

Application-level Middleware to Proactively Manage Handoff in Wireless Internet Multimedia

Paolo Bellavista, Antonio Corradi, Luca Foschini

Dip. Elettronica, Informatica e Sistemistica, Università di Bologna,
Viale Risorgimento 2, 40136 Bologna, Italy
{pbellavista, acorradi, lfoschini}@deis.unibo.it

Abstract. New deployment scenarios tend to consider the requirement of session continuity for service provisioning, especially multimedia streaming, to limited heterogeneous portable devices roaming among wireless localities. In particular, multimedia streaming should not experience any interruption while clients roam in wired-wireless integrated networks based on the standard best-effort Internet. The paper proposes an application-level middleware approach to proactively overcome Wi-Fi handoff and maintain multimedia session continuity in the wireless Internet by exploiting mobile proxies running on the wired network. Mobile middleware proxies locally support resource-limited clients, avoid packet losses during handoffs, pre-fetch local buffers with multimedia contents before handoff occurrence, and possibly reconfigure/re negotiate ongoing sessions after handoffs. Experimental results show that, notwithstanding the application-level implementation, mobile proxies can avoid streaming discontinuities with good efficiency in wireless-wired integrated networks even if their pro-activity is based on simple and lightweight handoff prediction techniques.

1 Introduction

A more and more common deployment scenario is the *Wireless Internet* (WI), where wireless solutions enhance the accessibility to the traditional wired Internet and to its services via IEEE 802.11 Access Points (APs) that work as bridges between fixed hosts and wireless devices [1, 2]. The popularity of personal portable devices and the increasing availability of WI APs are suggesting the provisioning of distributed services to a wide variety of mobile terminals, even with heterogeneous and limited resources. Even if device and network capabilities are growing, the development of WI applications is still a very challenging task, in particular when dealing with continuous services, i.e., applications that distribute time-continuous flows with Quality of Service (QoS) requirements, such as in the case of audio/video streaming [3].

One of the most challenging issues in supporting WI continuous services is the avoidance of flow interruptions when clients roam from one wireless locality to another, e.g., during handoffs between different AP coverage areas. We distinguish two types of datalink handoff: i) hard handoff, where the destination cell AP takes over from the origin cell AP reactively in a relay mode, by minimizing signaling overhead, but by increasing latency and packet losses; and ii) soft handoff that proactively acti-

vates the new data path to the destination AP before client disconnection from the origin cell [4]. The Wi-Fi standard specification adopts hard handoff, thus complicating the development and deployment of continuous services [5].

All the above issues call for novel middlewares aimed at providing session continuity to multimedia services and capable of operating service management operations to fit the dynamic properties of deployment environments, in particular to overcome handoff-specific continuity issues. The paper proposes an innovative mobile middleware integrated with an original lightweight mechanism for Wi-Fi handoff prediction. One of the paper core claims is the opportunity of working at the application level to dynamically handle handoffs in WI multimedia: only middleware infrastructures that can access application-level information (characteristics of exchanged multimedia flows, user preferences, installed software, ...) can effectively take over client/server responsibility of application components and can perform the challenging operations of handoff handling on their behalf, thus facilitating the design and implementation of WI multimedia applications. In addition, we claim the suitability of middleware solutions based on mobile proxies that work as intermediate entities to assist and support client devices in scarcity of resources: this architecture guideline permits to introduce proxies whenever in need over the service path, to flexibly handle and adapt multimedia flow delivery without leaving the whole management burden to client/server end-nodes; mobile proxies can predict client handoff in advance and migrate to follow client movements, by exploiting pre-fetched data to sustain streaming until the completion of needed flow re-directions.

The paper describes a proactive handoff management solution integrated with our Mobile agent-based Ubiquitous multimedia Middleware (MUM)¹, an application-level proxy-based infrastructure to support both streaming quality adaptation and session continuity, independently of WI client roaming [6]. On the one hand, MUM performs quality of service management by exploiting mobile middleware proxies dynamically distributed to intermediate nodes along the activated service paths between clients and servers. On the other hand, MUM provides session continuity by predicting client movements and by proactively activating service path reconfiguration to accelerate application-level handoff. The paper also presents a thorough evaluation of the handoff management performance of the MUM middleware in a wide-scale simulated deployment environment, by taking into consideration the different handoff situations, i.e., micro/macro/global handoffs, as described in the following. The reported experimental results show that, notwithstanding the portable Java-based implementation, MUM prediction-based handoff management avoids streaming interruptions in most common deployment scenarios.

2 Design Guidelines for Handoff Management Infrastructures

We claim the need for application-level middlewares based on the introduction of mobile proxies to ease the design and implementation of WI continuous services by providing flexible solutions to crucial mobility issues, e.g., application-specific cach-

¹ Additional information, experimental results, and the prototype code of the MUM middleware are available at <http://lia.deis.unibo.it/Research/MUM/>.

ing and filtering, QoS management, and interoperable session control (see also the related work section) [7, 8]. Let us introduce some of the benefits that application-level support solutions can provide with an example. Consider two users, Alice and Bob, who have subscribed for the same video broadcasting service to watch daily news while commuting by bus. The video broadcasting service is delivered through a 4G network composed by Wi-Fi hotspots, deployed at the bus station and at bus stops, and by a UMTS infrastructure deployed over the traversed city districts. Bob accesses the service from his full-fledged laptop and has subscribed for gold quality, i.e., maximum resolution and best possible network QoS, while Alice exploits a Personal Digital Assistant (PDA) and has subscribed for a bronze quality level (small frame size and no QoS guarantees at all). One day, Alice and Bob are sitting on the same bus when a vertical handoff between the Wi-Fi and the UMTS networks occurs due to the bus leaving the Wi-Fi enabled station: application-level middlewares can react to the discontinuity in available bandwidth by properly adapting provisioning depending on the differentiated profiles of Alice and Bob: Bob's video frames should become smaller while he continues to access the broadcasting service; Alice should have her service downscaled to only-audio streaming.

We claim that middlewares should relieve client/server application components from the above, possibly complex, handoff handling operations. Moreover, handoff management should be realized at the application layer since only this abstraction level offers the needed flexibility and expressiveness. Let us state that lower-layer handoff management solutions could neither take service-dependent decisions nor perform management operations selectively, e.g., only for WI multimedia.

In addition, handoff management middlewares should be aware of the runtime characteristics of provisioning environments: for instance, Wi-Fi implements hard handoff, while other networks, as UMTS, implement soft handoff (during handoffs, UMTS clients can receive packets from both the old and the new AP); by focusing on Wi-Fi networks, handoff duration greatly depends on client card implementation [9]. The application level is the most suitable to provide middleware supports with full context awareness [10]. Moreover, handoff management can benefit from multimedia flow tailoring to fit provision-time changing characteristics of execution environments: when Alice and Bob roam from the Wi-Fi network to the UMTS one, the handoff management middleware should degrade the broadcast QoS to fit UMTS bandwidth. We claim that middleware components should mediate and manage service adaptation: only an infrastructure at the application-level can properly exploit applicative protocols such as RealTime Protocol (RTP), Session Initiation Protocol (SIP), and RealTime Streaming Protocol (RTSP), to adjust service provisioning by acting as a mediator between clients and servers [7, 11].

Moreover, we claim that handoff management should exploit the introduction of mobile proxies. The current Internet is already populated by several kinds of proxies (for caching, authentication, re-directing, ...) and we claim that proxy-based architectures will spread more and more in the future as the most suitable design choice in the WI era. The exploitation of intermediate proxies is an effective alternative to designing and implementing fat clients/servers, specialized for peculiar characteristics of the provisioning environment. Proxy-based architectures can act as the middleware glue to extend client/server capabilities with new facilities, as in the addressed case of proactive handoff management, which should otherwise be replicated for each spe-

cific application, with higher development and deployment costs.

In particular, by focusing on WI multimedia, we claim the suitability of middleware proxies running on nodes along the path between clients and servers, and of dynamically injecting the management logic necessary for handoff handling and continuous service adaptation. Proxy-based middleware solutions can enhance the traditional client-to-server interaction model by decoupling the two endpoints with the insertion of an active service path [6]. Active service paths show their effectiveness when there is the need for runtime reconfiguration of only limited path segments and for decentralizing management operations to intermediate nodes traversed by service flows, thus improving scalability [12, 13]. In open and dynamic scenarios, it is crucial to build active paths only where and when needed, depending on the location of client requests and on client mobility during service sessions. To this purpose we claim the need for innovative code mobility techniques to dynamically migrate, install, and discard middleware proxy components at provision time, thus avoiding unnecessary static pre-installation; mobile agents are an effective technology to enable the dynamic mobility of support components depending on application requirements [14].

By concentrating on WI multimedia streaming, mobile proxy-based solutions can also reduce client-to-server signaling when handoffs occur. Continuous services usually adopt connection-less protocols such as UDP, and manage data re-transmissions directly at the application level to react to high jitter and packet losses. In the fixed Internet, the deriving flow-control signaling is fairly limited and present only during network congestions/failures. On the contrary, Wi-Fi hard handoff causes relevant packet losses and is perceived as a network failure at the application level; in traditional multimedia systems, that may produce non-negligible client-to-server signaling and wrong perceptions of client situations at server side. For instance, Internet radios degrade the provided QoS level by switching to lower quality flows, while Video on Demand (VoD) services require client re-buffering after a handoff-driven congestion, thus interrupting flow visualization for the duration of client re-buffering. Middleware proxies located at wired network edges close to their served wireless clients can split the direct client/server connection and significantly reduce both signaling traffic on the service path and QoS degradation at client. It is crucial to have mobile proxies that can follow client roaming at provision time to maintain co-locality with their supported devices during service sessions.

3 Proxy-based Proactive Handoff Management

The section first introduces the different handoff types managed by the MUM middleware, and then focuses on the architecture of the proactive session handoff facility.

3.1 Handoff Types and Countermeasures for Multimedia Continuity

The WI consists of several networks that usually cover a large geographical area, each of them consisting of several sub-networks. It is possible to distinguish three types of handoff: micro, macro, and global handoff [4].

Micro handoff (intra-subnet handoff) relates to clients that roam between two dif-

ferent Wi-Fi cells without changing IP addresses, i.e., before and after handoff clients are attached to the same subnets. Nonetheless, given that Wi-Fi handoff is hard, micro handoff may produce packet losses when clients switch from origin to target APs. In particular, it includes three main phases: handoff detection, target AP search, and AP re-association [15]. The last phase is the fastest (usually a few ms) and is almost constant for different vendors [9]; the other two phases are longer and vendor-dependent, since IEEE 802.11 only specifies the mechanisms to implement them, by leaving unspecified their combination and durations. For instance, in a testbed composed by Cisco APs and Cisco/Orinoco Wi-Fi cards, different implementations of IEEE 802.11b probing algorithm largely influence AP search duration (about 400ms for Cisco and 60ms for Orinoco cards) [9]; the handoff discovery phase, i.e., the client decision to start or not a new AP search, is the longest and less standardized process, and may last from 1000ms to 1600ms depending on card implementation [15]. Therefore, the total micro handoff duration may even reach 2s, thus being incompatible with the jitter/packet loss requirements of usual WI continuous services.

Macro handoff (intra-domain handoff) refers to clients that move between two Wi-Fi cells attached to different IP subnets and includes network-layer handoff, i.e., client IP address change. In particular, macro handoff duration exceeds micro handoff of the time needed to get new IP addresses at target subnets. Widespread solutions for wired networks, such as Dynamic Host Configuration Protocol (DHCP), requires some seconds to complete address re-configuration and do not fit WI continuous services.

Global handoff (inter-domain handoff), instead, regards mobile clients that roam between two Wi-Fi cells attached to different Internet domains. This requires not only address change, but also some additional time to perform authentication, authorization, and accounting operations, usually necessary when entering a new access domain. Moreover, domain change could also require session re-configuration: for instance, in the new domain there could be a new server, functionally equivalent to the one used in the origin domain, with better QoS and/or lower pricing.

Client-side data buffering is a common solution in streaming over wired networks to smooth possible congestions and packet losses along client-to-server paths. Nonetheless, we claim the unsuitability of traditional buffering, i.e., only pre-fetching chunks of multimedia flows at clients, to support multimedia continuity during Wi-Fi handoffs. In fact, macro and global handoffs require service path reconfiguration to redirect provisioned flows to new client addresses and service reconfiguration time can be long (client-to-server roundtrip time, redirection request processing time at server, ...), thus imposing too heavy data buffering at clients. Moreover, if reconfiguration time is longer than handoff, all data transmitted by the server to client old addresses are wasted and must be re-sent, while they could be locally moved from the last client location to the new one until data redirection is completed. In addition, even in the case that client-side buffers could store enough data to sustain streaming continuity, long disconnection periods (up to 2s) could produce misleading signaling to the server and QoS degradation. Finally, when providing WI live streaming, it is not possible to command server flow rewinding and, consequently, long handoff durations necessarily produce streaming gaps as long as the difference between Wi-Fi disconnection time and client-side buffer duration. In any case, let us point out that WI client-side buffers should be as limited as possible since most clients are expected to have limited memory/storage capabilities.

3.2 MUM Proxies for Proactive Handoff Management

The architecture of the MUM infrastructure to dynamically activate intermediate nodes for quality control and adaptation of WI multimedia streaming has already been presented elsewhere [6]; in the following, the paper originally focuses on MUM proactive management of Wi-Fi micro, macro, and global handoff. MUM grants service continuity by performing proactive handoff management at the application level, with WI-multimedia-specific protocols that aim at minimizing server signaling and client requests for flow redirection. MUM provides a proxy-based middleware support for micro, macro, and global handoffs by splitting client-server direct interaction and managing handoff at intermediate nodes. In addition, MUM employs an original two-level buffering technique to store data flows at both clients and proxies. Two-level buffering interacts with our lightweight Wi-Fi handoff prediction technique (Received Signal Strength Indication-Grey Model - RSSI-GM) that distinguishes which type of handoff is likely to occur with an appropriate advance time depending on the handoff type; the extensive presentation of our RSSI-GM solution is in [16]. MUM proxies interwork with lightweight MUM client stubs, which are the only middleware components executing on client devices and in charge of transparently interfacing to even legacy multimedia players and of executing the RSSI-GM prediction algorithm.

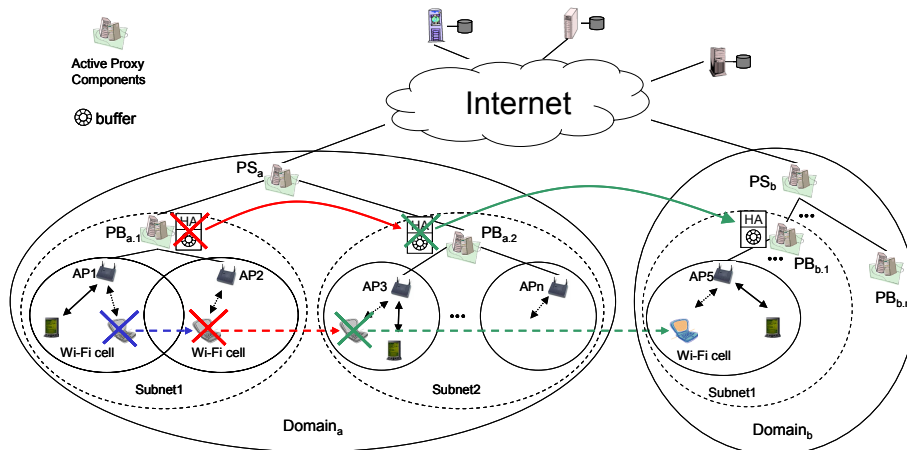


Fig. 1. MUM Handoff Management Architecture

In addition to client-side buffering, MUM originally proposes to exploit a second level of buffering proxies that use the storage resources of workstations on the wired network. Second-level buffers receive and store incoming flows during client disconnection, thus avoiding service interruptions, frame losses, and server re-transmissions. They are capable of promptly filling up client buffers at reconnection after handoff; they help in enabling streaming continuity also for clients with very limited memory resources; they interact with MUM handoff predictor to proactively increase their size only when handoffs are likely to occur. Besides, even in normal situations (far from handoffs) second-level buffers can improve visualization quality: they enable local retransmissions to client-side buffers in the case of packet losses due to the fragile

connectivity of last wireless hops. In addition, second-level buffers hosted on the wired network can widen the accessibility of multimedia services by alleviating power consumption on limited clients [17]. MUM devices also continue to host first-level buffers with limited size, as commonly employed to improve user-perceived quality, by smoothing network jitter effects: client-side buffers are crucial to provide multimedia frames to local players in the time interval of client disconnection during handoffs.

With a finer level of detail, MUM supports handoff management by activating a mobile proxy (the Handoff Agent – HA) at the Proxy Buffer (PB) node in the service path between client and server (see Figure 1). MUM automatically instantiates a new HA, implemented as a mobile agent, when a client starts a streaming session: once activated, HA performs proactive handoff operations and manages second-level buffering for its client for the whole session duration, by executing in its proximity and migrating to follow it in the case of subnet/domain change.

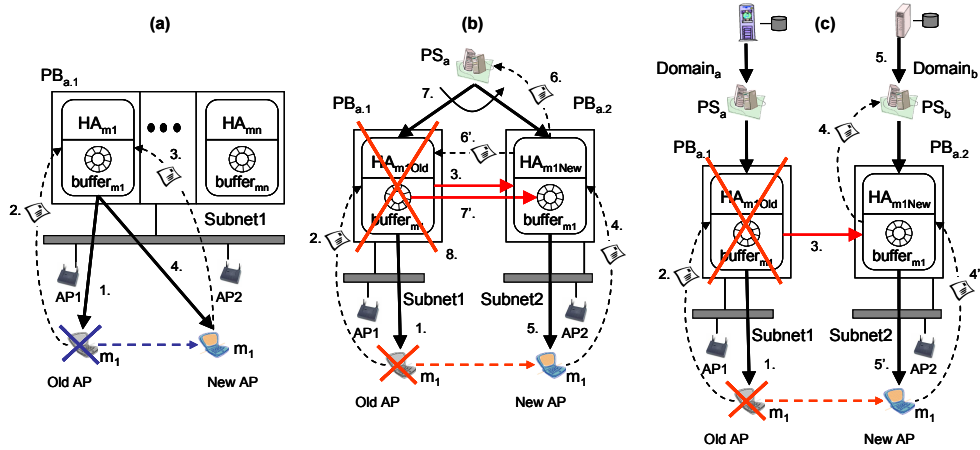


Fig. 2. Micro (a), macro (b), and global (c) handoff procedures

MUM micro handoff management starts when RSSI-GM foresees a client handoff. Triggered by handoff prediction notification (step 2 in Figure 2a), HA increases the size of its second-level buffer to store all data arriving from the server during micro handoff disconnection; second-level buffer size depends on micro handoff duration and, consequently, on several parameters related to client Wi-Fi card implementation. When the client is disconnected from its origin AP and not yet re-connected at the destination one, HA continues to buffer incoming streams as in normal conditions, while the player uses client-side buffered frames to sustain multimedia rendering. After reattachment at destination, the client stub notifies HA, thus causing the retransmission of lost frames (steps 3 and 4 in Figure 2a). Let us note that PB hosts HAs for all clients currently connected to the same subnet.

Macro handoff extends micro handoff with client address change, as shown in Figure 2b. By following the design principles of Section 2.2, MUM locally manages service reconfiguration and distributes handoff management load by exploiting code mobility. Local reconfiguration is made possible by introducing an indirection point

(the Proxy Switch – PS) in the client-to-server service path: PS works as an application-level router that forwards the multimedia flow to the HA current location. MUM micro and macro handoff management operations proceeds in the same way until handoff prediction. At the notification of a macro handoff, the old HA enlarges its second-level buffer, clones itself, sends its clone (new HA) with a copy of the current second-level buffer to the new PB (step 3 in Figure 2b), and continues filling up its buffer. Once attached to the new subnet and obtained a new IP address, the client stub triggers the retransmission of un-received multimedia frames due to temporary disconnection (step 4). As a consequence, the new HA starts refilling client buffer (step 5), commands PS to redirect the flow to itself (step 6), and asks the old HA for the multimedia data received after its migration (by indicating the last frame in the new second-level buffer, step 6'). After sending that data, the old HA terminates (steps 7' and 8); the new HA, instead, reconstructs the flow by merging data from old HA and PS, and forwards the merged stream to the stub. Let us note that the protocol requires no PS-side buffering and produces no signaling/transmissions unless clients connect to new subnets. Finally, to avoid DHCP long latency, MUM adopts Dynamic Registration and Configuration Protocol for IP address re-assignment [18].

If compared with macro handoff, the additional complexity of global handoff is the possibility and suitability of performing also non-local service path reconfiguration, possibly up to the server node. First of all, adding another indirection point would activate a new PS within the wired Internet core, while MUM aims at deploying its middleware components only at wired-wireless network edges. In addition, MUM has the goal of keeping separated different administration domains, to distribute proxy management responsibilities and to facilitate resource accounting. Global and macro handoffs proceed in the same way as far as the old HA sends its clone towards the new PB. Then, in global handoffs, while the new HA waits for its client, it immediately starts the server re-connection operations to reduce as much as possible reconfiguration time (step 4 in Figure 2c). After client reattachment (step 4'), the new HA refills client buffer while it continues waiting for the flow from the server to merge it with second-level buffer data, as in the case of macro handoff. Let us note that global handoff only requires minimum management between the two domains: similarly to [19], the only system management operation needed is to set logical correspondences between APs in reciprocal visibility at domain boundaries and PB host names, to enable the correct migration of HAs to their destination nodes.

4 Experimental Results

To thoroughly evaluate the effectiveness of the MUM handoff management infrastructure, it is necessary to test the MUM prototype in a wide-scale deployment scenario, i.e., a WI testbed composed by several domains, several subnets, dozens of Wi-Fi APs, and hundreds of served roaming clients. Such a large testbed is difficult to deploy and would require a large number of available mobile users to accomplish valuable experiments; therefore, as many other research proposals in the WI area, we have decided to exploit a simulator to feed the MUM prototype with realistic data about client roaming in a modeled wide-scale deployment scenario. Several wired-

network simulators are available, from both academy and industry; some of them also include wireless network modeling but, to the best of our knowledge, none addresses the specific problems of simulating different Wi-Fi card behaviors during handoff and of feeding pluggable software prototypes with simulation-generated handoff data. Thus, we have decided to develop a simple simulator that models user mobility and traversed tracks, calculates the RSSI values of mobile clients, feeds the MUM predictor with those values, and mimics physical/datalink/network-layer behaviors by temporarily interrupting streaming transmission during handoffs. The considered WI test-bed consists of 2 Internet domains and 6 subnets with 48 Wi-Fi cells; a mix of 600 wireless clients equipped with Cisco and Orinoco cards randomly move with variable speed between 0,6m/s and 1,5m/s; RSSI fluctuation has a 3dB standard deviation.

MUM components are implemented in Java and exploit the portable SUN Java Media Framework (JMF) for RTP-based video streaming and RTCP-based monitoring; in the experiments, we have provisioned H263-encoded VoD flows (frame size=176x144 pixels, and frame rate 8fps). Experimental results obtained in previous work have shown that HA migration and activation at new PBs takes less than 400ms; DRCP lasts about 100ms; macro handoff flow re-direction and merging require about 500ms; global handoff flow reconfiguration requires 340ms for PS interposition and 100ms for RTP streaming activation (with a JMF server); we assume 250ms as round-trip time [6]. In other words, handoff duration is between 440ms and 1500ms for micro handoff, about 1900ms for macro handoff, and about 2300ms for global handoff.

Among the different experimental results that can provide significant indications about middleware performance, the paper originally focuses on evaluating the impact of handoff prediction on proactive buffer management and, consequently, on the maintenance of streaming continuity. To this purpose, we have measured four primary performance indicators: *Efficiency* = $(PHO/PH)*100$, *Error* = $(NPH/PH)*100$, *Effectiveness* = $(SH/PHO)*100$, and *Continuity* = $(CV/SH)*100$. Predicted Handoffs (PH) is the number of handoffs foreseen by MUM predictor; Predicted Handoffs Occurred (PHO) is the number of PH corresponding to actual handoffs in the simulated environment; Non-Predicted Handoffs (NPH) is the number of actual handoffs occurred without an associated correct prediction; Successful Handoffs (SH) are PH occurred when the second-level buffer is correctly sized (for micro handoffs) or when new HA activation terminates before client re-attachment to target APs (for macro and global handoffs); and Continuous Visualizations (CV) represent how many times new HAs have completed merging operations before the termination of buffers from old PBs (this performance figure only applies to macro and global handoffs).

As reported in Table 1, Efficiency and Error experimental results show that short-term predictions are more challenging for the MUM middleware than long-term ones. That primarily depends on the characteristics of RSSI-GM prediction, which exploits past RSSI values to estimate future client position, and MUM stubs can maintain longer past RSSI sequences for global handoff prediction. Long-term predictions, however, have shown a higher standard deviation, i.e., the exact handoff time of global handoff is harder to predict. In general, Efficiency and Error are strictly related to required prediction advance time, while Effectiveness and Continuity mainly depend on correct second-level buffer sizing and prediction standard deviation. For instance, when early predictions occur, clients re-attach before the completion of the

MUM handoff procedure, thus decreasing Effectiveness and Continuity.

Table 1. Micro, macro, and global handoff performance indicators

	Efficiency	Error	Effectiveness	Continuity
Micro	90	10	99	-
Macro	93	8	98	99
Global	99	3	97	95

Micro handoff presents the worst Efficiency indicator, while Effectiveness is high, thus pointing out that MUM correctly dimensions second-level buffers in this case. Macro handoff limits useless HA migrations to 7% and achieves a good Error; Effectiveness is still high, i.e., almost all HA activations terminate before client re-attachment; Continuity is compromised only by sporadic late predictions. Finally, global handoff longer duration improves both Efficiency and Error; however, high standard deviation on predictions reduces Effectiveness and Continuity. Micro handoff Efficiency improvements are possible by anticipating handoff predictions and consequently paying an earlier second-level buffer enlargement, thus trading between Efficiency and storage overhead at proxies. Let us note that high standard deviation on predictions risks to produce cases of too late predictions, which also require second-level buffer refresh (at new locations) to renew obsolete data (sent in the meanwhile to clients at old locations). To overcome the problem, in global and micro handoff predictions, MUM commands HAs to continuously refresh the current and new second-level buffers until clients re-attach; in addition, MUM refreshes the buffers stored at new PBs until client re-attachments, and new HAs command PSs to forward flows to both old and new second-level buffers.

5 Related Work

The distribution of both traditional and multimedia services over the WI is introducing novel significant challenges. By focusing on handoff management, several research proposals have worked at the datalink and network layers of the OSI stack: [20] and [4] survey a number of micro, macro, and global mobility management proposals that tend to minimize handoff delays and packet losses in the context of Mobile IP. These general-purpose approaches suffer from the portability limitations typical of solutions at datalink/network layers and address handoff issues uniformly for any service developed on top of them, thus making impossible differentiated strategies for different application domains.

By concentrating on the few application-level approaches already available in the literature, [11] proposes to deploy a proxy-based infrastructure for service continuity in the specific case of macro handoffs by exploiting standard Linux tools to forward multimedia flows. More recently, the same researchers have enhanced their solution via SIP: the guideline is to employ multicast in a local domain and movement prediction to reduce packet losses during SIP session re-directions [7]. Their solution does not require ad-hoc infrastructures at client domains; however, it needs complex network management functionality, assumes a collaborative secure deployment environment, and does not include original handoff prediction mechanisms.

ment, and does not include original handoff prediction mechanisms.

Proxy Intelligent Modules for Adapting Traffic Efficiently (PRIMATE) proposes a proxy-based infrastructure to support micro, macro, and global handoff [8]. PRIMATE proxies are typically deployed at Wi-Fi APs, but can also install at border gateways connecting local networks to the rest of the Internet: similarly to MUM, they locally manage service reconfiguration without server intervention and predict client movements to adapt proxy-based buffering. However, PRIMATE prediction technique is centralized, scarcely scalable, and requires dedicated hardware; in addition, PRIMATE exploits kernel-level client modifications (data interceptors) that should be statically installed at each participant [21]. Another similar research proposal is [22], which describes a proxy-based mobile agent middleware for multimedia streaming: it proposes to employ proxies, called Virtual Servers, to assist roaming users but, differently from MUM, its proxies do not support proactive buffer movement and can be deployed only in intra-domain, ad-hoc, and statically installed network environments with specialized custom APs.

4 Conclusions and On-going Work

Middleware proxies working over the fixed network on behalf of their resource-constrained clients are demonstrating their suitability and effectiveness in the WI, especially when integrated with Wi-Fi handoff prediction. In particular, handoff prediction can help in realizing novel proactive proxy-based infrastructures that perform adaptive second-level buffering to eliminate/smooth the different discontinuity issues intrinsic to micro, macro, and global handoffs. The MUM research work is showing that it is possible to preserve WI multimedia streaming continuity, even to limited client devices, by adopting an application-level middleware approach that is portable and dynamically deployable over the standard WI.

The promising experimental results of the MUM prototype, both in the presented wide-scale simulated environment and in small WI deployment scenarios in our campus, are encouraging further research. In particular, we are extending our solution to support vertical handoff towards Bluetooth connectivity and to enforce different Service Level Agreements when client roam to highly populated congested Wi-Fi cells.

Acknowledgements

Work supported by MIUR FIRB WEB-MINDS and CNR Strategic IS-MANET Projects.

References

1. M.S. Corson et al.: Mobile and Wireless Internet Services: Putting the Pieces Together. IEEE Communications, Vol. 39, No. 6 (2001)
2. W. Stallings: Wireless Communications and Networks. Pearson Education (2001)
3. P. Ramanathan et al.: Dynamic Resource Allocation Schemes during Handoff for Mobile Multimedia Wireless Networks. IEEE Journal on Selected Areas in Communications, Vol.

- 17, No. 7 (1999)
4. D. Saha et al.: Mobility Support in IP: a Survey of Related Protocols. IEEE Network, Vol. 18, No. 6 (2004)
5. T. Kuang, C. Williamson: Hierarchical Analysis of RealMedia Streaming Traffic on an IEEE 802.11b Wireless LAN. Elsevier Computer Communications, Vol. 27, No. 6 (2004)
6. P. Bellavista et al.: MUM: a Middleware for the Provisioning of Continuous Services to Mobile Users. IEEE ISCC (2004)
7. A. Dutta et al.: Fast-handoff Schemes for Application Layer Mobility Management. IEEE PIMRC (2004)
8. J. Chan et al.: A Framework for Mobile Wireless Networks with an Adaptive QoS Capability, MoMuC (1998)
9. A. Mishra et al.: An Empirical Analysis of the IEEE 802.11 MAC Layer Handoff Process. ACM Computer Communication Review, Vol. 33, No. 2 (2003)
10. Debian: Tools for Linux Wireless Extensions, <http://packages.debian.org/stable/net/wireless-tools.html>
11. P. Hsieh et al.: Application Layer Mobility Proxy for Real-time Communications. 3G Wireless and Beyond (2003)
12. R. Braden et al.: Resource ReSerVation Protocol (RSVP). IETF RFC 2205 (1997)
13. J. Rosenberg et al.: SIP: Session Initiation Protocol. IETF RFC 3261 (2002)
14. P. Bellavista et al.: Application-level QoS Control and Adaptation for Video on Demand. IEEE Internet Computing, Vol. 7, No. 6 (2003)
15. H. Velayos, G. Karlsson: Techniques to Reduce IEEE 802.11b Handoff Time. IEEE ICC, (2004)
16. P. Bellavista et al.: Mobility Prediction for Mobile Agent-based Service Continuity in the Wireless Internet. IEEE MATA (2004)
17. G. Anastasi et al.: A Power-Aware Multimedia Streaming Protocol for Mobile Users. IEEE ICPS (2005)
18. A. Dutta et al.: Implementing a Testbed for Mobile Multimedia. IEEE GLOBECOM (2001)
19. J. Hillebrand et al.: Quality-of-service Signaling for Next-generation IP-based Mobile Networks. IEEE Communications Magazine (2004)
20. A.T. Campbell et al.: Comparison of IP Micromobility Protocols. IEEE Wireless Communications, Vol. 9, No.1 (2002)
21. J. Chan et al.: A QoS Adaptive Mobility Prediction Scheme for Wireless Networks. IEEE GLOBECOM (1998)
22. D. Bruneo et al.: VOD Services for Mobile Wireless Devices. IEEE ISCC (2003)