Enhancing Video Delivery in the LTE Wireless Access Using Cross-Layer Mechanisms

Michelle Wetterwald[1(✉)], Leonardo Badia[2], Daniele Munaretto[2], and Christian Bonnet[3]

[1] HeNetBot, Sophia Antipolis, France
Michelle.Wetterwald@henetbot.fr
[2] University of Padova, Padua, Italy
{badia,munaretto}@dei.unipd.it
[3] Eurecom, Sophia Antipolis, France
Christian.Bonnet@eurecom.fr

Abstract. The current evolution of the global Internet data traffic shows an increasing demand of video transmissions, which potentially leads to the saturation of mobile networks. To cope with this issue, this paper describes techniques to handle the video traffic load in the last hop, of the communication network, i.e., the wireless access. The general idea is to benefit from a cross-layer architecture for efficient video transport, where multiple wireless access technologies, represented by Wi-Fi and next generation cellular technologies (4G and beyond), interact with the upper layers through an abstract interface. This architecture enables the introduction of enhancements in the LTE-A wireless access: evolved Multimedia Broadcast and Multicast Services (eMBMS) extended with dynamic groupcast communications, video relay at the Packet Data Convergence Protocol (PDCP) level and a smart video frame dropping mechanism to provide mobile users with a satisfactory level of Quality of Experience (QoE). These video-aware mechanisms leverage the abstract interface and allow mobile operators to fine-tune their networks while coping with the upcoming mobile video traffic increase.

Keywords: Wireless access · LTE-advanced · Video transport · Mobile network operators · Cross-layer optimisation · eMBMS

1 Introduction

Recent market studies [1] and future technology forecast reports [2] show that the share of video in global Internet traffic is growing at a rapid pace. It already represents the majority of the Internet traffic and is going to become dominant in the near future. In parallel, due to the diffusion of smart mobile phones and tablets, users consume videos via wireless networks, either local or cellular. Mobile network operators face the growing challenge of providing wireless accesses tailored to the expected level of QoE at the user side when consuming Mobile TV, Video on Demand or user-generated content (upstreaming).

Taking this challenge into consideration, the objective of the MEDIEVAL project [3] was to enhance the existing network architecture to efficiently deliver video

R. Agüero et al. (Eds.): MONAMI 2014, LNICST 141, pp. 18–31, 2015.
DOI: 10.1007/978-3-319-16292-8_2

applications to the mobile users. The designed architecture is composed of four sub-systems, Video Services Control on top to provision the network services, then Transport Optimization (TO) to enhance video quality using transport and caching mechanisms and Mobility Management (MM) to allow video flow continuation when roaming [4] and finally, Wireless Access to optimise access network functions for video delivery in the last hop through heterogeneous wireless access technologies. Hence, novel mechanisms in the Wireless Access sub-system are designed and focus on enhanced access techniques which exploit cross-layer optimisations through the interaction with upper layers, e.g., application and transport layers. Contention-based techniques, such as the IEEE 802.11 standard for Wireless Local Area Networks (WLANs) [5], and coordination-based, e.g., the Long Term Evolution Advanced (LTE-A) of Third Generation Partnership Project (3GPP) cellular systems are covered.

As a main pillar of its global architecture, a wireless abstract interface guarantees a transparent interaction between the underlying wireless technologies and the video traffic-aware upper layers. This interaction is built upon the IEEE 802.21 standard, pictured in Fig. 1, which proposes three different Media Independent Handover (MIH) Services [6] and offers to the upper layer management protocols generic triggers, information acquisition and the tools needed to perform mobility. The Event Service (MIES) provides the framework needed to manage the classification, filtering and triggering of network events, and to dynamically report the status of the links. The Command Service (MICS) allows the upper layer management entities to control the behaviour of the links. The Information Service (MIIS) is distributes the topology-related information and policies from a repository located in the network. They result in a cross-layer architecture where the Media Independent Handover Function (MIHF) operates as a relay between the media-specific Link layer entities and the media-agnostic upper layer entities, e.g., MIH-Users. In the mobile terminal, the MIH-User is usually represented by a Connection Manager (CMGR) whose main role is to decide which path is best suited to reach the application server or the Correspondent Node (CN) located across the Internet [7].

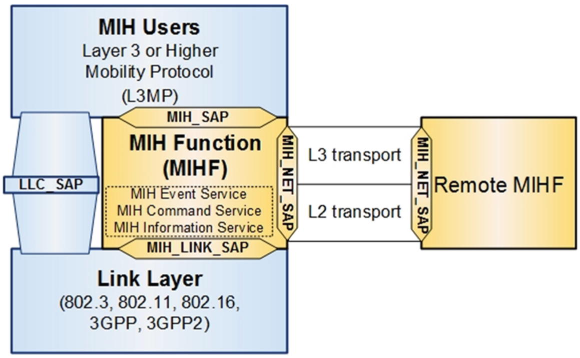

Fig. 1. IEEE 802.21 cross-layer model

Indeed, in the proposed architecture, a cross-layer relationship is established with upper components, i.e., the MM and TO subsystems, to exchange information about the capabilities of the components at the lower layers, as well as to configure them [8]. This interface, pictured in Fig. 2, is optimised by a central abstraction layer. This layer operates at both the Mobile Terminal (MT) and the Point of Attachment (PoA) to the network, which corresponds to the access point in WLAN and the base station, or eNodeB, in LTE-A networks. The associated functions are split into two main streams, as shown in Fig. 2. A Monitoring function dynamically retrieves the information related to the access networks availability and quality in order to provide it to the upper layers through the abstract interface. Moreover, it senses the environment searching for new available access networks; whenever they are found, it analyses their capacity, bandwidth usage, and available resources. The MM is mostly interested in the wireless signal events, while the TO considers the traffic measurements allowing a more precise estimation of the wireless cell load. Secondly, a Dynamic Configuration function takes into account the requests from the upper layers and the characteristics of the video flows to setup the network interface or establish radio channels to accommodate an upcoming data flow. It works by defining a utility function which makes it possible to allocate resources by providing the optimal set of parameters describing each technology. Generic Quality of Service (QoS) parameters for link throughput, classes of service, or error rates, are defined and translated through the interface into technology-specific parameters. The MM is thus able to activate and deactivate the network interfaces and resources due to the roaming of the user or some optimization decision made in the core network entities. The TO is able to command that some video frames, marked at the Internet Protocol (IP) packet level, become prioritised before leaving the PoA, avoiding deep packet inspection and thus preserving the user QoE.

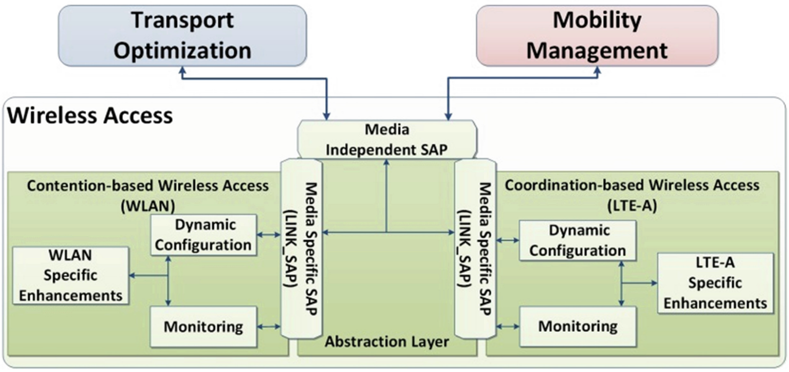

Fig. 2. Wireless Access sub-system architecture

As a consequence, the Wireless Access (WA) subsystem of the architecture is split into three main functional blocks. The abstraction layer component provides the generic interfaces between video specific functions (i.e., transport and mobility), while the wireless components include the features and mechanisms designed to further enhance the video flow transmission over the air. In fact, besides being tightly coupled with the monitoring and dynamic configuration functions, the wireless components have been enriched with technology-specific functionalities benefiting from the cross-layer architecture. Video applications are characterised by high throughput, i.e. large bandwidth to ensure good visual quality, and a strong sensitivity to jitter. Novel features and techniques should address these constraints. The focus of this paper is on the work performed from a system view on the upper layers of the LTE-A radio interface, contained in the "LTE-A Specific Enhancement" block shown on the right of Fig. 2. The enhancement applied to the cellular system covers group communications based on the 3GPP evolved Multimedia Broadcast and Multicast Services (eMBMS) standard. It further extends the cell capabilities and coverage thanks to the introduction of a relay at Layer 3 level between the eNodeB and the User Equipment (UE) and finally, when these methods are not sufficient, smartly drops part of the video traffic to ensure a target quality to the users. All three techniques can be used independently or complement one another.

The objective of this paper is to describe the enhancements achieved by the project for the upper layers of the LTE radio interface and provide directions to help the network operators better deliver video traffic in their cellular networks. The discussion is organised as follows. Section 2 discusses the optimization of group communications in the cellular LTE technology, i.e., the improvements proposed for the eMBMS multicast support. In Sect. 3, relays operating at the Packet Data Convergence Protocol (PDCP) level, just below networking layer, are introduced. Their impact on the quality of the video transmitted in the cell is analysed and evaluated. In Sect. 4, we propose a mechanism to smooth the load in the cell and avoid visual degradation of the video. Finally, we conclude the paper by assessing these different techniques, highlighting their benefits and suitability for future mobile networks.

2 Introducing Dynamic Groupcast Communications in the LTE Cell

The first enhancement applied to the LTE-A system addresses group communications. Since video content uses a large amount of the available transport capacity, distributing the same data to several users located in the same area wastes radio resources. Conversely, multicasting or broadcasting the service allows saving the resources that would be used if unicast Data Radio Bearers (DRB) were established for other users and/or purposes. Multicast communications allow sharing the resources on the wireless hop when a geographically-close and potentially large group of mobile listeners watches the same program. In LTE-A, the services broadcast by eMBMS are enhanced to support dynamic multicast sessions together with user mobility.

In the cellular part of the WA architecture, multicast is optimised by supporting and extending the eMBMS bearer service specified in the 3GPP standards [9, 10]. Its objective is to enable point-to-multipoint communications (p-t-m) over the radio interface (or Access Stratum), allowing resources to be shared in the network. The MBMS support has been subject to serious revisions within the 3GPP standardization, with the inclusion of new tools and procedures to improve its performance. Actually, the handling of multicast flow has disappeared in the transition between the initial and evolved versions of this standard, mostly due to business causes, costs and complexity of deployment. In the LTE and LTE-A systems, only broadcast sessions are proposed.

The Multicast-Broadcast Single Frequency Network (MBSFN) areas, pictured in Fig. 3, hosting the eMBMS, are configured semi-statically. When the network is built, some eNodeBs are set-up in order to support point-to-multipoint transmissions, while others, pertaining to reserved cells in the same area, do not offer that service. The MBMS configuration is beaconed over the related cells in two different messages (or System Information Blocks, SIB), independently of the number of listening mobile users in the cell. To avoid the allocation of broadcast resources (MBMS Radio Bearer or MRB) when the number of users is low, the eNodeB implements a counting procedure, where the connected MTs in the cell are invited to signal themselves back to the base station in uplink. This procedure is used to perform admission control and allocation of the MRB resources. In more recent advances, mobile nodes are able to inform the network of their interest and have the capability to receive MBMS sessions from a certain set of frequencies of the MBSFN, allowing the network entities to further enhance resource allocation in the cell. This information is transferred to the target eNodeB during the handover preparation phase within a specific MBMS context associated to the MT.

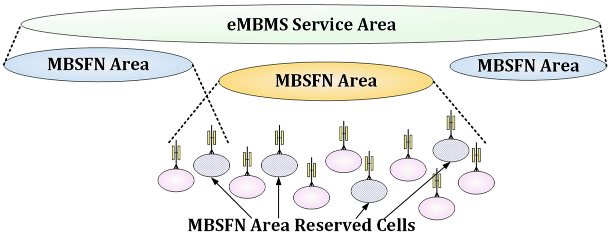

Fig. 3. eMBMS Areas

We extend these features to improve the semi-static broadcast support provided by the MBSFN. By using a cross-layer operation between the LTE-A component and the MM sub-system through the dynamic configuration function, the eMBMS can include the dynamic multicast resource allocation in a manner similar to what earlier planned by the standard. In our study, we simplify the Multicast session start and stop procedures at the eNodeB and their notification to the mobile. Another important feature is the counting of listening mobiles in each cell by the eNodeB. This information is used to trigger the multicast session if needed or move the flow back to a point-to-point bearer if only one user in the cell is listening. To avoid interference over other types of traffic (e.g., voice calls) that could take place simultaneously, it is important to establish a coordinated control of unicast and multicast communications in a cell providing the MBMS service.

When the connected MT joins a multicast session, an MBMS context is created in the network entities. Whether it happens while being attached to the LTE cell or during a handover, the procedures that enable it to receive the session are executed in the wireless access modules, as shown in Fig. 4. A request to activate the multicast resources is received by the LTE-A Access module in the PoA. If relevant, and based on internal algorithms taking into account the resources already allocated, it triggers the MBMS Session Start procedure, establishing a new MRB and informing the MT. The "MBMS session start" is executed dynamically in the eNodeB, upon the request from the MM sub-system, which removes the constraint to allocate resources when the network is built. If it happens during a handover, the MT still connected to its old PoA receives this information during the preparation phase and is thus able to configure and receive the MBMS service as soon as it attaches to the target cell.

The MT joins the service only once, as long as the context can be transferred between PoAs. This is another feature of the MM [11], which avoids the constraint of the MT self-signalling during the whole data reception period, whether in mobility or not. This enhancement allows a smooth support of the counting procedure, but with the eNodeB capable of identifying by itself the attached MTs that own a multicast context in the cell. It can then adapt the resource allocation to the real bandwidth consumption and the actual number of mobile listeners in its cell. When it detects that a flow, identified by a specific flow label and source address, marked as "multicast-enabled" is received simultaneously by several MTs, it transfers the video data into an MRB, even if the core network is not multicast-enabled. This improves sharing resource in the wireless access. Table 1 provides a summary of the mechanism traces recorded at the eNodeB.

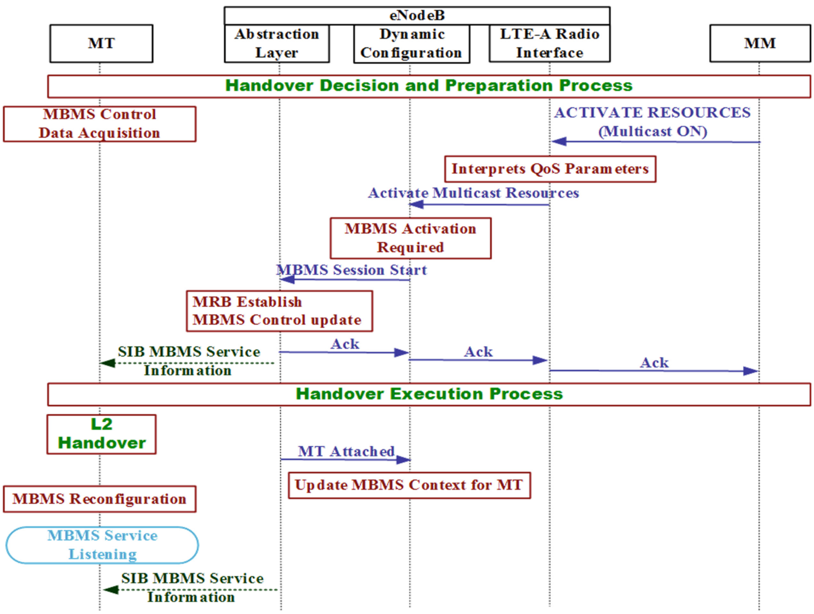

Fig. 4. Enhanced eMBMS Session Start during handover

Table 1. Traces obtained in eNodeB when applying the dynamic Session Start (time in ms.)

Steps	Start	End
- Final step of MT arrival in the cell (MT connected)		0.000
- MBMS context for service 97 established for the MT	0.011	0.070
- Successful MBMS context setup in the lower layers	2.673	2.666
- First multicast packet from IP to be sent to the MT	14858.290	14858.297
- eMBMS session start triggered	14858.301	14858.343
- Procedure on-going, packet sent as unicast, which prevents it from being delayed	14858.344	14858.360
- Notification: successful completion of the procedure	14922.571	14922.586
- IP multicast packet forwarded on the MBMS bearer	15859.957	15859.983
- IP multicast packet forwarded on the MBMS bearer	16857.716	16857.743

Another impact is expected also on the configuration of the radio access when taking into account the spectrum usage and the resource allocation. Multicast flows require bandwidth reservation based on the dedicated eMBMS Bearer parameters received from upper layers and the worst Channel Quality Indicator (CQI) of Multicast clients measured in the lower layers. This results in a bad spectrum usage because users with a robust link underutilise the bandwidth resources. Our solution combines H.264/SVC (Scalable Video Coding) together with cross-layer optimization to dynamically increase/decrease the video quality perceived by each user according to the different channel feedback messages, using mechanisms similar to those described in Sect. 4.

This is of particular interest for the Personal Broadcast Service [12] studied by the project and that is expected to gain momentum in the coming years. Here, user generated video content is distributed to a group of mobile listeners. When they are located in the same area, an eMBMS session can be activated. A typical use case is a group of tourists receiving personalised information from their guide during a visit [13] or the dissemination of a road hazard event in a cooperative vehicular system.

3 Relaying the Video Traffic at PDCP-Level

The eMBMS can be coupled with another feature introduced in the project. An LTE-A relay, operating on top of Layer 2, is able to improve the coordination between the unicast and the multicast transmissions in the cell by offloading the eMBMS sessions from the regular user traffic. This is made possible thanks to the flexibility provided by the cross-layer architecture to start the session dynamically in the LTE PoA.

The relaying scheme is introduced at the PDCP level in the LTE access network. It is worth noticing that in parallel to this work, Layer-3 relays were also being studied within 3GPP, and included in the LTE-A architecture at stage 2 level (i.e. high level design) [14]. The work achieved in the standard focuses on a new interface, the Un, between a dedicated eNodeB (called the Donor eNodeB) and the Relay. Moreover, as we mainly focus on video transmissions, we decide to assess the impact of the delay introduced on video streams by the relaying architecture.

Relaying techniques offer an interesting method for extending and improving wireless networks capabilities [15]. These techniques have been selected as part of the enhancements introduced to the LTE-A architecture. Their effectiveness has been investigated in the literature, showing good results in terms of both network coverage beyond the eNodeB and overall capacity. Outdoor measurements have shown in [16] that a time-shared LTE relaying system with 20 MHz bandwidth can both achieve 60 Mbps of data rates and cover the coverage holes in urban macro environments with a diameter of 300 m. A performance evaluation has been also accomplished in [17] via simulations showing interesting trade-offs between transmission power of both eNodeB and Relay Node (RN) and their positions. Two different LTE relay deployments are proposed in [18] considering the following criteria: early deployment (i.e., compatibility with current LTE Evolved Packet Core, or EPC, architecture), system complexity and traffic performance. The architecture complexity has been reduced considering packet aggregation of multiple UE flows with the same QoS requirement. Finally, header compression and stripping under the *Un* interface are added. Generally, the approach used in the literature focus on a very tight set of aspects of the LTE architecture, due to the complexity of the overall system. Here, on the contrary, we study the problem of LTE relay from an architectural point of view, considering all the aspects involved in a real deployment, from the IP level to the wireless access.

Relaying mechanisms usually operate on the LTE radio interface and can be performed at several levels: physical layer, link layer or just below the IP protocol stack. At physical layer level, the relay only repeats the received Radio Frequency (RF) signal. Such technology has been in operation for some time because it is very cheap and relatively simple. However, it increases the level of interference in the system, both propagating the inter-cell interference already present in the RF signal and introducing an additional contribution from the backbone signal to the relayed signal. Layer 2 relays introduce additionally demodulation, decoding, encoding and modulation, thus eliminating the noise. The Layer 3 relay operates on top of the PDCP level. It benefits from all the error correction mechanisms and transmission quality brought by the link layers, since the IP packet is extracted from the RN radio bearer and forwarded onto another UE radio bearer. However, this operation may have a cost in terms of QoS, which we evaluate for video applications. This relaying scheme has also an impact on the signalling flows and procedures for the attachment, detachment and coordination of resources management functions, both at the relay and mobile nodes.

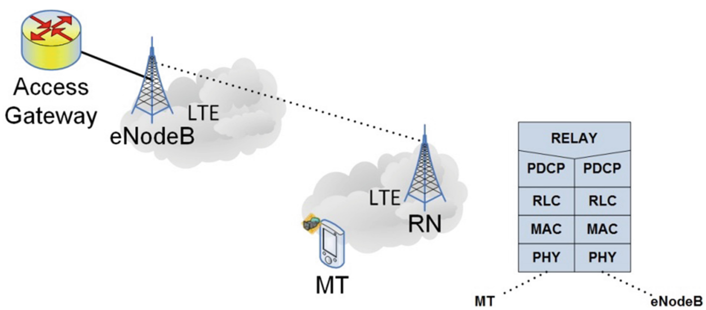

Fig. 5. LTE Relay Node in the wireless access

Figure 5 depicts how the RN plays a role in the wireless access architecture. A radio configuration similar to 3GPP is adopted. The eNodeB and RN signals at the physical layer level are assumed to be differentiated either by operating each link on a different frequency or by time-division multiplexing. The control plane analysis we perform mostly focuses on the impact on the latency and on the radio interface procedures for network attachment, session setup and tear down, and detachment of the mobile node or the LTE relay from the network. We consider here that the LTE relay serves as an extension of the network to increase its capacity and thus is not moving. The analysis also involves the wireless abstract interface, which allows the upper control layers to be agnostic from the specifics of the LTE technology.

At the initialization phase, the LTE module triggers the attachment of the RN to the LTE eNodeB, signalling that it is actually a RN. When the procedure is over, the RN starts broadcasting the system information in its cell. When a MT connects to the network, the RN informs the eNodeB that a new MT has appeared and retrieves its new cell configuration parameters, differentiating those related to the link with the eNodeB from those related to the link with the MT. A similar but reversed procedure is triggered when the connection has to be reconfigured because a new video session has started at the MT. In the data plane, the RN receives the packets from the PDCP layer on one side and forwards them to the opposite path. It can accommodate eMBMS sessions in an identical manner, potentially providing a different PoA for those MTs that are interested in receiving the multicast communications and alleviating the impact of eMBMS on other types of sessions.

The impact on the control plane turns into additional latency for establishing signalling and data radio bearers during session setup or when executing a handover. Execution traces, recorded by one of our partners in an operational network during the attachment of a MT, show that a radio reconfiguration takes only a very few milliseconds (less than 4 ms) compared to a total attachment time of 1.33 s. It can thus be accounted that in the control plane the impact of adding a relay at PDCP level will be minimal.

The theoretical analysis of the impact of the LTE relay on data traffic can be split into two parts. Firstly, the impact of the forwarding in the LTE Relay itself and secondly, the impact of adding a second radio link before the delivery of packets to the MT. The second radio link doubles the burden of radio transmissions on the traffic flow. It increases the effect of the Relay-to-eNodeB radio link on the QoS metrics for the delay or the jitter, but can be compensated by an adaptation of the coding and modulation techniques and parameters used on each link. Packet loss is compensated by the fact that the relay operates at PDCP level and that Layer 2 recovery mechanisms are fully operational.

In order to evaluate the resulting performance of such a scheme, we implemented a simple scenario within a network simulation performed with the open-source simulator ns-3 [19]. There, we show the improvement in terms of throughput achieved in a cellular network when relay nodes are enabled to help the eNodeB deliver the packets to multiple users. In this scenario, we first place 20 users in the coverage area of an eNodeB (transmission power of 30 dBm, bandwidth 5 MHz), using the Friis propagation loss model. In a second phase, we place 2 relays at few km from the base station. The base station sends 500 packets of 1024 bytes every 20 ms to each node. The simulation runs do not take into account signalling, which was studied independently as aforementioned and we assume that the channel between eNodeB and relays is ideal. This simplification can be justified by the fact that the RN is considered static with an optimised radio link towards the eNodeB.

Figure 6 shows the comparison of data reception for the different nodes according to their distance to the eNodeB. The blue points show the reception in the case without relays, whereas the magenta squares show the simulations with two relays. The figure confirms that in all cases the situation of the worst nodes, i.e., that suffered from losses in the standard case, has been improved to a large extent.

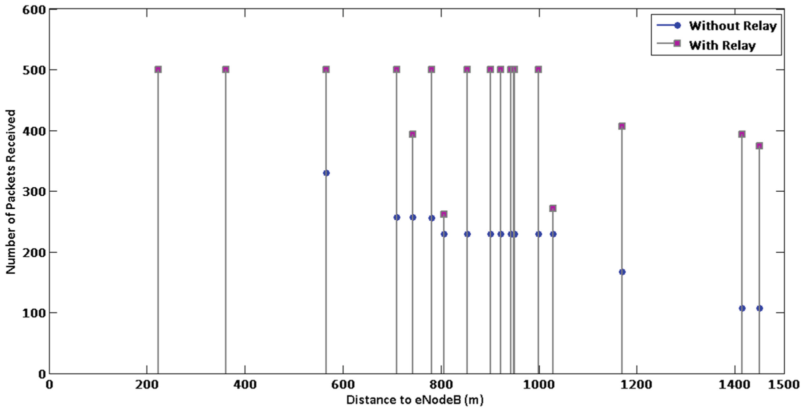

Fig. 6. Comparison of coverage with and without relay

This functionality permits to extend the network coverage while still benefiting from the transmission quality and error recovery present in the link layer protocols. MTs closer to the RN than to the eNodeB can access the cell while still obtaining a good communication quality. A larger number of users can be accommodated through the same eNodeB by distributing their load between several relay nodes, hence improving the scalability performance of the wireless access. The traffic passing through the eNodeB can be increased, compared to a standard MT-eNodeB attachment, since the transmission between the Relay and the eNodeB is expected to be of good quality and can use modulation and coding schemes with low redundancy. The results obtained prove that this type of relay has a moderate impact on the general control plane procedures, while improving drastically the transmission and coverage of the LTE cell, which benefits network operator and users.

Even though such relays had been under specification in 3GPP since the beginning of the project, our study has shown how they could positively impact the video traffic delivery. Beside enlarging the coverage and improving the reception quality in the related cells, we propose that such relays are used to separate the eMBMS groupcast listeners from the regular users with unicast traffic, which would put aside current limitations faced by operators to deploy the eMBMS. One of the major reasons for not deploying MBMS in previous releases of 3GPP was its radio impact on other types of communications when sharing the same cell. Coupling an LTE RN to the eNodeB to handle specifically the MBMS traffic allows a dedicated node with differentiated physical and medium access parameters to serve as MBMS PoA for video delivery. Users listening to MBMS broadcast or multicast sessions can be attached to the LTE RN while the others remain attached to the eNodeB (or another LTE Relay attached to it) and are unaffected.

4 Smart Video Frame Dropping

In the previous sections, mechanisms were introduced to extend the capability of the LTE-A cell. However, there are cases when this is not sufficient and sudden heavy traffic load conditions have to be handled. The simple, yet very unpopular, solution consists in denying access to new users or even breaking some existing communications. Accepting all data traffic means that part of the data packets will not be able to go through, being dropped in a random fashion at the link layer, which may generate a temporary degradation or even stalling of the image on the screen [20].

The last mechanism outlined in this paper to improve the transport of video applications in the LTE-A cells selects instead specific video frames in the eNodeB to address overload in the last hop. We propose a cross-layer mechanism where we try to resolve the issue of high occupancy of Radio Link Control (RLC) buffers, by reporting it through the abstract interface to the TO. The upper layers can mark the priority of the IP packets according to their video content (e.g., SVC video layer). The lower priority packets can be dropped based on parameters transferred through another cross-layer interaction in the eNodeB.

A cross-layer Video Frames Selection function performs this temporary rate adaptation on the last hop, yet avoiding deep packet inspection. It classifies and filters the received video frames according to a dedicated mark previously introduced in the IP packet header. When a congestion is detected in the network, the data packets are marked for prioritisation by the TO. The lower priority packets can then be dropped before the video frames are actually handled by the Link layer protocols, according to the receiver capabilities. This reduces the bandwidth occupation and loosens the level of traffic load in the last hop. The process initially designed performs the full process inside the PoA itself: detect the congestion, decide on the filtering and drop the packets. However, considering that a global SVC layer optimization algorithm exists in the TO, an alternative solution has been adopted that keeps the decision and marking update of the IP packets in the TO, based on the results of its algorithms, while the decision is executed in the LTE-A specific wireless component. This last operation, restricted to the overloaded cell, is accomplished in the eNodeB, after the packets coming from the Core Network have been decapsulated from the General Packet Radio Service (GPRS) Tunnelling Protocol-User (GTP-U) tunnel and before they get encapsulated in the PDCP protocol.

Figure 7 indicates with a (*) the components of the implementation involved in this mechanism. New functions have been introduced in the RRC (Radio Resource Control) and LTE-A specific wireless components at the eNodeB that retrieve the measurement of buffer occupancy from the RLC layer and signal an event to the upper layers through the abstract interface when this occupancy reaches a certain threshold corresponding to heavy load conditions. In the case of the initial solution, where the whole process is performed in the eNodeB, a classifier located at the Non-Access Stratum (NAS) driver above the PDCP layer is able to drop silently the least significant video frames, based on the marking of the packets arriving from the IP protocol stack. The classifier operates by comparing the Differentiated Services Code Point (DSCP) field of the IP packet header with an active mask, thus avoiding deep packet inspection of other header or even data fields in the classifier, and of the network layer fields in the wireless access layers. In the alternative solution, on request from the TO, some measurements of the planned Physical Resource Blocks (PRB) and total data volume from the MAC layer are reported through the abstract interface, enabling the TO to drop the least important packets directly in the core network. The implemented process affects the eNodeB only, and is split between the LTE radio interface protocols (RRC and MAC layers), and the LTE-A specific component which retrieves and analyses the measurements, then executes the required actions.

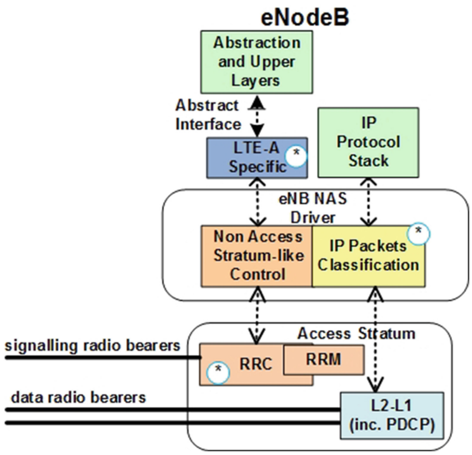

Fig. 7. Implementation of the video frames selection in the eNodeB

Functional results could be obtained with a local testing system. This successful test has been performed on a small testbed focusing mostly on network measurements and congestion detection. Another part of the testbed complemented this evaluation, taking care of the packet dropping as reported in [21]. The test performed here allowed validating the correct operation of the LTE-A specific module in cooperation with the radio interface protocol layers and the abstract interface. The traces obtained are summarised in Table 2. From a functional point of view, the correct execution of the following features has been verified: detecting the congestion situation in the eNodeB according to the specified threshold, triggering notification about the high load event to the TO, returning link traffic parameters on request from the upper layers and finally stopping the specific measurements when the situation has returned to normal condition, in order to reduce the mechanism overhead on the control plane.

Table 2. Traces recorded at the eNodeB during a congestion event (time in s.)

Event	Time
- LTE-A module receives an event subscription for congestion notification.	0.000
- It polls periodically the lower layers to check the cell correct operation.	35.991
- Congestion detected (RLC buffers for MT0 above threshold); a notification is sent to the upper layers.	41.296
- Upper layer (TO) requests periodic measurement retrieval	41.297
- Link traffic parameters forwarded through the L2.5 Abstraction Layer.	44.421
- After the problem resolution by the TO, the measures fall back to normal conditions.	72.975
- Request received from the TO to stop forwarding the measurements	72.976
- Request executed by the LTE-A module	75.579

5 Conclusion

This paper has described several enhancements proposed by the MEDIEVAL project to mobile network operators in order to help them more efficiently distribute the video traffic in the wireless cells. Our objective is to reduce the load imposed by this specific type of applications, which are undertaking a huge growth in the coming future. Under this objective, we have focused on next generation wireless networks where we aim at providing video-friendly optimizations. Towards that goal, we have based our architecture on three main pillars: cross-layer abstraction, access network monitoring and network interface dynamic configuration. They have served as a basis to the development of innovative features that should improve the current design of operator networks in the last hop. The first concept was based on group communications. We have enhanced the eMBMS to configure dynamic multicast sessions, with better performance for the session setup procedure, benefiting from the cross-layer design which allows receiving the eMBMS parameters at the eNodeB ahead of the session start. We have evaluated the impact of introducing eNodeB relays operating at the PDCP level on the QoS and cell coverage extension, including for separating eMBMS traffic from legacy service. Finally, we have implemented a cross-layer mechanism to selectively drop IP packets containing lower priority video frames in order to handle heavy load conditions in a specific cell and potentially avoid congestion or access rejection. This filtering applies in the eNodeB, at the junction between the GTP-U tunnel and the PDCP protocol. From these enhancements, we have demonstrated that the abstract interface introduced between the upper layer control entities and the wireless access modules provides additional capabilities to efficiently manage the network traffic and to introduce novel network mechanisms in a video-optimised way. Moreover, the combination of enhanced link-specific mechanisms allows the wireless link access to go beyond a simple wireless transmission of data.

Acknowledgments. The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7-ICT-2009-5) under grant agreement n. 258053 (MEDIEVAL project).

References

1. Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2013–2018. http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.html/. Accessed May 2014
2. Celtic Plus Purple book, March 2012. http://www.celtic-initiative.org/PurpleBook+/Purplebook.asp. Accessed May 2014
3. http://www.ict-medieval.eu
4. Costa, R., Melia, T., Munaretto, D., Zorzi, M.: When mobile networks meet content delivery networks: challenges and possibilities. In: ACM MobiArch, August 2012
5. IEEE Standard for Information Technology-Telecommunications and information exchange between systems-Local and metropolitan area networks-Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications, IEEE Std. 802.11, 2007
6. Piri, E., Pentikousis, K.: IEEE 802.21. Internet Protoc. J. **12**(2), 7–27 (2009)
7. Kassar, M., Kervella, B., Pujolle, G.: An overview of vertical handover decision strategies in heterogeneous wireless networks. Comput. Commun. **31**(10), 2607–2620 (2008)
8. Corujo, D., Bernardos, C.J., Melia, T., Wetterwald, M., Badia, L., Aguiar, R.L.: Key function interfacing for the MEDIEVAL project video-enhancing architecture. In: Pentikousis, K., Aguiar, R., Sargento, S., Agüero, R. (eds.) MONAMI 2011. LNICST, vol. 97, pp. 230–243. Springer, Heidelberg (2012)
9. Lecompte, D., Gabin, F.: Evolved multimedia broadcast/multicast service (eMBMS) in LTE-advanced: overview and Rel-11 enhancements. IEEE Comm. Mag. **50**, 68–74 (2012)
10. 3GPP TS 23.246: Multimedia Broadcast/Multicast Service (MBMS); Architecture and functional description, Release 12
11. Figueiredo, S., Wetterwald, M., Nguyen, T., Eznarriaga, L., Amram, N., Aguiar, R.L.: SVC multicast video mobility support in MEDIEVAL project. In: Proceedings of Future Network and Mobile Summit 2012, Berlin, Germany, 4–6 July 2012
12. 3GPP TR 22.947: Study on Personal Broadcast Service (PBS), Release 10
13. Badia, L., Bui, N., Miozzo, M., Rossi, M., Zorzi, M.: Improved resource management through user aggregation in heterogeneous multiple access wireless networks. IEEE Trans. Wireless Commun. **7**(9), 3329–3334 (2008)
14. 3GPP TS 36.300: Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Overall description; Stage 2, Release 10
15. Quer, G., Librino, F., Canzian, L., Badia, L., Zorzi, M.: Inter-network cooperation exploiting game theory and Bayesian networks. IEEE Trans. Commun. **61**(10), 4310–4321 (2013)
16. Wirth, T., Venkatkumar, V., Haustein, T., Schulz, E., Halfmann, R.: LTE-advanced relaying for outdoor range extension. In: Proceedings of VTC Fall, September 2009
17. Beniero, T., Redana, S., Hämäläinen, J., Raaf, B.: Effect of relaying on coverage in 3GPP LTE-advanced. In: Proceedings of VTC Spring, April 2009
18. Huang, X., Ulupinar, F., Agashe, P., Ho, D., Bao, G.: LTE relay architecture and its upper layer solutions. In: Proceedings of IEEE GLOBECOM, December 2010
19. NS-3 simulator. http://www.nsnam.org/. Accessed May 2014
20. Quality of Experience for Mobile Data Networks: White Paper, Citrix, 2013
21. Fu, B., Kunzmann, G., Wetterwald, M., Corujo, D., Costa, R.: QoE-aware traffic management for mobile video delivery. In: Workshop on Immersive and Interactive Multimedia Communications over the Future Internet, IEEE ICC 2013, Budapest, Hungary, 9–13 June 2013

http://www.springer.com/978-3-319-16291-1

Mobile Networks and Management
6th International Conference, MONAMI 2014, Würzburg, Germany, September 22-26, 2014, Revised Selected Papers
Agüero, R.; Zinner, Th.; Goleva, R.; Timm-Giel, A.; Tran-Gia, P. (Eds.)
2015, XII, 458 p. 201 illus., Softcover
ISBN: 978-3-319-16291-1