Mobility-Aware Connectivity for Seamless Multimedia Delivery in the Heterogeneous Wireless Internet

Paolo Bellavista, Antonio Corradi, Carlo Giannelli
Dip. Elettronica, Informatica e Sistemistica - Università di Bologna
Viale Risorgimento, 2 - 40136 Bologna - ITALY
Phone: +39-051-2093001; Fax: +39-051-2093073
{pbellavista, acorradi, cgiannelli}@deis.unibo.it

Abstract

The diffusion of wireless terminals with multiple communication interfaces, e.g., IEEE 802.11, Bluetooth, and UMTS on the same device, is pushing towards the necessity of middleware solutions to dynamically and seamlessly select the proper connectivity technology to exploit at any time. That selection should consider several elements, at very different abstraction layers, from application bandwidth to energy consumption requirements, from connectivity costs to user preferences, i.e., it should be context-dependent. The paper presents our context-aware MAC middleware for multi-interface wireless terminals that enables the dynamic determination and selection of the most suitable interface and connectivity provider among the available ones. MAC novelty is primarily in two crucial challenging elements. On the one hand, it considers not only infrastructure-based connectivity providers, e.g., UMTS base stations, but also peer nodes, e.g., neighbor nodes accessible via Bluetooth and connected via Wi-Fi to the Internet infrastructure. On the other hand, MAC can evaluate both infrastructure and peer connectivity providers not only based on usual parameters such as available bandwidth and energy consumption, but also taking into account innovative and crucial indicators such as the degree of mobility, even relatively to mobile connecting clients.

1 Introduction

Two major trends are manifest in the evolution of the mobile computing area in the last decade: the growing availability of processing/memory resources at mobile nodes and the widespread diffusion of wireless communication technologies. Those trends push towards considering a larger and larger set of services, from traditional Internet applications to novel location-based services, that can be accessed by wireless clients independently of their mobility at provisioning time.

In the following, let us call *interfaces* the wireless network interfaces available at a mobile node and that the node can typically exploit to connect to the Internet infrastructure, e.g., IEEE 802.11 and Bluetooth wireless client cards. Instead, we will use the *connector* term to indicate a device, e.g., an IEEE 802.11 Access Point (AP) that provides Internet access acting as a bridge between a mobile node and the traditional fixed network infrastructure. In other words, interfaces model the wireless hardware equipment available at client side, while connectors are the entities providing real access to the fixed network via wireless communications with client-side active interfaces. We call the integrated networking scenario with fixed Internet hosts, mobile nodes with wireless interfaces, and connectors in between, as the Wireless Internet (WI).

The WI is a definitely usual deployment scenario nowadays, but forces to consider several novel challenges for service provisioning. In fact, a mobile node may move outside the coverage range of its currently used connector(s) and that could produce connectivity requests to newly discovered connectors (handover between origin and destination connectors). That requires properly handling the handover process by reducing the time needed for connector change and avoiding packet losses during that interval. Service provisioning becomes even more challenging in the case of continuous services, i.e., applications that distribute time-continuous flows of information to their requesting clients, such as audio and video streaming [1]. Indeed, the support to WI continuous services should address the very challenging issue of avoiding any temporary flow interruption during client handovers.

Several relevant solutions have recently emerged in order to support WI service provisioning [2, 3]. All these proposals must provide at least two crucial support functions: i) an evaluation procedure to quantitatively measure the suitability of available connectors, e.g., depending on currently available bandwidth and connection costs, and ii) continuity management mechanisms to optimally select when and to which connector to perform a handover while minimizing service interruptions, as better detailed in the following. In most common WI support solutions, mobile nodes can exploit only one interface at a time; handovers are horizontal, i.e., they involve origin and destination connectors exploiting the same interface (in particular, the terms used in the literature are intra-horizontal and inter-horizontal to indicate, respectively, the cases where origin and destination connectors belong to either the same or different administrative network domains). We call the above deployment scenarios *homogeneous WI*, to recall the fact that connectors exploit the same interface type in these cases.

We claim that in the near future there will be the need to consider more complex and flexible deployment scenarios than the homogeneous WI for the provisioning of continuous services. In particular, we focus on envisioned, novel, and challenging scenarios where i) mobile nodes are equipped with and able to simultaneously exploit several heterogeneous interfaces, and ii) connectors include both infrastructure-based equipment, e.g., IEEE 802.11 or GPRS APs, and mobile nodes providing Internet connectivity in a peer to peer fashion, i.e., *peer connectors*. In the following we indicate the above scenario as *heterogeneous WI*, because it simultaneously involves different wireless technologies and different types of connector.

The paper presents the primary design and implementation choices of our Mobility-Aware Connectivity (MAC) middleware for supporting continuous service provisioning in the heterogeneous WI. MAC permits to dynamically exploit at best multiple interfaces and both infrastructure-based and peer connectors, by switching among them by respecting continuity constraints during service sessions. To that purpose, MAC proposes a novel context-based connector evaluation process, specifically designed to deal with heterogeneous WI issues: it dynamically evaluates all the available connectors not only based on usual parameters such as currently available bandwidth and estimated energy consumption, but also taking into account innovative and crucial indicators such as connector mobility, even in relation with the mobile clients they are offering connectivity to. In addition, MAC provides original solutions for continuity management by adopting a proactive context-dependent approach to perform streaming pre-fetching with the needed advance time. To that purpose, MAC exploits innovative mechanisms for handover prediction, which are out of the specific focus of this paper and can be found in [4, 5].

The rest of the paper is structured as follows. Section 2 describes the needed background about homogeneous WI, by demonstrating the motivations of significantly improving/extending it towards the envisioned heterogeneous WI. Section 3 outlines the primary characteristics of our MAC middleware, while Section 4 positions our novel approach with regard to the state-of-the-art about evaluation process and peer-based connectivity. Ongoing research work and conclusive remarks end the paper.

2 Background and Motivations

As briefly stated before, it is possible to identify two major phases in a handover procedure: evaluation process and continuity management. The former is in charge of gathering information about the currently accessed connectors (and possibly the other available ones) and of evaluating their current suitability, e.g., depending on the currently provided QoS level. The latter is in charge of exploiting the evaluation process result to choose when to perform a handover and to which connector. Moreover, continuity management should provide support mechanisms for seamless handovers for continuous services, e.g., by temporarily broadcasting packets to both origin and destination connectors to minimize packet loss [6].

By considering the notable example of the widespread IEEE 802.11, the evaluation process is embedded in interface firmware and is usually based on Received Signal Strength Indication (RSSI) or Signal to Noise Ratio (SNR), monitored for both origin and destination connectors. The assumption is that lower RSSI and SNR values correspond to limited network performance. Continuity management for intra-horizontal handovers is mainly realized via AP signaling messages to update mobile node location (represented as the currently accessed AP). Continuity management for inter-horizontal handovers is not standardized, but can exploit some partial support mechanisms, either standard such as Mobile IP or special-purpose [4, 5].

Mobile IP is usually sufficient to keep connections alive while performing handovers. However, that is not enough for continuous services where the user may experience interruptions at handover occurrences, e.g., because during a handover between IEEE 802.11 APs there is a time interval in which a mobile node receive packets from neither the origin nor the destination AP. To counteract this issue, [4] and [5] propose mobile agent-based shadow proxies that run at the edges between wireless networks and the wired Internet, close to the served mobile clients, to offer adaptively-sized pre-fetched buffers depending on handover predictions.

Today the above homogeneous WI is the most spread service provisioning scenario. However, the growing resource availability on mobile nodes pushes towards the adoption of a more complex and valuable heterogeneous WI scenario characterized by client nodes with multiple heterogeneous interfaces, possibly activated simultaneously, and heterogeneous connectors, both infrastructure-based and peer ones. In fact, mobile nodes are currently equipped with several wireless interfaces with very differentiated characteristics: for instance, Bluetooth has limited coverage, medium bandwidth, limited power consumption and typically peer-to-peer connections, while IEEE 802.11 has medium coverage, larger bandwidth and connections established in either infrastructure or peer-to-peer way. In addition, the increasing capabilities of mobile nodes suggest new deployment scenarios where clients can also play the role of connectivity providers (peer connectors). For instance, a mobile node with both UMTS and Bluetooth interfaces may decide to play the role of peer connector by exploiting its UMTS interface to connect to the Internet and by offering itself as a Bluetooth modem for neighbors with Bluetooth capabilities.

Notwithstanding the complexity of supporting them, heterogeneous WI scenarios can provide several relevant benefits and advantages. First of all, peer connectors can significantly extend the connectivity opportunities for mobile clients. For instance, in an area with only cellular connectivity to the Internet, a peer connector with both UMTS and Wi-Fi interfaces can open the Internet access even to nodes with only the IEEE 802.11 interface. In addition to offering connectivity to otherwise disconnected nodes, peer connectors can provide alternative connectivity ways, possibly more suitable according to specified evaluation metrics than directly exploiting infrastructure-based connectors. For instance, a peer connector with flat-rate UMTS connectivity can offer itself as a Bluetooth modem for free in the time intervals when it is not interested in accessing Internet services. Anyway, the exploitation of peer connectors could be far more complex than infrastructure-based usage and requires novel support approaches to effectively tackle newly introduced issues: for instance, peer-based connectivity tends to be less reliable, also in the case of a non-moving client node, since peer connectors can move out of client radio range or abruptly revoke connectivity provisioning.

In such a complex and dynamic scenario the evaluation process cannot be based only on raw monitoring data from the physical layer, such as RSSI and SNR. Moreover, it is not feasible to statically determine the metric to apply once for all and to embed it in interface firmware. On the contrary, the evaluation process should also consider more expressive context information, at different levels of abstraction, to take handover decisions. Gathered context should include the static/dynamic characteristics of available interfaces and, for each interface, of the available connectors, mobility degree of peer connectors (see the following section), user preferences, node resource availability (from battery level to available memory), and application-specific quality requirements.

3 The Mobility-Aware Connectivity Middleware for Continuous Services in the Heterogeneous WI

Differently from first middleware proposals emerged for the heterogeneous WI, we have concentrated our attention to the support of continuous services. To that specific purpose, we claim the need for innovative evaluation processes with the primary goal of maximizing connectivity durability, e.g., to minimize the number of handover processes unless quality requirements cannot be met with currently used connectors. As a secondary but crucial goal, these evaluation processes should work to maximize useful bandwidth for the served services while minimizing power consumption at client nodes. In addition, for the sake of performance and scalability, evaluation metrics should be primarily based on context data that can be directly gathered at mobile nodes, with no need of continuous interaction with the network infrastructure, thus reducing the imposed monitoring overhead. About context awareness, let us notice that we do not aim to provide here a general-purpose framework for the gathering and disclosure of high-level context information. The paper only aims to identify and exploit a limited subset of relevant context information, e.g., dynamically changing degree of node mobility, to effectively address mobility-related multimedia management issues.

Based on the above guidelines, we have designed and partially implemented the Mobility-Aware Connectivity (MAC) middleware, which specifically targets the support of continuous services in heterogeneous WI scenarios. In particular, the paper presents how MAC models any possibly available connector and how it performs its original evaluation process based also on connector classification. In fact, there are two primary novelty aspects in MAC: the support to very heterogeneous type of connectors, by considering their differentiated characteristics, and the wide set of context information, at different levels of abstraction, exploited by the evaluation process. On the one hand, MAC considers both infrastructure-based and peer connectors, either fixed, e.g., an Internet-wired desktop PC working as connector for nearby Bluetooth devices, or mobile, e.g., a UMTS PDA working as connector for dynamically encountered Wi-Fi nodes while walking together in a guided downtown tour. On the other hand, MAC adopts a flexible context-dependent evaluation process that, for example, exploits different metrics depending on the considered connector type (infrastructure or peer), on client node mobility state (still or motion), and on peer connector mobility state (transient or joint), as better detailed in the following.

3.1 MAC Connectors

MAC originally supports different types of connectors and classifies them according to the taxonomy depicted in Figure 1. MAC supports infrastructure-based connectors with IEEE 802.11 and GPRS interfaces. They are considered always reliable, i.e., MAC assumes that this type of connectors always try to forward the traffic of associated clients (and are expected to succeed, apart from dependability issues due to traffic congestion) and do not endanger user privacy (no traffic auditing). Security issues are not the primary focus of the MAC research project and are not addressed in the paper. In addition, MAC supports peer-based connectors with IEEE 802.11 and Bluetooth interfaces. The reliability of these connectors is dynamically determined depending on connector runtime behavior (based on user preferences, past interaction history, and client location).

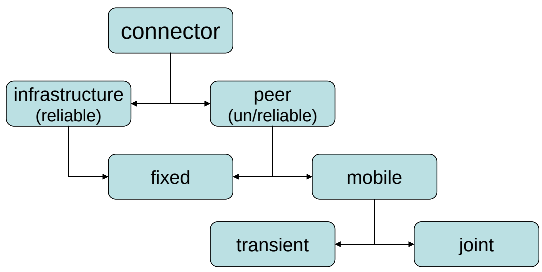

Figure 1. Types of connectors supported in MAC.

In addition to reliability, MAC considers other crucial aspects that deeply differentiate the runtime behavior of peer and infrastructure connectors. First of all and most important, MAC distinguishes fixed and mobile connectors. Understanding whether a peer connector is fixed/mobile is crucial (and a challenging issue) because it directly impacts on the stability of offered connectivity: mobile peers usually become unavailable with higher probability because more easily they can exit the client radio range.

The classification of connectors into reliable/unreliable and fixed/mobile sub-classes is a rather static decision, expected not to change during a service session. MAC evaluates the above classification for the eligible connectors of a given client only at the beginning of its service session. In addition to that static context, MAC also considers an innovative and highly dynamic context indicator: the mutual degree of mobility between a mobile client and the associated mobile peer connector. In particular, MAC classifies mobile connectors as either *transient* or *joint*, depending on the fact that, respectively, the connectors move with either different or the same speed (in both module and direction) of the associated mobile client. The transient/joint sub-class obviously depends on mobility behaviors at runtime and its correct dynamic determination/update is a key point for MAC effectiveness in terms of limited overhead and durable evaluation process results. In fact, MAC supports the exploitation of both transient and joint mobile connectors, but transient ones usually have higher probability of becoming rapidly unavailable, e.g., because the transient mobile peer is a PDA carried by a user walking on the same sidewalk with opposite direction. On the contrary, joint mobile peers, such as a PDA connector sharing the same train wagon with its client, can probably provide a more suitable connectivity offer with greater durability.

3.2 MAC Evaluation Process

The MAC evaluation process exploits both static context, e.g., average power consumption and interface nominal bandwidth, and dynamic context, e.g., mobility state of clients and peer connectors. Those context data are effectively monitored and taken into account to provide the best tradeoff between connection durability and quality for the set of applications running at a client node, as better detailed in the following.

The MAC evaluation process is based on the home-office-home mobility pattern assumption. Users are expected to move among well-known locations, by alternating movement and still phases. In particular, users usually perform still phases in well-known locations, e.g., at the office, at home, or at the coffee machine, while perform movement phases as long as they depart from a well-known location and arrive to another well-known location, e.g., while walking in the corridor between an office and a classroom.

In many everyday-life situations, user requirements and surrounding environment vary in relation to user mobility state, either *still* or *motion*. For instance, a client in the still phase is expected to find infrastructure connectors available, with good quality and short/medium coverage range, e.g., IEEE 802.11 APs. In addition, while still, clients tend to require high reliability. On the contrary, in the motion phase it is possible to assume only the presence of infrastructure connectors with large radio coverage, e.g., UMTS base stations. Moreover, users in motion usually decrease their network performance requirements, e.g., because they only perform background network activities such as email transferals. Therefore, we claim the suitability of middleware solutions for the heterogeneous WI that dynamically change their evaluation process also depending on user mobility state.

MAC thoroughly follows that design guideline. When a user is fixed (*still* state), MAC assumes that there is high probability of the availability of limited/medium coverage range infrastructure-based connectors (Bluetooth and IEEE 802.11) or fixed peer-based connectors (such as the nearby desktop PC of a user currently working on her PDA). In both cases, the primary objective of enduring connections is easily achieved. In fact, these types of connector usually provide sufficient QoS, are reliable and inexpensive. MAC should simply select the connector belonging to the above types that maximizes bandwidth while minimizing power consumption. Connectors of the above types should in any case be preferred to both GPRS/UMTS base stations (with variable quality and often non-negligible costs) and mobile peers (unreliable).

On the contrary, when a user is in motion (*motion* state), it is reasonable to first choose infrastructure-based connectors with medium/large radio coverage, thus increasing the probability of connection durability. When infrastructure connectors are not available, MAC also considers mobile peer connectors with joint mobility with regards to the requesting clients. If even joint connectors are not available, MAC additionally explores the opportunities offered by transient connectors with at least medium range wireless coverage, e.g., IEEE 802.11 mobile peers. In fact, transient peers with limited coverage range should be avoided, if possible, because their exploitation may produce frequent hand-overs, thus introducing non-negligible overhead and power consumption. Table 1 summarizes the above considerations, depending on the motion/still state of client nodes.

Table 1. The MAC evaluation process depends on client motion/still state.

Client mobility state	Wireless interface	Connector type
Still	Bluetooth ≈ 802.11 >> UMTS	fixed connector >> joint peer
Motion	802.11 ≈ UMTS >> Bluetooth	infrastructure ≈ joint peer >> transient peer

(≈ equivalent, > better, >> much better)

Based on these guidelines, given a client, MAC selects one type of connector, if any available, and limits its evaluation process to the set of connectors of that type. MAC quantitatively evaluates every eligible connector by determining *ConnectorValue* in the [0, 1] range (0=worse choice, 1=best choice) for each of them. To that purpose, MAC exploits the evaluation function:

$$ConnectorValue = EnduranceValue + QualityValue$$

where *EnduranceValue* estimates the expected connector durability and *QualityValue* its expected quality in terms of bandwidth and energy consumption.

The evaluation of each addend dynamically changes depending on the fact that MAC considers the connector either fixed or mobile. For a fixed connector:

$$EnduranceValue = CMob \bullet Range$$

$$QualityValue = (1 - CMob) \bullet (\alpha \cdot Bandwidth + \beta \cdot Energy)$$

while in the case of a mobile connector:

$$EnduranceValue = (1 - Joint) \bullet Range$$

$$QualityValue = Joint \bullet (\alpha \cdot Bandwidth + \beta \cdot Energy)$$

CMob and *Joint* are values in the [0, 1] range and model the MAC estimates for, respectively, client and connector mobility. In particular, *Joint* represents the probability of connector movement with regards to the associated client. MAC determines the original *CMob* and *Joint* indicators as detailed in Section 3.3. The *Range* parameter (always in the [0, 1] interval) concisely models the radio coverage of a connector and only depends on the associated interface. For instance, MAC associates IEEE 802.11 AP connectors with a *Range* value of 0.7 while *Range* for Bluetooth peer connectors is 0.3. *QualityValue* functions for both fixed and mobile connectors consider the *Bandwidth* and *Energy* parameters (always in the [0, 1] interval), which model the ratio between, respectively, available/required bandwidth and available/required energy. α and β ($\alpha, \beta \geq 0$ and $\alpha + \beta \leq 1$) concisely model user-specified application-level requirements, thus adapting the relative relevance between bandwidth and energy user preferences.

Let us rapidly note that the second equation pair does not include *CMob*: in fact, based on home-office-home assumption, MAC considers mobile peers only if there are not fixed connectors in visibility; with no fixed connectors in visibility, it is impossible to estimate *CMob* (see the following). Finally, reliability and economic costs are not included in the above evaluating functions because MAC automatically excludes connectors not compliant with user requirements about reliability/cost from the set of eligible connectors.

3.3 MAC Context Gathering

To apply the above evaluation function, MAC has to gather several context data at different levels of abstraction. To that purpose, MAC requests users to express their application-level requirements at registration time. User requirements are assumed not to change during a service session and include: bandwidth, energy consumption (power saving or maximum performance), maximum affordable cost, and required reliability. In addition, MAC gathers context information at lower layers of abstraction. First of all, MAC estimates *CMob* in terms of probability for the client node to be in motion. Then, for each eligible connector, MAC monitors context information to understand whether the connector is infrastructure or peer, fixed or mobile, reliable or unreliable, free or charged, which is its *Joint* indicator, the average coverage range, the nominal bandwidth, and the average power consumption (the last three indicators are simply inferred from the associated interface type).

As already stated, MAC can collect both static and dynamic context information. Here, for the sake of briefness, we only focus on how MAC obtains the most interesting and challenging dynamic context indicators, i.e., *CMob* and *Joint*. To estimate these values, MAC monitors the execution environment and collects RSSI data about any eligible connector.

Delving into finer details, for each interface MAC determines the list of available connectors and collects RSSI sequences for each connector. Then, for each fixed (mobile) connector *CMob* (*Joint*) is set linearly depending on the variability of the RSSI sequence for that connector. To estimate RSSI sequence variability, MAC monitors the evolution of the module of the first harmonic of RSSI sequences, obtained via the application of the Discrete Fourier Transform (DFT) to 4s-long RSSI sequences. The DFT adoption permits MAC to low-pass filter RSSI sequences by effectively reducing RSSI variability due to signal noise, in order to identify only RSSI modifications due to actual mobile node movements.

We have experimentally validated how *CMob* and *Joint* depend on RSSI variability and the values used in MAC are the result of these experimental evaluations. For instance, only to mention some practical configuration examples, our experiments in IEEE 802.11 testbed environments suggest us to set *CMob* to 0 when the first harmonic module is <= 1.3, to 1 when the module is >= 2.3, to a linearly dependent intermediate value otherwise. Further details about how MAC gathers and models RSSI sequences and detailed experimental results of filtering procedures are out of the scope of the paper and can be found in [5]. Additional information about MAC implementation, the downloadable code of the mobility estimator prototype, and extensive performance results obtained via both in-the-field and simulation tests are available at:
`http://lia.deis.unibo.it/Research/MAC/`

About gathering IEEE 802.11 raw RSSI data, due to the relative novelty and high heterogeneity of wireless technologies, there is currently no standard specification, widely accepted by vendors and available in most common operating systems, of an application-level API to achieve full RSSI visibility. To achieve maximum portability, MAC includes different implementations of RSSI monitoring mechanisms, which are automatically selected depending on the specific characteristics of served client devices. In particular, for Linux client nodes MAC exploits an original and operating-system-transparent Java API to obtain the list of all APs in current visibility and their related RSSI information. MAC accesses those monitoring data via the standardized virtual directory `/proc/net/wireless` provided by the Linux Wireless Extensions [7]. For Windows CE /.NET clients, instead, MAC transparently binds the same Java API to the monitoring functionality provided by the Microsoft Network Driver Interface Specification User-mode I/O (NDISUIO), which is platform-dependent but portable among different wireless interface implementations [8]. For instance, MAC exploits the NDISUIO function `DeviceIOControl()` to query the `OID_802_11_`

BSSID_LIST_SCAN object to retrieve the complete list of currently reachable IEEE 802.11 APs.

Let us note that MAC only performs local monitoring at client nodes. In that way, it achieves a twofold benefit. First, MAC exploits only local information that is available despite mobile clients are currently connected to the heterogeneous WI. Then, it does not require any external special-purpose support component, e.g., monitoring components working on the infrastructure side, thus enabling the potentially immediate MAC adoption in any heterogeneous WI scenario by only deploying MAC components at mobile nodes.

However, even the only local monitoring of network interfaces is a power consuming process [9]. Therefore, to minimize power consumption, MAC performs "aggressive" context gathering only when required (client in *research* state), while performing "lazy" monitoring otherwise (client in *connected* state). In other words, a mobile node in research state requires to get new connectors as soon as possible, e.g., because the *ConnectorValue* of a connector in use is decreasing or that connector has abruptly become unavailable; in that case MAC performs aggressive monitoring. For a client in connected state, instead, MAC reduces monitoring overhead only to understand if relevant events occur, e.g., by simply estimating *CMob* via the observation of RSSI variations of eligible fixed connectors. More precisely, MAC performs aggressive monitoring only when i) a connector currently in use becomes unavailable or does not meet anymore user-specified quality requirements, or ii) client state passes from still to motion or vice versa.

To understand whether a client is in the still/motion and research/connected states, MAC exploits *CMob* monitoring and its time evolution, as depicted in Figure 2. MAC switches the client state from still to motion whenever *CMob* becomes greater than 0.6, while it performs the inverse switch when *CMob* passes below the 0.4 threshold. The adoption of different thresholds for the two state transitions has been decided to prevent from bouncing effects. In fact, frequent switching between still and motion states would impose repeated perturbations in connector selection, by possibly causing frequent and expensive connector changes and by consequently degrading continuous service quality. By focusing on the research state, in this case MAC monitors the RSSI of fixed connectors to infer *CMob*. In particular, when the state is research&still, MAC estimates the *ConnectorValue* for fixed connectors; when the state is research&motion, it also monitors the RSSI of mobile connectors to infer their *Joint* and *ConnectorValue* indicators.

On the contrary, MAC performs lazy monitoring for clients in connected state. When a client is connected&still, MAC monitors only the RSSI of connectors currently in use, which are likely to be fixed, to evaluate both *CMob* and *QualityValue*. When the state is connected&motion, instead, MAC gathers context information about connectors in use only if they are fixed; if there are mobile connectors in use, for them MAC evaluates *Joint* and *QualityValue* indicators. In the latter case, MAC should also monitor the RSSI of possibly available fixed connectors to infer *CMob*; if there are no fixed connectors, *CMob* is set to 0.9.

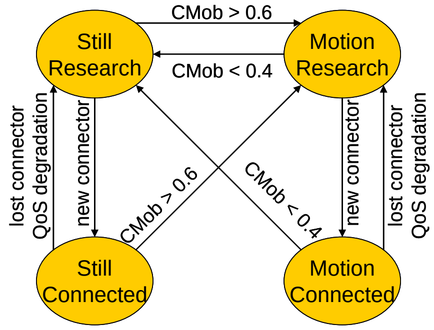

Figure 2. MAC adapts context monitoring based on a state transition algorithm for client nodes.

4 Related Work

Some recent research contributions have proposed evaluation metrics for heterogeneous wireless technologies. The evaluation process function in Terminal Management System (TMS) defines quality and cost indicators for eligible connectors depending on a user-specified priority order among connectivity providers and interfaces [10]. TMS applies its evaluation function to any available connector: each function term is weighted according to a weight set based on user-specified priorities, which can change at service provisioning time. Also the Vertical Handoff Decision Function (VHDF) provides users with the capability to specify a priority order among different network characteristics, by defining a proper weight set [11, 12]. The VHDF evaluation function exploits the weights to calculate a linear combination of current network conditions, network performance, service cost, power requirements, security, and pro-activity of the exploited handover process. Both TMS and VHDF neither consider peer connectors nor address the challenging issues stemming from node mobility.

Policy-based system to ROam Transparently among Overlay Networks (PROTON) is a more sophisticated evaluation solution based on policies, i.e., event-condition-action rules that specify the management actions to execute whether defined conditions apply, with condition evaluation triggered by event notification [13]. In particular, PROTON permits to define tautness functions that evaluate how tautly a condition fits to an event: the closer is the returned value to 0, the tauter a condition is to a specific event. However, the relevant flexibility of the policy-based PROTON evaluation process is achieved at the expenses of the complexity and intrusiveness of the support system, which require for instance the deployment of special-purpose components on both the infrastructure- and the client-side, by imposing an additional non-negligible monitoring overhead. Let us observe that MAC has decided to limit the flexibility of its evaluation function and not to be policy-based in order to avoid that monitoring overhead: MAC only evaluates the suitability of available connectors at the beginning of each service session and when client nodes are in the research phase state, also depending on their degree of mobility.

Only a few middleware solutions have already considered the challenges of exploiting both infrastructure and peer connectors. [14] proposes a two-hop-relay architecture, based on Relay Gateway (RG) nodes, that can behave both as mobile clients and as cellular gateways. They can seamlessly switch interfaces depending on network availability. In addition, they can improve wireless LAN coverage by exploiting cellular interfaces where wireless LAN connectivity is not available. Non-RG client nodes close to RGs may ask for connectivity through them, in a non-transparent way. A similar approach based on two-hop paths is presented in [15]. Other solutions propose more flexible and complex multi-hop organizations for peer connectors. For instance, in [16] a peer connector, called Proxy Client, can interwork with both cellular and IEEE 802.11 ad-hoc networks. Differently from [14, 15], in [16] mobile nodes can interact with Proxy Clients not only directly but also via another intermediate layer of peer connectors, namely Relay Clients, in a multi-hop ad-hoc manner. If compared with MAC, these proposals do not provide any quantitative evaluation metric that takes into account specific characteristics of peer-based connectivity, such as the joint mobility degree.

5 Conclusions

The advances in device miniaturization and wireless communications are pushing the migration towards the heterogeneous WI, where both evaluation process and continuity management are crucial to fully take advantage of the novel opportunities offered by this challenging deployment scenario. The paper demonstrates the need for novel middleware supports that evaluate interface and connector suitability considering not only traditional parameters but also more expressive context information, adaptively gathered only when strictly required in order to reduce monitoring overhead. MAC shows the viability of the adoption of middleware solutions for service continuity in the heterogeneous WI, by pointing out the need to consider also the connector type and the mobility degree of clients and peer connectors in order to establish a proper evaluation metric.

The encouraging results obtained by the first MAC prototype are encouraging and further stimulating our on-going research activities. In particular, we are currently evaluating the MAC performance over a wide deployment scenario with dozens of Wi-Fi/Bluetooth infrastructure-based and peer connectors, to validate our middleware capability to support continuous services with even strict handover requirements, such as multimedia streaming. In addition, we are extending the current MAC prototype to include the support to additional interfaces such as UMTS and WiMax.

Acknowledgements

Work supported by the MIUR PRIN MOMA and the CNR Strategic IS-MANET Projects.

References

[1] P. Ramanathan, K.M. Sivalingam, P. Agrawal, S. Kishore, “Dynamic Resource Allocation Schemes during Handoff for Mobile Multimedia Wireless Networks”, *IEEE Journal on Selected Areas in Communications*, Vol. 17, No. 7, pp. 1270-1283, Jul. 1999.

[2] V. Ghini, P. Salomoni, G. Pau, "Always-best-served Music Distribution for Nomadic Users over Heterogeneous Networks", *IEEE Communications*, Vol. 43, No. 5, pp. 69-74, May 2005.

[3] Yi Pan, M. Lee, J.B. Kim, T. Suda, "An End-to-end Multipath Smooth Handoff Scheme for Stream Media", *IEEE Journal on Selected Areas in Communications*, Vol. 22, No. 4, pp. 653-663, May 2004.

[4] P. Bellavista, A. Corradi, C. Giannelli, "Mobile Proxies for Proactive Buffering in Wireless Internet Multimedia Streaming", *IEEE Int. Conf. Distributed Computing Systems (ICDCS) Workshops*, pp. 297-304, June 2005.

[5] P. Bellavista, A. Corradi, C. Giannelli, “Adaptive Buffering based on Handoff Prediction for Wireless Internet Continuous Services”, *Int. Conf. High Performance Computing and Communications (HPCC)*, LNCS 3726, pp. 1021-1032, 2005.

[6] N. Shenoy, R. Montalvo, "A Framework for Seamless Roaming across Cellular and Wireless Local Area Networks", *IEEE Wireless Communications*, Vol. 12, No. 3, pp. 50-57, June 2005.

[7] Debian - Tools for Manipulating Linux Wireless Extensions, http://packages.debian.org/stable/net/

[8] Microsoft Software Developer Network - NDIS Features in Windows CE, http://msdn.microsoft. com/library

[9] E. Ferro, F. Potorti, "Bluetooth and Wi-Fi Wireless Protocols: a Survey and a Comparison", *IEEE Wireless Communications*, Vol. 12, No. 1, pp. 12-26, Feb. 2005.

[10] E. Adamopoulou, K. Demestichas, A. Koutsorodi, M. Theologou, "Intelligent Access Network Selection in Heterogeneous Networks - Simulation Results", *Int. Symp. Wireless Communication Systems (ISWCS)*, pp. 279-283, Sept. 2005.

[11] A. Hasswa, N. Nasser, H. Hossanein, "Generic Vertical Handoff Decision Function for Heterogeneous Wireless", *IFIP Int. Conf. Wireless and Optical Communications Networks (WOCN)*, pp. 239-243, Mar. 2005.

[12] N. Nasser, A. Hasswa, H. Hassanein, "Handoffs in Fourth Generation Heterogeneous Networks", *IEEE Communications Magazine*, Vol. 44, No. 10, pp. 96-103, Oct. 2006.

[13] P. Vidales, et al., "Autonomic System for Mobility Support in 4G Networks", *IEEE Journal on Selected Areas in Communications*, Vol. 23, No. 12, pp. 2288-2304, Dec. 2005.

[14] Hung-yu Wei, R.D. Gitlin, "Two-hop-relay Architecture for Next-generation WWAN/WLAN Integration", *IEEE Wireless Communications*, Vol. 11, No. 2, pp. 24-30, Apr. 2004.

[15] A. Lera, A. Molinaro, S. Polito, G. Ruggeri, "End-to-end QoS Provisioning in 4G with Mobile Hotspots", *IEEE Network*, Vol. 19, No. 5, pp. 26-34, Sept.-Oct. 2005.

[16] H. Luo, R. Ramjee, P. Sinha, L.E. Li, S. Lu., "UCAN: a Unified Cellular and Ad-hoc Network Architecture", *ACM Mobicom*, pp. 353-367, 2003.