This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the WCNC 2008 proceedings.

Architectures for Seamless Handover Support in Heterogeneous Wireless Networks

Marco Miozzo*, **Michele Rossi***† and **Michele Zorzi***†
*Consorzio Ferrara Ricerche, via Saragat 1, 44100 Ferrara, Italy
†DEI, University of Padova, Dept. of Information Engineering, via Gradenigo 6B, 35131 Padova, Italy
Email: {mzzmrc}@unife.it {rossi,zorzi}@dei.unipd.it

***Abstract*—In this paper we study the performance of the Ambient Networks (AN) access selection architecture. We consider heterogeneous wireless networks, where mobile terminals (MTs) own multiple radio technologies and need to remain connected while on the move. We would like to provide the MTs with seamless IP services, such as video/audio streaming, so that changes in their point of attachment will not affect the experienced streaming quality. In the first part of this paper, we present the AN architecture for access selection, along with its functional entities (FEs), their interrelations and the algorithms that are to be run within each FE. Hence, we describe our ns2 simulation framework and detail two simulation scenarios. We finally discuss, through extensive simulation results, the effectiveness of the AN architecture in reaching the above goals.**

I. INTRODUCTION

Advances in micro-fabrication, wireless technologies and embedded micro-processors are enabling a massive integration of different radio technologies in the same handheld device. In addition, multiple technologies often coexist within the same geographical area: cities are now covered with cellular networks (e.g., GSM and UMTS), WiFi hot-spots are becoming common and WiMAX is also expected to be widely deployed in the near future. This heterogeneity of accesses will empower services and create new business models and opportunities.

From a technical standpoint, exploiting these opportunities demands the development of novel algorithms and architectures. An open problem in this sense is the efficient exploitation of such a rich population of wireless accesses. A possible solution would be to take advantage of the growing capabilities and computational power on today's wireless terminals to decentralize the decision making procedures involved in the access selection. In fact, there are reasons for (partially) incorporating access selection algorithms within the mobile terminal. First, it is aware of the radio technologies in its physical surroundings, second, it can monitor the perceived quality for each available access and, third, it is aware of the quality currently experienced by the running applications as well as of the user's profile. In this paper, we partially address these issues by presenting (and validating) a framework for joint handover execution and access selection in multi-technology wireless networks. Our primary goal is to support mobile terminals (MTs) such that they can move seamlessly across networks, maintaining a given target quality of service (QoS). This should happen irrespective of their location as well as of the radio technology in use. The philosophy we adopt in doing so is that of the Ambient Networks project (AN) [1], which consists of enhancing cooperation among networks, while completing them with the intelligence needed to make complex decisions. According to the terminology used in this project, one of the key ingredients to enable such a cooperation is the so called *composition* [2]. Composition can occur among networks or between networks and MTs. It uniforms the communication across the various interfaces of the architecture (at any level of abstraction, e.g., between users and operators) and the functionalities needed to negotiate QoS and establish business relations. Fundamental to this framework is the Composition Agreement (CA), which is used to define the cooperation among different actors (i.e., networks providers, service providers, users and "third parties" such as traffic aggregators). CAs can be used to modify business relations with access providers (APs), change (and set up) users' security profiles and their QoS requirements. In summary, this framework allows the integration between networks and terminals so that MTs can fully exploit the functionalities offered by the visited networks. In its simplest form, CA may be seen as a message exchange between MT and network which is executed on the fly as the former joins the latter. This exchange may involve, in order, network discovery, establishment of basic connectivity, negotiation of connection parameters (including link layer, security, billing and QoS requirements) and setup of network and transport layers. Hence, composition enhances the integration among networks via negotiations at potentially all layers.

Previous work on heterogeneous access selection can be found in a number of papers [3]–[5]. In [3] the authors study QoS provisioning in *always best connected* (ABC) networks. According to their view, the IP protocol will have a determinant role in supporting end-to-end paths inclusive of fixed and wireless networks. The focus in [3] is more on architectural issues, especially focusing on how IP can be used to connect all the pieces of the architecture together. The management of channel access and link layer issues is, however, left to the wireless islands placed at the border of the network, which usually provide direct access to MTs. The work in this paper is complementary in nature as we study architectures and algorithms to efficiently implement access selection in these islands. The work in [4] gives a comprehensive treatment of the mechanisms that are to be orchestrated for the realization of handovers at the IP level in heterogeneous and wireless networks. Finally, reference [5] evaluates the load balancing performance of several access schemes for the provisioning of voice over IP in heterogeneous wireless networks.

In this paper we continue this line of research through the evaluation of the handover performance of an Ambient Networks architecture in multi-technology environments. Towards

This work has been supported by the WWI Ambient Networks Project. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the Ambient Networks Project.

1525-3511/08/$25.00 ©2008 IEEE

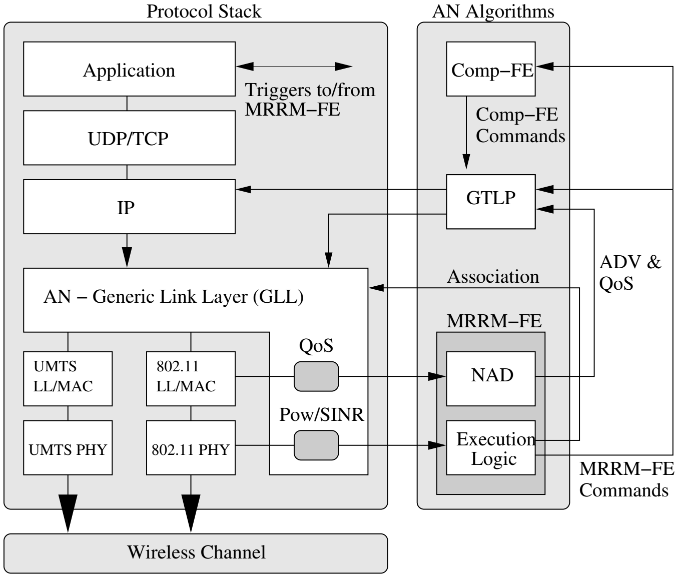

Fig. 1. Diagram of the implemented access selection architecture.

this end, we extended the Network Simulator 2 through the design of a Multi InteRfAce Cross Layer Extension (MIRACLE) [6]. In fact, the standard distribution of ns2 does not natively support multiple radio interfaces. MIRACLE enables their integration in the simulator and, at the same time, offers an efficient and embedded engine for handling cross-layer messages. This is achieved, as much as possible, through the exploitation of the original ns2 core and modules. The reader is referred to [7] for a detailed description of the simulator.

The rest of the paper is organized as follows. In Section II we give an overview of the Ambient Networks architecture for access selection. In Section III we describe the algorithms which are implemented within the architecture and their interactions. In Section IV we detail the simulator scenarios and we subsequently report simulation results that demonstrate the effectiveness of our approach. Finally, in Section V we draw our conclusions.

II. AMBIENT NETWORKS ARCHITECTURE

In this section we present the AN architecture for access discovery, selection and handover management. The inherently dynamic nature of the network scenario under analysis introduces quite a few challenges for the support of MTs. In fact, we must guarantee a seamless mobility management through a number of endpoints of possibly different administrative and technological domains. Several functional entities (FEs) are defined within the architecture to make decisions at different layers of the protocol stack. These FEs embed the actual system intelligence and communicate with each other through dedicated control channels. In this paper we only consider those FEs covering a main role in the access selection process. It should be observed that many other FEs have been defined for the support of services such as, e.g., security, billing, etc. The reader is referred to [1] for a comprehensive treatment of the matter.

Composition FE (Comp): the role of this functional entity is to dynamically manage the composition agreements (CAs) between MTs and the network operator. Such a management involves the negotiation, the realization and the deletion of CAs. Once a CA is established among two entities, they are in the same ambient networks control space (ACS). This implies their full integration in terms of service provisioning.

Generic Link Layer FE (GLL): GLL is an adaptation layer. Its main role is to enhance technology specific link layer (LL) and physical layer (PHY) functionalities. Also, it provides upper layers with a technology independent interface towards underlying modules while expressing the measured link quality in an abstract and general manner. This allows a fair comparison of different radio technologies. For instance, GLL may be in charge of monitoring LL/PHY performance and modifying several parameters such as, e.g., size of retransmission buffers, retransmission limit, FEC type at LL and signal strength, PHY-layer FEC and modulation at PHY. In addition, GLL sends triggers [8] to the execution logic (i.e., to MRRM, to be defined below) as QoS requirements are not met and establishes link layer connectivity upon request.

Multi Radio Resource Management FE (MRRM): we may consider it as the heart of the AN architecture. It provides an advanced joint management of radio interfaces. It keeps track, for each available interface/channel, of the perceived QoS, link performance and energy consumption. This information is then related to the user profile (as well as to the constraints defined within the ACS about, e.g., operator and user policies) in order to make access selection decisions. MRRM functionalities may be distributed for improved performance, e.g., across MTs and the operator. The MT is in fact aware of its own perceived QoS, whereas the operator has information about its traffic load, the load of neighboring cells, etc. The role of the network advertisement and discovery (NAD) subblock is twofold. The NAD residing in the MT collects advertisements from the access points (APs) within radio range. The NAD on the network side sends advertisements on a periodic basis. These may contain information about the load of the corresponding AP, its security features, billing, etc. Additionally, the NAD can send network discovery messages, even though this functionality is not studied in this paper.

Generic Transport Layer Protocol (GTLP): this is the transport protocol used to reliably connect FEs across different systems of the same ACS, i.e., to exchange messages between peer FEs.

A sketch of the AN architecture we implemented in the simulator is given in Fig. 1. The reader is referred to [7] for a simulation oriented description of the architecture and to [1], [9] for further details on the AN access selection framework.

III. DISCUSSION ON THE HANDOVER PROCEDURE IN THE AN ARCHITECTURE

This paper focuses on the evaluation of the handover performance of the AN architecture. In particular, we want to understand whether seamless handovers can be supported, despite the possible burden of control messages and the lack of reactivity due to network congestion/channel impairments.

The main idea in our framework is to hide, as much as possible, the attachment procedures needed for an MT to change its point of attachment to the network. Note that setting up a new connection involves the execution of several basic operations, most of which are technology dependent (i.e., WiFi Connectivity in IEEE802.11, UTRAN Packet-Switched setup in UMTS, IP connectivity, etc.). Moreover, for a seamless service during handovers, flow management functionalities should also be exploited so as to provide the necessary changes in the flow path from the old AP to the new one. Finally, MRRM should trigger the handover execution well before the quality of the current connection drops below a certain threshold. In fact, seamlessly moving across points of attachment requires that the old connection remains active during the establishment of the new link.

The realization of the above goals entails the evaluation of all available radio accesses and the estimation of the corresponding expected QoS at the MAC, link and physical layers. This is done by the GLL, which periodically collects packet error rate (PER) statistics. Instantaneous PER values are subsequently processed through, e.g., a sliding window averaging procedure, and triggers are sent to MRRM when the averaged PER falls below a given reference value, PER_{th}. In our simulations we consider $PER_{th} = 10^{-3}$, as we found it adequate for a make before break handover approach. We note that the quality of the currently used (active) radio interface can be reliably measured as we can monitor the performance (in terms of jitter, errors, etc.) of the flows we transmit over it. The quality of other radio accesses should instead be estimated from incoming advertisements or monitoring the status of control channels (when available). This may lead to a less reliable estimation of the expected QoS.

In addition to the above measurements, the quality perceived at the application layer (APP) is also monitored. In our simulations, the application layer sends a specific trigger to the MRRM when it detects an insufficient QoS. This, as shown later in the simulation results, helps MRRM in making correct handover decisions. We found this type of trigger particularly effective when the network (or the AP) is overloaded, i.e., the signal quality is good but the actual bandwidth available to the user is insufficient, given his flow requirements. In our specific simulations, we consider IP video-streaming services in downlink. Accordingly, the application layer monitors errors in the received packet sequence as well as its jitter performance.

Thanks to the triggers described above and to the AP advertisement mechanism, MRRM should have enough information to predict when it is necessary to either change the point of attachment to the network or to use a different radio interface.

Next, we describe the access selection policy we implemented at the MRRM. We kept the handover rules as simple as possible in order to isolate the effect of the various protocol components. Simulation results for this policy are given in the next section. In detail, as MRRM receives a trigger, from either GLL or APP, it starts evaluating all the discovered APs still within radio range (i.e., collecting fresh link statistics). This trigger is usually due to an unacceptable drop in the current QoS. After this, the AP with the highest effective bandwidth is selected as the new point of attachment. For effective bandwidth we mean the raw LL bandwidth scaled down to account for protocol overhead and packet error rate at the link layer. Both raw bandwidth and error rate are determined from sliding window averaging. This gives good estimates for the actual bandwidth provided by the new AP. Unfortunately, estimation of the jitter performance towards this AP is not possible at the MT due to the absence of active flows using the corresponding link at this time. This aspect could be enhanced by a network-assisted estimation, which we however did not implement in our current simulation setup. A time hysteresis threshold is used to protect the MT against ping-ponging between APs.

Upon selecting a new AP, the MRRM starts the attachment procedures. In order, these involve the (technology dependent) establishment of basic link connectivity, the AN attachment procedure (ANAP) and the establishment of a composition agreement (CA). ANAP is used to join the operator's AN control space (ACS) [1] and consists of a 4 way handshake, whereas CA implies, in the best case, a 6 way handshaking [2]. The actual handover between the current and the selected AP may only start upon completion of the CA. From this description of the AN handover process, it is clear that seamless connectivity is only possible if the time taken to complete the above steps is sufficiently short. This depends on the goodness of the predictions made by the MRRM, on the user's speed and on the promptness of the new AP in serving the new user. In the following Section IV, we evaluate the above handover mechanism in a scenario with multiple technologies and APs.

IV. PERFORMANCE EVALUATION

In this section we evaluate, through ns2 simulations, how the AN access selection architecture performs in the presence of multiple access technologies as well as user mobility. Two legacy cases, to be detailed below, are also defined for comparison. Towards this end, we extended standard ns2 thanks to the MIRACLE framework [6], [7].

Next, we present the general performance evaluation setup. We accurately simulate two radio technologies, namely, IEEE 802.11g and UMTS, modeling the behavior of the wireless channel as well as the multi-user interference (in terms of signal to interference plus noise ratio, SINR), which is used to assess the correctness of each packet we receive for both wireless technologies. This modeling is very important as it allows a precise evaluation of the network capacity. See [7] for further details on the structure of the simulator. MTs at the beginning of the simulation are randomly placed within a simulation area of 300×300 square meters. A single UMTS AP is placed in the center of this area and provides coverage to all nodes. The position of IEEE 802.11 APs depends on the simulation scenario under investigation, as we detail shortly. All MTs are equipped with both radio technologies and move according to the Gauss-Markov mobility model [10], which is configured to mimic a pedestrian and a moderate mobility behavior (the model's speeds are 2 and 15 km/h, respectively, and the α parameter is 0.8). The application running at the MTs is an H.264 video streaming. The source of the streaming flows is a so called mobile network operator (MNO), placed in the fixed Internet portion of the network. MNO is connected to UMTS and 802.11 APs via error free wired links having a fixed delay of 200 ms. Users' data flows in the downlink direction (APs → MTs), whereas standard (e.g., AP association, ARP, etc.) and AN signaling messages use both uplink and downlink channels. Next, we detail the two simulation scenarios we used for our performance evaluation.

A. Simulation Scenario 1

This scenario is considered to measure the performance improvements with respect to a simple legacy access selection

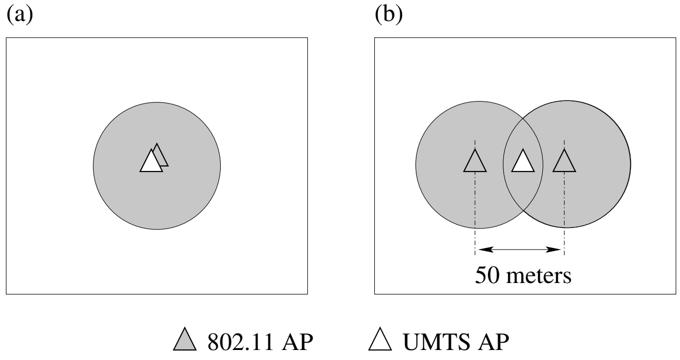

Fig. 2. Simulation scenarios.

strategy. In addition to what we introduced above, a single IEEE 802.11 AP is placed in the center of the simulation area, providing coverage to the MTs located within a distance of about 100 m (see Fig. 2(a)). According to the legacy policy, the point of attachment to the network is selected at the beginning of the connection and kept unchanged for its full duration. In particular, the AP providing the highest effective bandwidth is picked as the serving AP. With this scenario we would like to emphasize the advantages provided by the selection of the technology at runtime, through the continuous assessment of the available accesses.

B. Simulation Scenario 2

In addition to the above general settings, two IEEE 802.11 APs are placed so that their coverage areas partially overlap as we show in Fig. 2(b). This introduces some additional diversity in the available points of attachment. The legacy access selection mechanism works as follows. As for the previous simulation scenario, MTs pick their serving AP at connection setup time. Subsequently, MTs can only perform intra-technology handovers, i.e., transitions between UMTS and IEEE 802.11 are not allowed. With this scenario we would like to measure the performance improvements made possible by inter-technology handovers.

C. Simulation Results

The results reported in the following graphs are obtained averaging over 100 simulations for each simulation point we show in the figures. This gives sufficiently tight confidence intervals.

As a first result, in Fig. 3 we show the packet error rate (PER) at the application layer for simulation scenarios one (SC1) and two (SC2). Application layer packets are 148 bytes long, the user's speed is 2 km/h so as to mimic a pedestrian mobility behavior and the streaming data rate is 64 kbps. The AN access selection architecture provides better PER performance at all system loads (number of MTs). For simulation scenario number one, this gain is due to the possibility of dynamically changing the point of attachment (UMTS → 802.11 and vice versa) during the service. This gives a substantial gain which corresponds to the difference between the two curves with square bullets (SC1 legacy vs. SC1 AN). Also, the performance of both access selection techniques further improves for scenario two, i.e., considering an additional 802.11 AP. For the legacy scheme, such a performance gain is due to the improved access selection rule (i.e., intra-technology handover). In the AN case, instead, it is due to the AN capacity of effectively exploiting the additional access point.

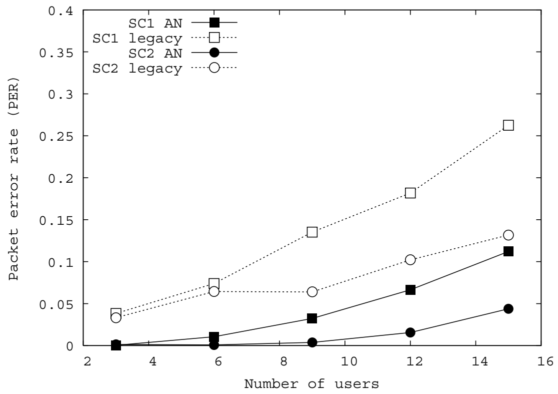

Fig. 3. PER in the two simulation scenarios (64 kbps, 2 km/h).

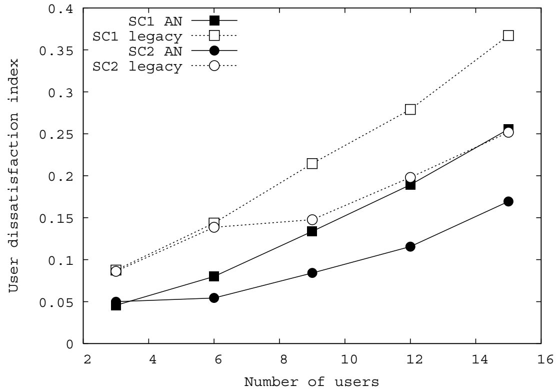

Fig. 4. User dissatisfaction index in the two simulation scenarios (64 kbps, 2 km/h).

In Fig. 4 we look at the *user dissatisfaction index*. This metric is defined as the fraction of time during which the QoS perceived by the application is below certain quality requirements. In detail, we say that a user is dissatisfied when at least one of the following conditions is verified: 1) the instantaneous inter-packet delay (jitter) at the application layer is longer than 100 ms and, 2) the data rate estimated at the application through sliding window averaging is smaller than a target data rate (64 kbps in this figure). The duration of this time window is one second. These QoS specifications are in line with those usually required by, e.g., a video streaming application, in order not to have any perceivable degradation in the video quality. Hence. we say that an MT moves seamlessly across different radio technologies/APs if the above QoS requirements are met at all times. We note that the PER (Fig. 3) can also be seen as a satisfaction index giving the percentage of time during which the MT receives correct packets. With this interpretation in mind, we can directly compare Fig. 3 and Fig. 4. As a result of this comparison, we can conclude that jitter and rate requirements quantitatively matter as the dissatisfaction index in Fig. 4 is higher than the corresponding PER at the application layer (Fig. 3). In addition, the difference in performance between the AN and legacy case is slightly reduced in Fig. 4. However, we shall note that this reduction

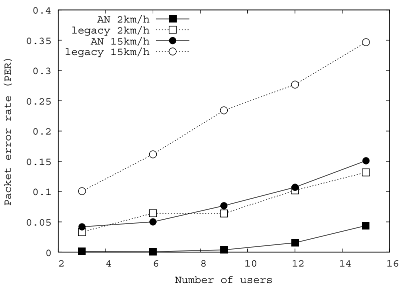

Fig. 5. PER vs. system load (scenario 2, 2 and 15 km/h, 64 kbps, legacy and AN).

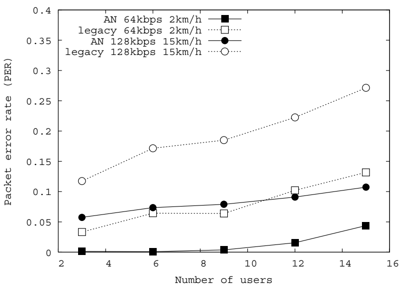

Fig. 6. PER vs. system load (scenario 2, 2 and 15 km/h, 64 and 128 kbps, legacy and AN).

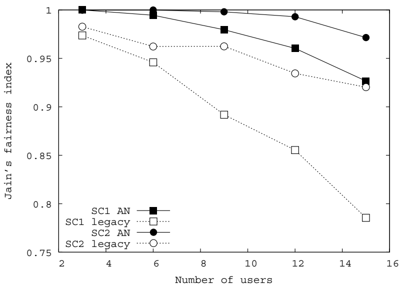

Fig. 7. FI vs. system load (scenarios 1 and 2, 2 km/h, 64 kbps, legacy and AN).

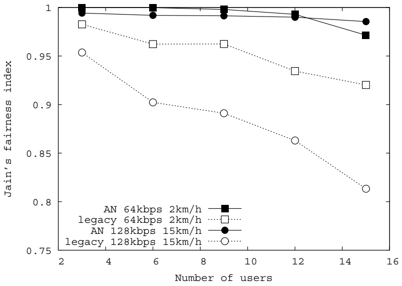

Fig. 8. FI vs. system load (scenario 2, 2 and 15 Km/h, 64 and 128 Kbps, legacy and AN).

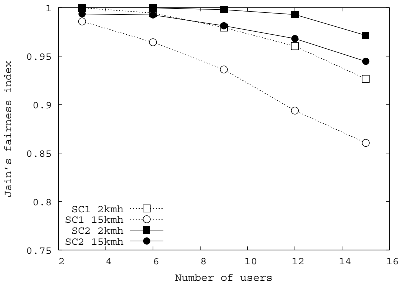

Fig. 9. FI vs. system load (scenarios 1 and 2, 2 and 15 km/h, 64 kbps, AN only).

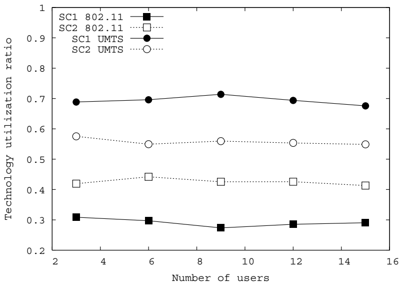

Fig. 10. Technology utilization time (scenarios 1 ad 2, 2 km/h, 64 kbps, AN only).

is nearly constant over the whole range of considered system loads. Also, the curves in Fig. 4 are well approximated by a rigid shift of those in Fig. 3. While not being a rigorous proof, these facts provide evidence that the PER metric, in the considered scenarios, is representative of the actual QoS perceived at the application. Finally, Fig. 4 can be used as an aid to system design, i.e., to infer the maximum number of users that can be supported for given QoS requirements.

Fig. 5 shows the PER as a function of the user's speed for simulation scenario number two. The performance improvements provided by the AN access selection architecture are larger for increasing speed. Also, the legacy case at 2 km/h performs nearly as well as AN at 15 km/h. This suggests that AN for a given PER target can support MTs moving at higher speeds. In Fig. 6 we show the PER metric for two extreme cases. In the former, the data rate is 64 kbps and MTs move according to a pedestrian mobility behavior, i.e., at 2 km/h. The latter has a higher data rate of 128 kbps and the speed is 15 km/h. As expected, higher data rates as well as higher speeds degrade the system performance for both access selection schemes. However, we note that the performance gap between these two settings is larger for the legacy case. Moreover, we note that AN for 128 kbps and 15 km/h performs closely to the legacy case for 64 kbps and 2 km/h. Hence, the more intelligent AN access management strategy allows the support of more demanding applications in more severe mobility conditions.

In Figs. 7–9 we show Jain's fairness index (FI). If x_i is the average throughput measured at the application layer by MT i, FI is calculated as $FI = [\sum_{i=1}^{N_u}(x_i/B_{th}^i)]^2/[N_u \sum_{i=1}^{N_u}(x_i/B_{th}^i)^2]$, where N_u is the number of MTs in the system and B_{th}^i is the target bandwidth for MT i. In Fig. 7 we first compare simulation scenarios one and two: the AN access selection mechanism gives very good performance in all cases in terms of fairness. This is not true for the legacy case, see especially scenario number one. Moreover, for the AN scheme the fairness improves as we look at the second simulation scenario, which again confirms its effectiveness in exploiting additional access opportunities. The fairness performance is further studied in Fig. 8 where we only plot results for scenario number two. Simulations are again run considering the two extreme cases in terms of data rate and speed which we discussed above. AN performs well in both cases, whereas the performance of the legacy scheme is substantially impacted in the more demanding scenario (128 kbps, 15 km/h).

In the following Figs. 9–13 we evaluate the performance of the AN access selection architecture. With Fig. 9 we show the impact of the simulation scenario on its fairness performance. In practice, having more access points leads to an increased fairness. This means that the AN access selection algorithm properly balances the traffic among the available APs.

Fig. 10 shows the fraction of time during which MTs use the two radio technologies. The most important observation here is the effect of adding a further AP (scenario one → scenario two). The AN architecture is in fact able to exploit the new AP such that the utilization ratio tends to 0.50, i.e., MTs on average use each radio technology for 50% of the time. This effect is further emphasized at higher speeds: we do not report these results here due to space constraints. We further observe that the technology utilization ratio does not depend on the number of MTs. This is a desirable fact as it means that the system resources are used in a fair manner, irrespective of the system load.

As discussed in Section III, each FE in the architecture (including the application) periodically monitors the experienced

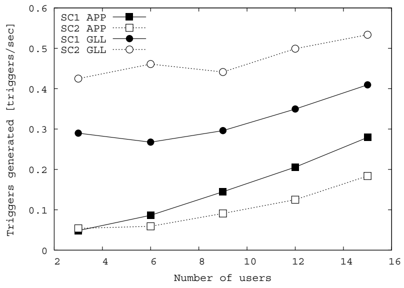

Fig. 11. Triggers generated from GLL and APP (scenarios 1 and 2, 2 km/h, 64 kbps, AN only).

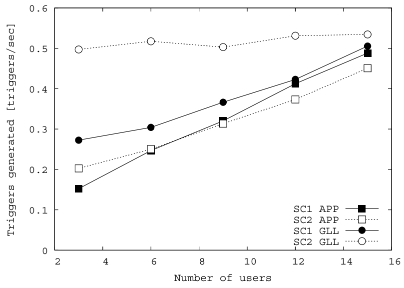

Fig. 12. Triggers generated from GLL and APP (scenarios 1 and 2, 15 km/h, 64 kbps, AN only).

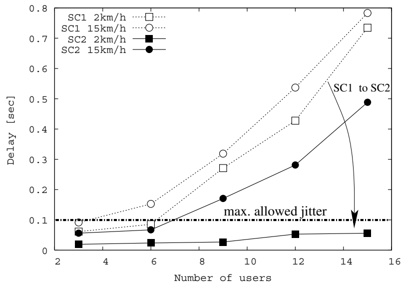

Fig. 13. Handover delay vs. system load (scenarios 1 and 2, 2 and 15 km/h, 64 kbps. AN only).

QoS. If the measured quality drops below some predefined threshold, the FE in question informs the MRRM through certain signals that are referred to as *triggers* [8]. Triggers may also be generated upon the detection of a new AP, network or service. These signals are subsequently evaluated by the MRRM which, based on its view of the QoS experienced at the different layers of the protocol stack, decides whether they should be ignored or taken into account. In the latter case, triggers may eventually lead to handing over to a different AP or to using a different radio technology. In Figs. 11 and 12 we plot the triggers per second generated by GLL and APP. First of all, we note that the trigger generation rate is reasonably low in all cases (smaller than 1 trigger every two seconds). Also, as the number of available accesses increases (i.e., scenario one → scenario two) the GLL generation rate increases as well, whereas APP reduces the number of triggers it generates per second. This is correct as a larger population of radio accesses inherently implies more frequent QoS evaluations at the link layer (GLL). In addition, if these accesses are properly exploited, the QoS at the application will be more stable and, in turn, the APP trigger generation rate will decrease. In Fig. 12 we show the effect of increasing the user's speed. In this case, APP generates more triggers as the QoS experienced at this layer is more likely to be endangered by the frequent changes in the point of attachment to the network. The GLL trigger generation rate is instead almost unaffected as it rather depends on the status of the surrounding wireless links (especially in terms of multi-user interference). In fact, the dynamics of these links are usually faster than those at which the MT changes its point of attachment.

Finally, in Fig. 13 we plot the average delay experienced during a handover between APs. As expected, the number of MTs in the system heavily impacts the delay performance. This is especially due to the multi-user interference, which makes the attachment procedure, as well as the corresponding CA, longer. We note that the delay is substantially reduced in scenario number two, through the addition of a further 802.11 AP: this, in the pedestrian mobility case (2 km/h) allows to meet the delay requirements for a larger range of system loads. The user's speed considerably affects the delay performance as well. In this case, in fact, an MT using a short range technology (802.11) has a shorter time to complete the attachment procedure before the connection with the old 802.11 AP breaks.

V. CONCLUSIONS

In this paper we presented and evaluated the Ambient Networks access selection architecture for heterogeneous wireless networks. We first presented this architecture along with its functional entities (FEs), their interrelations, as well as the algorithms that are to be run within each FE. We subsequently described our ns2 simulation framework. We finally discussed, through extensive simulation results, the effectiveness of this architecture in providing the users with seamless connectivity as they move across networks. The chosen access selection strategy was intentionally simple to pinpoint the inherent limitations of our approach. According to this policy, handovers are in fact triggered when the QoS drops below predefined thresholds. More advanced mechanisms are however possible. For example, inter-technnology handovers may be initiated even when there is no apparent degradation in the experienced QoS, according to user defined preferences, e.g., always pick the cheapest connection with a given QoS. Alternatively, network side algorithms may invite users to move across systems to achieve load balancing. These topics are left for future research.

REFERENCES

[1] N. Niebert, A. Schieder, J. Zander, and R. Hancock, *Ambient Networks*. John Wiley & Sons, 2007.
[2] C. Kappler, P. Pöyhönen, M. Johnsson and S. Schmid, "Dynamic Network Composition for Beyond 3G Networks: a 3GPP Viewpoint," *IEEE Network*, vol. 21, no. 1, pp. 47–52, Jan/Feb. 2007.
[3] G. Fodor, A. Eriksson, and A. Tuoriniemi, "Providing Quality of Service in Always Best Connected Networks," *IEEE Commun. Mag.*, vol. 41, no. 7, pp. 154–163, July 2003.
[4] W. Zhang, J. Jaehnert, and K. Dolzer, "Design and Evaluation of a Handover Decision Strategy for 4th Generation Mobile Networks," in *Proceedings of IEEE VTC*, Jeju, Korea, Apr. 2003.
[5] A. P. da Silva, F. R. P. Cavalcanti and R. A. de O. Neto, "VoIP Capacity Analysis of Wireless Multi-access Networks using Access Selection Schemes," in *Proceedings of IEEE PIMRC*, Athens, Greece, Sept. 2007.
[6] ns-MIRACLE: Multi InteRfAce Cross Layer Extension for ns-2, http://www.dei.unipd.it/ricerca/signet/tools/nsmiracle.
[7] N. Baldo, F. Maguolo, M. Miozzo, M. Rossi and M. Zorzi, "ns2-MIRACLE: a Modular Framework for Multi-Technology and Cross-Layer Support in Network Simulator 2," in *Proceedings of NSTools*, Nantes, France, Oct. 2007.
[8] J. Mäkelä and K. Pentikousis, "Trigger management mechanisms," in *Proceedings of ISWPC*, San Juan, Puerto Rico, USA, Feb. 2007.
[9] K. Pentikousis, R. Aguero, J. Gebert, J. A. Galache, O. Blume and P. Pöyhönen, "The Ambient Networks heterogeneous access selection architecture," in *Proceedings of M2NM*, Sydney, Australia, Oct. 2007.
[10] T. Camp, J. Boleng, and V. Davies, "A survey of mobility models for ad hoc network research," *Wireless Communications and Mobile Computing*, vol. 2, pp. 483–502, 2002.