needs to have multiple contemporary registrations with multiple terminals (e.g. fixed phones, additional mobile terminals, PC with software phone, video conferencing station etc.). Each of these terminals will send its registrations to the SIP registrar (where a proper service logic will coordinate the call processing). On the other hand the SBC will be seen by the SIP Registrar as one single terminal and the SBC will hide to the SIP Registrar the Mobile Terminal roaming across different access networks. The MMC is configured to use the SBC as nexthop/outbound proxy.

The SBC will take care of NAT traversal, so that the Mobile Terminal can be reached by SIP signaling and can send/receive media flows even beyond a NAT. As for SIP signaling, the MMC in the MT and the MMS in the SBC implement the SIP extension described in [13] which allows the MMC to receive SIP responses on the same port where it sent corresponding SIP requests. Then a "keep-in-touch" mechanism is needed to keep the pinhole in the NAT open. Various techniques can be used [1] such as dummy packets (from the MMC to the MMS or vice-versa) or well formed SIP messages, we use periodic SIP register messages from MMC to MMS. The "keep-in-touch" packets are sent every 30 seconds, so they use a very limited amount of resources.

A. Initial Registration

The initial registration procedure is initiated by the User Agent, when it sends a registration request to the SIP Registrar. The message is sent by the UA to the MMC. The MMC forwards it to the MMS which creates the association between the MT and its point of contact. Acting on behalf of the Mobile Terminal, the MMS will forward the registration to the SIP Registrar,. which needs to update the contact address associated with the user's "address-or-record" (that is the public user identifier) so that it points to the SBC. To this aim, the SBC modifies the Contact header field inserting its own address. Another solution would be to use the Path header field [13], if the targeted Registrar Server supports the Path extension. From now on, only the MMS will keep track of the Mobile Terminal movements, while the SIP Registrar will just believe that the MT location is the IP address of the SBC. The sequence diagram of this procedure is described in the upper part of Fig. 4.

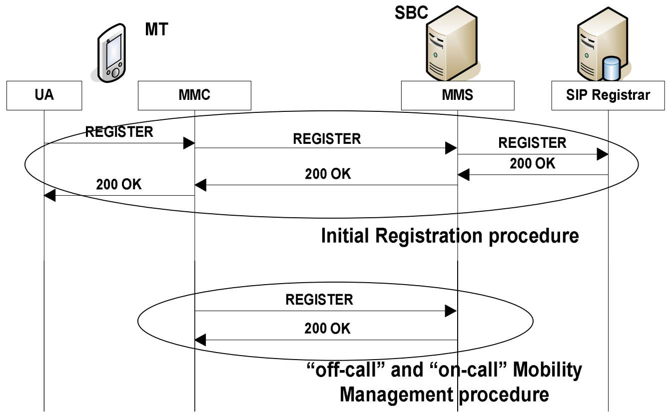

Fig. 4. Registration and Mobility Management procedures

B. Off-call Mobility

The "off-call" mobility management procedure is performed after the Initial Registration procedure while the terminal is not engaged in a call. This procedure updates the MT point of contact in the SBC, so that incoming session establishment requests are routed over the active and selected interface.

The procedure is activated by the MMC whenever the need for a handoff is identified, provided that both the involved wireless interfaces have been assigned an IP address. The procedure consists in a SIP re-registration procedure over the "target" network followed by the activation of the "keep-in-touch" mechanism.

As illustrated in the lower part of of Fig. 4, the procedure is executed only between the MMC and the MMS. The MMC sends over the new wireless interface a SIP REGISTER request directed to the SBC. The SBC processes the request in order to update the previously stored contact information associated with the user and then sends back the 200 OK message. The "keep-in-touch" mechanism is then switched off on the old network and activated on the new network.

C. Session Establishment

The session establishment procedure consists in a standard SIP session setup procedure. Session establishment messages from/to the Mobile Terminal are sent/received over the selected wireless interface as identified in the SBC by the Initial registration procedure and the off-call mobility management procedure. The signaling flow between the MT and the CT involves the SBC which is in charge of proxying the SIP messages. However before relaying the INVITE request from the caller and the corresponding 200 OK response from the callee the SBC modifies the SDP bodies of these messages in order to act as RTP proxy for media flows in both directions. This is needed to correctly handle NAT traversal in the path towards the MT, and it is done by exploiting the symmetric RTP approach as in existing Session Border Controller.

Once the session is established, the media packets start to flow over the selected wireless interface. In principle, there is no need to send anything on the unselected active interfaces, that should be used only when an “on-call” mobility procedure occurs. On the other hand our practical experience suggested that starting sending the packets on the 3G interface introduces an initial delay that can be quite large and can cause noticeable disruption in the voice communication during the handoff. Therefore we introduce a “keep-alive” mechanism between MMC and MMS during the call phase: the MMC sends dummy UDP packets to the MMS over the unselected wireless interface. The MMS will take care of discarding these packets.

D. *On-Call Mobility*

The on-call mobility management procedure takes place when the UA identifies the need for handoff during an ongoing VoIP session. A “forward” handover procedure has been defined, i.e. the handover signaling messages will be exchanged on the target network. Therefore the handover can be performed even if the communication on the old network is interrupted abruptly. The handover procedure is MT initiated. The mobile terminal sends an “on-call” REGISTER message over the target network interface addressed to the MMS in the SBC. Differently from an “off-call” REGISTER request, the “on-call” REGISTER request contains in the message body the reference to the active session to which the handover is referred.

In the same time, the Mobile Terminal starts duplicating the outgoing media packets on both interfaces (unless the old interface has gone down). As soon as the MMS in the SBC receives the REGISTER message, it will:

- start accepting packets coming from the new interface and discarding the ones coming from the old interface for the media flows corresponding to the session ID contained in the REGISTER message body;
- send back the SIP 200 OK message to the MMC in the Mobile Terminal;
- start sending the media packets directed to the Mobile Terminal using the new interface.

Thanks to the fact that the terminal has already started sending the packets on the new interface, the duration of the handover is minimized.

The most critical issue is that the “on-call” REGISTER message could be lost for any reason, delaying the handoff procedure. The standard SIP procedure foresees that the client performs a set of retransmission of the REGISTER if the 200 OK is not received back. The SIP standard suggests a default value of the retransmission timeout equal to 500 ms, that is doubled on each retransmission. However this is not compatible with a reasonable performance of the handover in case of the loss of the REGISTER message. Therefore we mandate that for the “on-call” REGISTER message a different duration of the retransmission timer is used. REGISTER messages are sent with a fixed interval of 200 ms until the 200 OK is received or a transaction timeout occurs. We set the transaction timeout is set to 3 s corresponding to a maximum of 15 retransmissions.

On the terminal side, the MMC will stop duplicating the packets on both interfaces as soon as the 200 OK is received or the first media packet is received on the new interface. Note that if the media packet is received, but no 200 OK message, the MMC will still continue sending the REGISTER message until the register transaction expires.

E. *Comparison with canonical SIP based mobility*

In classical SIP based mobility, when the MT moves to a new access network changing its IP address during a call it re-invites the remote CT in order to re-establish a new media streams (the handover is handled by the the remote CT). Then the MT has to register the new address to the SIP registrar. Instead, in the proposed solution the two functions are tied in just one registration procedure between the MT and the SBC, while the corresponding terminal is let completely unaware of the MT movement. This increases the handover performances, increases the compatibility with legacy remote terminals that might not handle correctly the re-invitation procedure, guarantees a better privacy (since the position and movement of the MT are hidden by the SBC).

F. *Handoff Criteria*

We want to mention that several criteria can be used (also in combination) to drive the handover decision. Monitoring of the quality of the received signal (power, S/N ratio) is one option. The problem is that signal quality information is only local and does not take into account the load of the wireless cell nor the load of the network between the cell and the SBC. This suggests to implement some kind of IP level measurements of the QoS in the path between the MT and the SBC over the different wireless networks. An overall policy should also be supplied, such as “always use a WiFi access when available” or “always use the best available network connection”.

V. Implementation and Measurements

We have implemented the proposed solution and realized a testbed across commercial WiFi and 3G networks. The Mobile Terminals has been implemented using laptops with Windows XP, the SBC and the SIP Registrar are implemented on a standard PC (both Windows XP and Linux can be used). MMC and MMS have been implemented in Java using (and modifying) the open source MjSip Java SIP stack [14]. As SIP

User Agent we used the Xlite software client [15]. The laptop is equipped with an internal WiFi card and with a PCMCIA card for 3G access. As WiFi access network we used both an our own Access Point connected to the Campus Fixed Lan, and a WiFi network in our labs which is connected to Telecom Italia backbone. As 3G network we used the Vodafone network. The SBC/SIP Registrar was located in our campus LAN and given a public IP address. As Correspondent Terminals we experimented both a PC in our campus LAN and a PC using an ADSL access.

In the Mobile Terminal, the MMC interacts with the operating system by checking the status of the interfaces with the "ipconfig" command. In the current prototype the MMC does not implement the monitoring of the signal quality of the interfaces nor it takes IP level measurements (the implementation of these features is for further work). The MMC offers a simple Graphical User Interface which shows the currently active interfaces and allows to control the handover by choosing the "selected" interface among the active ones. Fig. 5 below depicts the testbed layout.

A. Evaluation of handover performance

We analyzed the performance of the handover by capturing the media and signaling packets on the MT and on the SBC, using the Ethereal passive measurement tool [16]. We did not consider the path between the SBC and the Correspondent Terminal, as it does not impact the performance of the handover. The GSM codec at 13 kb/s was used. We have recorded the departure and arrival times of voice packets at the MT and at the SBC. We analyzed both the uplink flow (MT→SBC) and the downlink flow (SBC→MT) and we considered the handovers from WiFi to 3G and vice-versa (in total we have 4 scenarios).

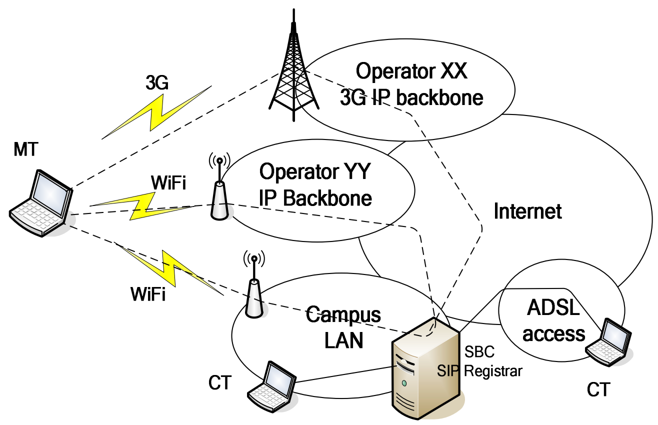

Fig. 5. Testbed layout

Looking at the 4 graphs in Fig. 6, in the x axis we put the departure time of packets from the originating interface, while in the y axis we put the arrival time of the packets at the destination interface. As the clocks are not synchronized, the time is relative to the first sent or received packet on the interface and we are not able to measure the absolute "one-way delay". This is not a problem, as we are interested in the differential delay among arrived packets. For the different scenarios we will discuss: 1) the impact of the difference in the one-way delay between the WiFi and the 3G network during the handover; 2) the handover completion time, i.e. the time elapsed from when the MMC starts the handover procedure and when the procedure is completed and the voice in both directions is flowing on the target interface.

Let use define as U_{up} and U_{dn} the one way delay for the 3G network in the uplink (MT→SBC) and downlink (SBC→MT) direction. These delay do not only cover the 3G network, but all the path from MT to SBC, crossing the 3G network (see Fig. 5). Similarly we define W_{up} and W_{dn} for the WiFi network. In the performed experiments, the measured round trip time between the MT and the SBC for the 3G access (i.e. U_{up}+U_{dn}) was in the range of 350-400 ms, while for the WiFi access (i.e. W_{up}+W_{dn}) was in the range of 120-250 ms .

Fig. 6 reports the results for the 4 scenarios. From the diagrams related to uplink (a and b) we can give an estimate of the difference in the "one-way delay" for the 3G access and for the WiFi access in the "uplink" (i.e. U_{up}-W_{up}). As the packets are duplicated, the difference in the y axis between the arrival of the same packet sent on the WiFi and on the 3G interface is the delay difference. It turns out that at the time of our tests, uplink one-way delay experienced in the 3G access is 60 to 100 ms higher than the one experienced in WiFi.

Let us consider the handover from WiFi to 3G, in the uplink case. The temporal diagram in Fig. 7 shows the sequence of events. When the MT takes the handover decision, it sends a SIP register message (on the target interface, 3G in this case) and start duplicating the packets. As in our case the WiFi is faster, some packets will arrive to the SBC before than the REGISTER message and will be forwarded. When the REGISTER arrives, the SBC starts discarding the packets from the WiFi interface (they are marked with a square in Fig. 6-a). A set of packets will arrive from the 3G interface which are the copies of the already arrived packets. This packets, marked with a circle in Fig. 6-a, will be forwarded and received by the CT as duplicated packets. The duration of this burst of duplicated packets is equal to the difference of the "uplink" one way delay between WiFi and 3G (U_{up}-W_{up}). As for the handover completion time, it roughly corresponds to the round trip time on the target interface (3G in this case: U_{up}+U_{dn}). We measured 370 ms for the interval between the REGISTER and the 200 OK in the test shown in Fig. 6-a. As a confirmation, we can see that in Fig. 6-a the packets are duplicated from t=250 ms to t=620 ms.

In case of the handover from 3G to WiFi (always in the uplink case), the MT sends the SIP REGISTER message on the "faster" WiFi network and starts duplicating the voice packets. When the SBC receives the REGISTER it will start accepting packets sent on the WiFi interface and discarding

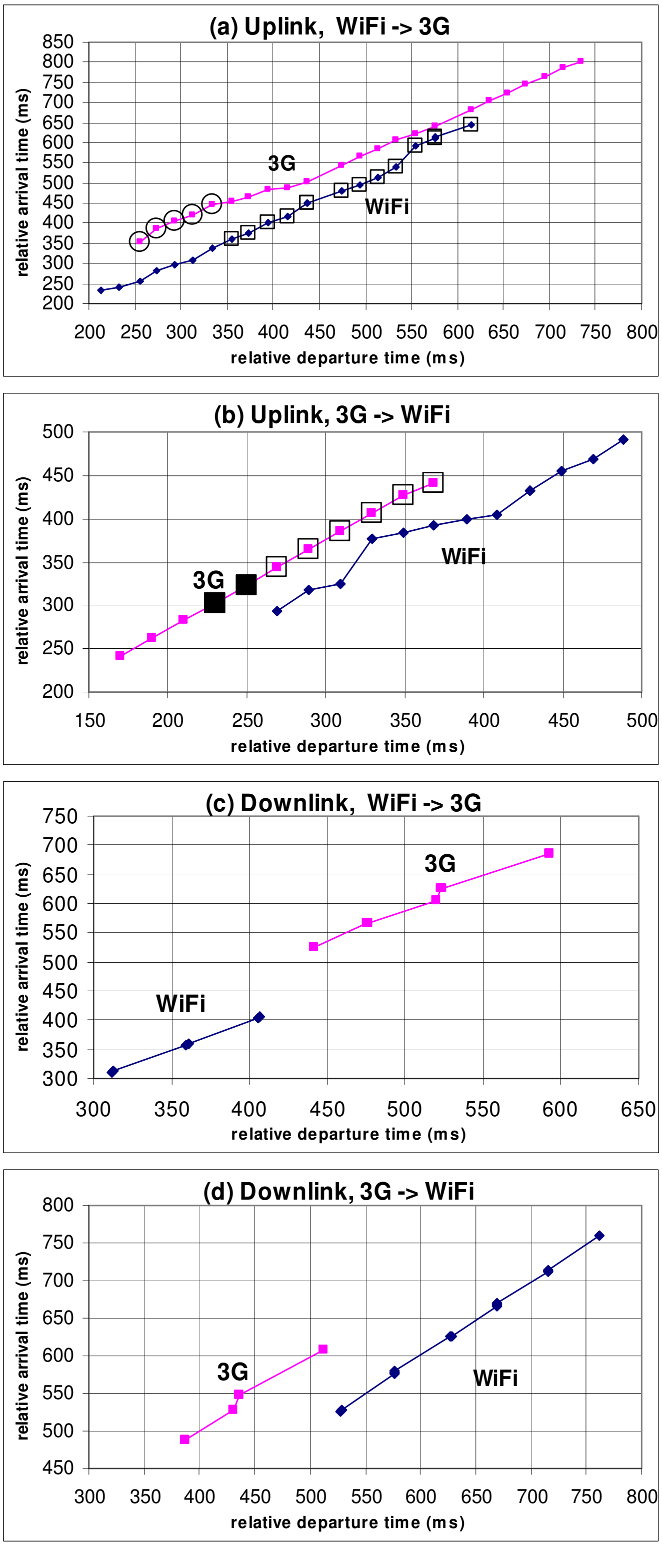

Fig. 6. RTP arrival patterns during handovers in 4 scenarios

those sent on the 3G interface (marked with a square in Fig. 6-b). The first received packets sent on the WiFi interface will have an higher sequence number than the last one received coming from the 3G interface, as the packets sent on the WiFi interface have "overcome" the ones sent on the 3G interface. A number of packets will be lost, and these packets are marked with the solid square in Fig. 6-b. The duration of the burst of lost packets is again equal to the difference in the uplink one way delay between 3G and WiFi network. As for the handover completion time, we measured 130 ms for the interval between the REGISTER and the 200 OK in the test shown in Fig. 6-b. In fact, the packets are duplicated from t=250 ms to t=380 ms.Coming to the downlink flows, let us consider the handover from WiFi to 3G (Fig. 6-c). The SBC will stop sending packets towards the WiFi network and start sending them towards the 3G network when the REGISTER message is received.

The first packet sent towards the 3G network will experience an additional delay equal to the difference in the "downlink" one way delay between 3G and WiFi network (U_{dn}-W_{dn}).

Finally, let us consider the handover from 3G to WiFi for the downlink flow (Fig. 6-d). As soon as the REGISTER message is received by the SBC the packets will be sent towards the WiFi interface and will arrive at the MT in advance with respect to packets with lower sequence number previously sent towards the 3G interface. The duration of the advance is equal to the difference in one-way delay.

Similarly to what we have done in the uplink, from the data reported in Fig. 6-c and d we can evaluate the difference in one-way delay for the downlink U_{dn}-W_{dn}. In our test we measured that 3G one-way downlink delay was from 80 to 100 ms higher than WiFi.

The results show that the different delay between the WiFi and 3G network is a critical factor. If the differential delay is reasonably low the voice decoder is able to hide the handoff. In our tests, where U_{up}-W_{up} and U_{dn}-W_{dn} are in the order of 100 ms, the handovers are not perceived.

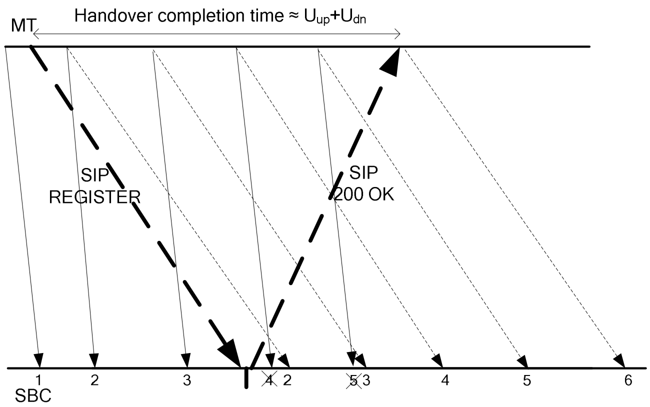

Fig. 7. Temporal diagram for WiFi → 3G handover

It is interesting to compare the results shown in Fig. 6 with the corresponding measurements without using the keep-alive mechanism introduced in section IV.C. As we can see in Fig. 8, which reports the uplink measurement for the WiFi to 3G handover, the initial differential delay between the 3G and WiFi is in the order of 2,8 s. Correspondingly, we have that the duration of the handover (during which all packets are duplicated) is in the order of 3 s. This is due to the fact that starting to transmit over a 3G interface requires a considerable amount of time. Just for comparison, we have reported the diagram of Fig. 6-a in an arbitrary position in the left part of Fig. 8. It is possible to appreciate the difference in terms of handover duration and of the distance between 3G and WiFi packet arrival time. The conclusion is that the keep-alive mechanism is needed to support seamless handover.

The results shown in Fig. 6 consider the favorable case in which both interfaces remain active during the handover. It can happen that the old interface goes down suddenly and does not allow to transmit packets during the handover. In our solution, the uplink flows are not affected, as the MT starts sending packets on the new interface immediately. On the contrary, the downlink flows are affected, as the SBC will start transmitting packets towards the new interface only after receiving the handover request from the MT. We have analyzed this case with temporal diagrams similar to the one shown in Fig. 7 and we have repeated the measurements of handover performance. The full results are not shown here for space constraints, anyway we have found theoretically and measured from the testbed that on the downlink flow we have an impairment in the order of $U_{up} + W_{dn}$ for the handover 3G→WiFi and in the order of $W_{up} + U_{up}$ for the handover WiFi→3G.

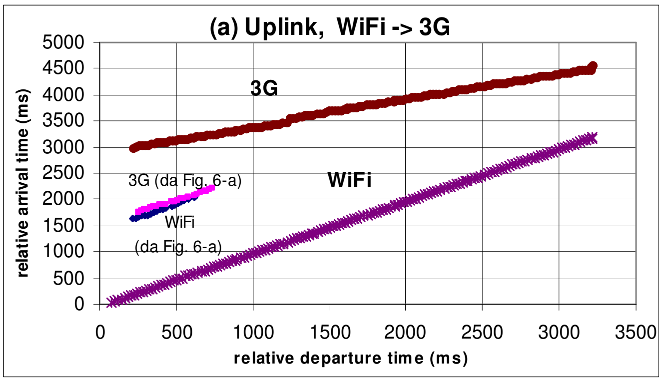

Fig. 8. RTP arrival pattern without keep-alive on the unselected 3G interface

VI. Conclusions And Future Work

In this paper we have presented a solution for seamless vertical handover between heterogeneous networks like WiFi and 3G based on SIP. The novelty of the solution is that it is strongly coupled with the NAT traversal features provided by the so called Session Border Controllers. Assuming that the Mobile Terminals will be mainly roaming on networks with private IP numbering, the mediation of an SBC is already needed to use SIP based services, therefore we straightforwardly propose to enhance SBC functionality to support the mobility. The proposed solution can be exploited in the short term by a 3G operator willing to extend its services to WLAN, by a VoIP provider that uses the 3G network as IP transport and by an enterprise that wants to directly manage its voice services. In the long term this kind of approach will likely need to be included in 4G networks, which aim to support communication over heterogeneous networks (including legacy networks) in a seamless way.

Ongoing work concerns the realization of the mechanisms to drive the handover decision (both collecting the signal quality information from the network cards and making IP level measurements during the call active phase.

We recall that our solution relies on the fact that the MT is able to authenticate and get an IP address on the different visited networks, and how this is done was out of the scope of this work. The IRAP (International Roaming Access Protocols) [17] forum has exactly the goal to support establishing safer, simpler, seamless connectivity at PWLAN hotspots and is addressing integration with cellular networks (2G and 3G) from the authentication perspective. We are now working on including IRAP specifications in our prototype.

References

[1] J. Rosenberg *et al.*, "SIP: Session Initiation Protocol," IETF RFC 3261, June 2002.

[2] H. Schulzrinne and E. Wedlund, "Application-Layer Mobility Using SIP," *Mobile Comp. and Commun. Rev.*, vol. 4, no. 3, July 2000.

[3] C. Perkins et al. "Mobility support ofr IPv4" RFC 3344.

[4] A. Campbell et al., "Design, implementation and evaluation of cellular IP", IEEE Personal Communications Magazine, vol. 7, pp. 42-49, Aug 2000

[5] Ashutosh Dutta, Sunil Madhani, Wai Chen, Onur Altintas, Henning Schulzrinne "Fast-handoff Schemes for Application Layer Mobility Management", PIMRC 2004, Spain.

[6] J. Rosenberg and H.Schulzrinne, "An Extension to the Session Initiation Protocol (SIP) for Symmetric Response Routing", IETF RFC 3581, August 2003..

[7] D. Yon and G. Camarillo, "TCP-Based Media Transport in the Session Description Protocol (SDP)", IETF RFC 4145, September 2005.

[8] J. Cumming "Sip Market Overview" http://www.dataconnection.com/network/download/whitepapers/sipoverview.pdf

[9] M. M. Buddhikot, G. Chandranmenon, S. Han, Y.-W. Lee, S. Miller, L. Salgarelli, "Design and Implementation of a WLAN/CDMA 2000 Interworking Architecture", IEEE Communication Magazine, November 2003.

[10] Giuseppe Ruggeri, Antonio Iera, Sergio Polito: "802.11-Based Wireless-LAN and UMTS interworking: requirements, proposed solutions and open issue". Computer Networks 47(2): 151-166 (2005).

[11] A.Dutta, B.Kim, T.Zhang, S.Baba, K.Taniuchi, Y.Ohba, H.Schulzrinne "Experimental Analysis of Multi Interface Mobility Management with SIP and MIP", IEEE Wirelesscom 2005.

[12] M. Handley and V. Jacobson, "SDP: Session Description Protocol," IETF RFC 2327, April 1998

[13] D. Willis and B. Hoeneisen, "Session Initiation Protocol (SIP) Extension Header Field for Registering Non-Adjacent Contacts", IETF RFC 3327, December 2002.

[14] Mjsp open source JAVA SIP stack, http://www.mjsip.org

[15] Xlite SIP User Agent, http://www.counterpath.com/

[16] G. Combs *et al.*, "Ethereal: A Network Protocol Analyzer", http://www.ethereal.com/

[17] International Roaming Access Protocols (IRAP), http://www.goirap.org/