

Support of multiple content variants in the multimedia broadcast / multicast service

George Xylomenos, Konstantinos Katsaros and Vasilis Tsakanikas

xgeorge@aueb.gr, ntinos@aueb.gr and vasilistsakanikas@yahoo.gr

Mobile Multimedia Laboratory, Department of Informatics

Athens University of Economics and Business

Patision 76, Athens 104 34, Greece

Abstract—The Multimedia Broadcast/Multicast Service (MBMS) was designed to enable the mass distribution of multimedia content in 3rd Generation and beyond cellular networks. If such services are to become commercially viable, they must be able to efficiently support widely heterogeneous user requirements, for example, due to terminal limitations and cost constraints. This paper presents an MBMS extension that allows multiple variants of the same content to be economically distributed to heterogeneous receivers, explicitly taking into account the possibility of using either dedicated or common radio channels. We describe our extended multiple content variant MBMS model by explaining the modifications that it imposes on the standard MBMS model, as well as the manner in which it can be combined with layered coding. We also present an analytical evaluation of our approach against alternatives based on the standard MBMS in terms of control and user plane overhead, and compare the analytical predictions with detailed simulation results. Both the analysis and the simulations indicate that our proposal can indeed satisfy heterogeneous user requirements, while consuming considerably lower resources than the standard-based alternatives.

Index Terms—MBMS; UMTS; multimedia; multicasting

I. INTRODUCTION

The deployment of 3rd Generation networks has made cellular systems feasible platforms for distributing multimedia content to mobile users. While the resource requirements of services such as video streaming make them expensive for individual users, these costs can be dramatically reduced when multiple users desire to receive the same service. This can be achieved either by *broadcast*, where all users receive the service, or by *multicast*, where only selected users receive the service. To this end, the *Universal Mobile Telecommunications System* (UMTS), specified by the *3rd Generation Partnership Project* (3GPP), has introduced the *Multimedia Broadcast/Multicast Service* (MBMS) [1]. MBMS enables resource sharing for multicast and broadcast throughout the network, including over the resource-constrained radio links, for services ranging from media streaming to file downloads [2]. In this paper we specifically focus on multicasting, as it is more appropriate for targeting users that have subscribed to, and possibly paid for a service.

The MBMS multicasting mode is similar to IP multicasting [3]: it delivers the same content to all receivers participating in a group identified by an IP multicast address, using the same *Quality of Service* (QoS) parameters for the entire




Fig. 1. The components of MBMS.

multicast distribution tree. For multimedia services to become commercially viable, they must attract enough subscribers to leverage the savings due to radio link sharing [4]. When the receivers are heterogeneous however, for example, terminals with different screen resolutions or users with different budgets, it is difficult to select for distribution an appropriate variant of the content: a high-quality variant may not work with low-end terminals, while a low-quality variant may not satisfy users prepared to pay more for the better service; in both cases, potential users and revenues are lost.

To address this issue, we designed an MBMS extension that supports the distribution of multiple variants of the same content to different subsets of receivers, without introducing redundancy into the data stream; we refer to this extension as *Multiple Content Variant* (MCV) MBMS. The desired variant is dynamically selected by each receiver, based on the terminal capabilities and/or user preferences. Our approach takes into account the intricacies of MBMS, and in particular the possibility of using either dedicated or common radio channels. In this paper, we describe how MCV MBMS is derived from the standard MBMS and evaluate its performance against some standard-based alternatives via analysis and simulation, showing that our scheme greatly increases the number of satisfied users with a moderate increase in the resource consumption, thus making MBMS more attractive to operators and subscribers alike.

The outline of the rest of this paper is as follows. In Section II, we describe the standard MBMS model, along with its state management and signaling procedures, while in Section III we describe our extended MCV MBMS model. In Section IV, we present an analytical comparison of MCV MBMS against some standard-based alternatives, while in Section V we present a simulation based comparison. In Section VI, we discuss the related work and clarify our contributions, presenting our conclusions in Section VII.




Fig. 2. Standard MBMS Bearer and UE contexts.

II. THE STANDARD MBMS MODEL

An outline of a UMTS network supporting MBMS is shown in Figure 1. A new functional entity, the *Broadcast/Multicast Service Center (BM-SC)*, controls MBMS service provisioning. Both internal and external content sources are allowed, but data can only enter the network via the BM-SC. The *Core Network (CN)*, consisting of *Gateway GPRS Support Node (GGSN)* and *Serving GPRS Support Node (SGSN)* elements, the *Radio Access Network (RAN)*, consisting of *Radio Network Controller (RNC)* and *Node-B* elements, and the *User Equipment (UE)* must also be modified for MBMS.

In MBMS, a multicast group is identified not only by a class D IP address, as in IP multicast, but also by an *Access Point Name (APN)*, which essentially identifies a GGSN serving a specific UMTS network. While in IP multicasting anyone can send to and receive from a group, in MBMS a UE must first subscribe to a group, using some mechanism external to MBMS, in order to be later allowed to join it. In addition, only the GGSN may transmit data to a group. These are the ideal properties for commercial services where receivers are charged [1]. Finally, while IP multicasting supports arbitrary topologies, a UMTS network is a logical tree with the GGSN as the root and the UEs as the leaves.

Each network node needs to maintain additional state and support additional signaling procedures for MBMS. First, forwarding state is required so that an internal node may determine which of its children should receive packets addressed to a multicast group. Second, user state is required so that the network may charge multicast group members. Each node therefore maintains an *MBMS Bearer Context (MBC)* for each multicast group to facilitate forwarding, and an *MBMS UE Context (MUEC)* for each UE that is currently a member of that group to facilitate charging [5], as shown in Figure 2.

The MBC contains group-specific information, such as its IP multicast address and QoS parameters. A table in the MBC indicates which downstream nodes should receive packets addressed to that group. For example, in Figure 2 child #1 should receive packets (marked 1) but child #2 should not (marked 0). When a multicast packet arrives, the node examines the MBC for the appropriate group and forwards the packet to all children marked 1. The MUEC contains UE specific information for a particular group. The network uses the MUECs linked to an MBC via their multicast address to charge the UEs belonging to that group.

When an MBMS service is about to be offered, its attributes are first administratively entered into a new MBC at the BM-




Fig. 3. MBMS signaling procedures.

SC. Additional MBCs and MUECs are dynamically created and deleted in the network based on the UE-initiated signaling, as shown in Figure 3; solid arrows in the figure denote per UE signaling messages, while dotted arrows denote per group signaling messages. A UE desiring to join an MBMS group sends an *Internet Group Management Protocol (IGMP)* [6] join message to the GGSN serving the UE, thus triggering the MBMS multicast *activation* procedure (signaling phases 1 and 2 in Figure 3). Note that MBMS uses a join/leave mode of IGMP, which is more suitable for UMTS networks than the standard query/response mode of IGMP [7]. Signaling phase 1 results in the UE receiving from the BM-SC the APN of the GGSN that acts as the data source, which may differ from the GGSN that originally received the join, so that signaling phase 2 may continue toward the appropriate GGSN. In signaling phase 2, the SGSN, GGSN and BM-SC each create a MUEC for the UE, shown as a filled square in Figure 3.

When the first MUEC for a group is created at a GGSN or SGSN during signaling phase 2, that node sends a *registration* message to its parent (BM-SC or GGSN, respectively). The parent marks the corresponding entry in its MBC so as to start forwarding data to that child. Using the information provided in the response the child also creates an MBC for the group, shown as a filled circle in Figure 3; recall that the BM-SC is already administratively configured with the MBC for the group. The result of the activation and registration procedures is the creation of a multicast distribution tree from the GGSN towards all UEs participating in an MBMS service. In signaling phase 3, a *session start* procedure establishes radio

bearers before data transmission starts. After data transmission completes, in signaling phase 4 a *session stop* procedure releases these radio bearers.

The multicast *deactivation* procedure (signaling phases 5 and 6 in Figure 3) is initiated by the UE sending an IGMP leave message to the GGSN serving the UE; it is shown after the session stop procedure, but it can be triggered at any time. This procedure essentially reverses the activation procedure: in signaling phase 5 it redirects the UE towards the GGSN that acts as the data source, while in signaling phase 6 it causes the MUEC for the UE to be deleted from the appropriate nodes, shown as an unfilled square in Figure 3. When the last MUEC for a group is deleted at a GGSN or SGSN during signaling phase 6, that node sends a *deregistration* message to its parent and deletes the MBC for the group, shown as an unfilled circle in Figure 3. On receiving a deregistration, the parent marks the corresponding entry in its MBC so as to stop forwarding data to that child.

When the session start procedure indicates to an RNC that it should establish radio bearers, the RNC must decide what type of channel it should employ in each cell. Two options exist: either separate *Point to Point* (PtP) links towards each participating UE using the *Dedicated Transport Channel* (DCH), or a common *Point to Multipoint* (PtM) link towards all participating UEs using the *Forward Access Channel* (FACH). With the FACH, a single transmission reaches all UEs, regardless of their number and position, but the FACH must transmit at a high enough power level to reach the edge of the cell. With the DCH, a separate transmission is needed for each UE, but at the minimum power level required to reach it.

As the energy required to reach a UE scales faster than linearly with distance, for a few UEs it is normally more economical to employ multiple DCHs rather than a single FACH. As the network only tracks UEs with active signaling connections however, the RNC must estimate how many UEs participate in the group in each cell. This is achieved via a procedure known as *UE Counting* [8], whereby the RNC asks the UEs participating in a multicast group to establish a signaling connection with a specific probability; based on the number of UEs establishing such connections, the RNC estimates the total number of UEs belonging to the group. If this number is lower than an administratively defined threshold T , individual DCHs are established, otherwise a common FACH is established. Ideally, T should be set so that the average power required to transmit T copies of a packet over the DCHs would be roughly equal to the power required to transmit it over the FACH. The optimal threshold T depends on the position of each UE participating in the group. However, tracking all UEs as they move is uneconomical, therefore a simple estimated threshold T is used.

III. THE EXTENDED MBMS MODEL

Our MCV MBMS model extends the standard MBMS by allowing a single MBMS service to offer different variants of the same content to different subsets of receivers, so as to maximize the number of satisfied users [9]. This can be




Fig. 4. Extended MBMS Bearer and UE contexts.

achieved without inflating the original data stream, provided that the content variants have the property that a lower quality variant can be directly derived from a higher quality one. From a business perspective, the variants to offer must be chosen to match the specifications of common terminals, but their number should not be too large, as this would cause multicast groups to degenerate to individual receivers. In the following discussion we assume three variants, numbered 1 (low quality, LQ), 2 (medium quality, MQ) and 3 (high quality, HQ).

In MCV MBMS a UE indicates its preference for a content variant by simply including in its IGMP join message the variant's number; this number is also included in the MBMS request messages shown in Figure 3 as part of the activation and registration procedures (signaling phases 1 and 2). The UE may later modify its preference by sending a new IGMP join, without the need to first leave and then join the group again; if the service employs MBMS security, this obviates the need to generate and distribute new keys, as group membership remains the same [10]. For example, the user may request higher quality audio to better hear a passage or the terminal may request lower quality audio when the bandwidth at its location is limited.

In the MCV MBMS model, each node maintains additional information in the MBC and MUEC, shown in Figure 4 with a gray background. The MUEC is extended with the content variant requested by the UE during the activation procedure, so that the UE may be charged accordingly. The downstream node table in the MBC is extended with the number of the variant to forward to each child, for example, 0 (none) to 3 (high quality). Each internal node asks its parent during the registration procedure for the highest quality variant requested by any of its own children, as from this variant it can produce any lower quality ones. Hence, each node forwards the lowest possible amount of data to each of its children.

A node determines the highest quality variant requested by its children by storing in the MBC a set of counters for the number of UEs per variant, as shown in Figure 4. To maintain these counters, during the activation and deactivation procedures, at the point where in the standard MBMS a node would create or delete a MUEC (shown with filled and unfilled squares, respectively, in Figure 3), in MCV MBMS it must do one of the following:

- 1) If an MUEC was just created, the counter for its variant is incremented at the MBC.

- 2) If an MUEC was just deleted, the counter for its variant is decremented at the MBC.
- 3) If an UE just modified its variant by sending an updated join, the counter for its previous variant is decremented and the counter for its current variant is incremented at the MBC.

To maintain the forwarding table in the MBC, at the point where in the standard MBMS a node would create or delete an MBC (shown with filled and unfilled circles, respectively, in Figure 3), in MCV MBMS it must do one of the following:

- 1) If the first MUEC for the group was created, the MBC is created and the parent is asked to start forwarding data via the registration procedure.
- 2) If the last MUEC for the group was deleted, the MBC is deleted and the parent is asked to stop forwarding data via the deregistration procedure.
- 3) If the counter for a higher quality variant than the current one became nonzero, the parent is informed to forward that variant via the registration procedure.
- 4) If the counter for the current variant became zero, the next nonzero counter is found and the parent is informed to forward that variant via the registration procedure.

Counter updates add a fixed cost to message processing, while variant numbers add a fixed overhead to some signaling messages. The signaling flow is the same as in the standard MBMS; the only departure is that in MCV MBMS a UE or an internal node may change its current variant, thus triggering an updated activation or registration procedure.

As stated above, MCV MBMS requires lower quality variants to be directly derived from the higher quality ones. This can be achieved by layered coding [11], where the source encodes the lowest quality variant as the base layer and then encodes a series of successive enhancement layers. The first higher quality variant consists of the base layer and the first enhancement layer; each successive variant adds another enhancement layer. The source injects all layers to the multicast distribution tree and each node forwards to each child only the layers required to reconstruct the variant requested by that child. For example, in the MBC shown in Figure 4, the node needs to receive variant 2 from its parent, that is, the base layer and the first enhancement layer; it would forward only the base layer to child #1 and both the base layer and the first enhancement layer to child #3.

In our MBMS simulator, we implemented layered coding with three layers. Each packet contains data from a single layer and is tagged with a layer number, allowing nodes to easily drop redundant packets; no other processing is required at internal nodes. The resulting data stream consists of a base and two enhancement layers; users requesting the LQ variant only receive the base layer, users requesting the MQ variant receive the base and first enhancement layers, and users requesting the HQ variant receive the base and both enhancement layers.

Finally, in MCV MBMS during the session start procedure the RNC must first estimate the number of UEs that have requested each variant, and then decide whether to establish individual PtP bearers or a common PtM bearer for each layer. If the number of UEs requesting the base layer is smaller than

the administratively defined threshold T , then all layers will be transmitted over multiple DCH links, while if the number of UEs requesting all layers is not smaller than T , then all layers will be transmitted over a single FACH link; in these cases MCV and the standard MBMS behave alike. The third case is unique to MCV MBMS: when the number of UEs is not in the above regions, the lower layers may be transmitted over a common FACH link, while the higher layers may be transmitted over separate DCH links, depending on the number of UEs requesting each variant.

IV. PERFORMANCE EVALUATION: ANALYSIS

A. Definitions and Assumptions

For MCV MBMS to be useful, it must be more economical in terms of signaling load and data transmissions than the standard MBMS. In the following two sections, we compare the performance of MCV MBMS against three alternative options for satisfying the needs of heterogeneous UEs based solely on standard MBMS mechanisms. This section provides an analytical performance evaluation, while the following section adds simulation results. The *Single MBMS* approach uses a single MBMS service to transmit all layers; it incurs the minimum possible signaling overhead, but wastes transmission resources as it delivers data to UEs that have not requested it. The *Multiple MBMS* approach uses a separate MBMS group per layer, with each UE joining the groups required to decode the desired variant; it transmits the same data as MCV MBMS, but requires additional signaling since some UEs need to join multiple groups. Finally, *Simulcast MBMS* uses a separate MBMS group per variant, with each UE joining the desired group only, meaning that the per UE signaling load is similar to that of single MBMS, but the per group signaling load is similar to that of multiple MBMS; its main disadvantage however is that lower layers are transmitted multiple times, once for each group.

As each MBMS service is independent, we can study a single service with no loss of generality. We assume that N UEs participate in a given multicast MBMS service in a network where U UEs exist in total, and that the network is controlled by a single GGSN, S SGSNs and R RNCs. We also assume that the N participating UEs are uniformly distributed between the SGSNs and RNCs and denote the average number of UEs served by each SGSN or RNC as $N_S = N/S$ or $N_R = N/R$. If we further assume that the N participating UEs are distributed in C cells in proportion to the total number of users in each cell, then the expected number of users N_i in a cell i where U_i potential users exist is $N_i = (U_i/U)N$.

In order to provide multiple content variants, layered coding is used to generate layer 1, the base layer, layer 2, the first enhancement layer, and layer 3, the second enhancement layer. An UE requesting layer $n > 1$ must also receive layers 1 to $n - 1$ before decoding, that is, an LQ UE only needs layer 1, an MQ UE needs layers 1 and 2, and an HQ UE needs layers 1, 2 and 3. We assume that after layered coding B equally sized packets are injected into the network, with the probability that each packet is part of layer 1, 2 and 3 being l_1 , l_2 and l_3 , respectively, where $l_1 + l_2 + l_3 = 1$. We further

TABLE I
SUMMARY OF NOTATION.

Symbol	Description
U	Number of UEs in the network
N	Number of UEs participating in a service
S	Number of SGSNs in the network
N_S	Average number of UEs served by each SGSN
R	Number of RNCs in the network
N_R	Average number of UEs served by each RNC
C	Number of cells in the network
U_i	Number of UEs in cell i
N_i	Number of UEs participating in a service in cell i
B	Number of packets generated by the source
l_i	Probability that a packet belongs to layer i
v_i	Probability that a UE requests variant i
P_D	Average DCH transmission power for data packets
P_F	FACH transmission power for data packets
T	Switching threshold between DCH and FACH

assume that the participating UEs request the LQ, MQ, or HQ variant with probabilities v_L , v_M or v_H , respectively, where $v_L + v_M + v_H = 1$. Finally, we denote the average power required to transmit a data packet via the DCH by P_D , the power required to transmit it via the FACH by P_F and the switching threshold between the DCH and FACH by T . This notation is summarized in Table I.

To simplify the control plane analysis, we assume that at least one UE requesting the HQ variant is served by each SGSN; for Simulcast MBMS in particular, we assume that one UE requesting each variant is served by each SGSN. For the user plane analysis, we assume that at least one participating UE is located in each cell. When N is small, these assumptions cannot always hold, therefore our analytical predictions of signaling and user plane load will overestimate the actual values for small N . For MCV MBMS, we also assume that data packets exploit unused header bits to store their layer number, without increasing their size; for reasonably large data packets, this overhead would be negligible in any case.

Finally, two extreme assumptions may be made for the number of (de)registrations in MCV MBMS. In the best case, the first UE to join the service in each cell is an HQ UE, thus causing all layers to be requested with a single registration, and the last one to leave it is also an HQ UE, this causing all layers to be dropped with a single deregistration. In the worst case, first the LQ, then the MQ and then the HQ UEs join each service, thus causing each layer to be requested via a separate registration; for deregistrations this sequence is reversed, thus causing each layer to be dropped via a separate deregistration. As the actual sequence of joins and leaves is unknown, we simply assume that the average number of (de)registration packets required is the arithmetic mean of these two extreme cases.

B. Control Plane Analysis

As Single MBMS only provides a single content variant, comparing MCV MBMS, Multiple MBMS or Simulcast MBMS against it indicates the signaling overhead required to provide multiple content variants in the extended or in

TABLE II
PACKETS AND BITS PER SIGNALING PROCEDURE.

Symbol	Description
A_p/A_b	Number of packets/bits for each activation
E_p/E_b	Number of packets/bits for each deactivation
R_p/R_b	Number of packets/bits for each registration
D_p/D_b	Number of packets/bits for each deregistration
SA_p/SA_b	Number of packets/bits for each session start
SE_p/SE_b	Number of packets/bits for each session stop

the standard MBMS model. We examine the cost of each MBMS procedure using the notation summarized in Table II for the number of packets and bits required in the standard MBMS. In MCV MBMS, the request packets in the activation and registration procedures are inflated by 8 bits to hold the desired content variant, while the session start request packets must hold an additional 128-bit QoS profile for each enhancement layer; the standard QoS profile in these packets is used for the base layer. The response packets for these procedures do not need a content variant field and are therefore not inflated, and the same is true for both the request and the response packets for the session stop, deregistration and deactivation procedures. As a result, only half of the packets in the activation, registration and session start procedures is inflated in MCV MBMS.

Starting with the activation, in Single MBMS it requires $A_p N$ packets and $A_b N$ bits for the entire network. In MCV MBMS, the same number of packets is required as each UE joins one group, but half of these packets is inflated by 8 bits, therefore the total number of bits required is $A_b N + 8A_p N/2$. In Multiple MBMS, LQ, MQ and HQ UEs join one, two and three groups, respectively, meaning that all UEs join the first group, $v_M + v_H$ UEs join the second group and v_H UEs join the third group, therefore activation requires $A_p(1 + v_M + 2v_H)N$ packets and $A_b(1 + v_M + 2v_H)N$ bits. In Simulcast MBMS each UE also joins a single group, therefore the same number of packets and bits is required as in Single MBMS. Similarly, in Single MBMS the deactivation requires $E_p N$ packets and $E_b N$ bits; the same applies to MCV MBMS as deactivation packets are not inflated, and to Simulcast MBMS as only a single group is used. In Multiple MBMS, some UEs leave multiple groups, therefore $E_p(1 + v_M + 2v_H)N$ packets and $E_b(1 + v_M + 2v_H)N$ bits are required.

For registration, the amount of packets and bits required depends on the number of SGSNs serving participating UEs, the content variants requested by these UEs and (for MCV MBMS) the exact ordering of the activations. For Multiple MBMS and MCV MBMS we assumed that at least one HQ UE is served by each SGSN, while for Simulcast MBMS we assumed that at least one UE requesting each variant is served by each SGSN; in all cases, this means that all SGSNs will eventually register for all variants. In Single MBMS, only a single registration is performed by each SGSN, requiring $R_p S$ packets and $R_b S$ bits for the entire network. In Multiple MBMS each SGSN needs to register to all three groups, requiring three times as many packets and bits, and the same

TABLE III
COMPARISON OF SIGNALING PACKETS.

Procedure	Single MBMS	MCV MBMS	Multiple MBMS	Simulcast MBMS
Activation	$A_p N$	$A_p N$	$A_p(1 + v_M + 2v_H)N$	$A_p N$
Deactivation	$E_p N$	$E_p N$	$E_p(1 + v_M + 2v_H)N$	$E_p N$
Registration	$R_p S$	$2R_p S$	$3R_p S$	$3R_p S$
Deregistration	$D_p S$	$(R_p + D_p)S$	$3D_p S$	$3D_p S$
Session Start	$SA_p S$	$SA_p S$	$3SA_p S$	$3SA_p S$
Session Stop	$SE_p S$	$SE_p S$	$3SE_p S$	$3SE_p S$

TABLE IV
COMPARISON OF SIGNALING BITS.

Procedure	Single MBMS	MCV MBMS	Multiple MBMS	Simulcast MBMS
Activation	$A_b N$	$(A_b + 4A_p)N$	$A_b(1 + v_M + 2v_H)N$	$A_b N$
Deactivation	$E_b N$	$E_b N$	$E_b(1 + v_M + 2v_H)N$	$E_b N$
Registration	$R_b S$	$(2R_b + 8R_p)S$	$3R_b S$	$3R_b S$
Deregistration	$D_b S$	$(R_b + 4R_p + D_b)S$	$3D_b S$	$3D_b S$
Session Start	$SA_b S$	$(SA_b + 128SA_p)S$	$3SA_b S$	$3SA_b S$
Session Stop	$SE_b S$	$SE_b S$	$3SE_b S$	$3SE_b S$

holds for Simulcast MBMS. In MCV MBMS, in the best case the HQ UEs will join first, requiring a single registration, while in the worst case the LQ UEs will join first, then the MQ UEs and then the HQ UEs, requiring three registrations. The arithmetic average is therefore two registrations requiring $2R_p S$ packets and $2R_b S + 2(8R_p S/2)$ bits, since half of these packets are inflated by 8 bits. Similarly, in Single MBMS the deregistration procedure requires $D_p S$ packets and $D_b S$ bits, while in Multiple MBMS and Simulcast MBMS it requires three times as many packets and bits. In MCV MBMS, in the best case the HQ UEs will leave last, requiring a single deregistration, while in the worst case the HQ UEs will leave first, then the MQ UEs and then the LQ UEs requiring two registration updates and a deregistration. The arithmetic average is one registration and one deregistration requiring $R_p S + D_p S$ packets and $R_b S + 8R_p S/2 + D_b S$ bits.

For the session start, the amount of packets and bits required also depends on the number of SGSNs serving participating UEs, but due to our assumptions, all SGSNs will eventually receive a session start. In Single MBMS only a single session start is needed, requiring $SA_p S$ packets and $SA_b S$ bits. In Multiple MBMS and Simulcast MBMS a session start is needed for all three groups, requiring three times as many packets and bits. In MCV MBMS the same number of packets is needed as only a single group exists, but half of these packets is inflated by two more QoS profiles, therefore the total number of bits required is $SA_b S + 2(128SA_p S)/2$. Similarly, in Single MBMS the session stop procedure requires $SE_p S$ packets and $SE_b S$ bits, while in Multiple MBMS and Simulcast MBMS it requires three times as many packets and bits. In MCV MBMS the same number of packets and bits are required as in Single MBMS since no extra fields are required. The total numbers of packets and bits required for each signaling procedure are summarized in Table III and IV.

C. User Plane Analysis

In order to compare the various approaches with respect to their user plane overhead, we first observe that Multiple MBMS and MCV MBMS are equivalent in this respect, as we have assumed that in MCV MBMS the layer number is stored in unused header bits of user plane packets. For the *Core Network* (CN) we focus on the packet processing requirements of each alternative, therefore we study the number of user plane packets received by internal nodes. For the *Radio Access Network* (RAN) we focus on the energy requirements of each alternative, therefore we study the energy consumed in each cell to transmit user plane packets [12].

Starting with the CN, since we have assumed that at least one UE has joined the service in each cell, each SGSN and each RNC will receive all user plane packets in Single MBMS; each node will receive all layers, that is, B packets, therefore the total user plane traffic is $(S + R)B$. In MCV/Multiple MBMS we must also take into account UE preferences. Since at least one UE has joined the service in each cell, each SGSN will receive at least the base layer. The first enhancement layer will be received by an SGSN if at least one of the N_S UEs that it controls has asked for either the MQ or the HQ variants; this probability is the complement of the probability that all these UEs have asked for the LQ variant, or $(1 - v_L^{N_S})$. The second enhancement layer will be received by an SGSN if at least one of the N_S UEs that it controls has asked for the HQ variant; this probability is the complement of the probability that none of these UEs have asked for the HQ variant, or $[1 - (1 - v_H)^{N_S}]$. We can treat the RNCs identically by replacing N_S with N_R . Therefore, the user traffic received by all SGSNs is $\{l_1 + (1 - v_L^{N_S})l_2 + [1 - (1 - v_H)^{N_S}]l_3\}SB$ packets, while the user traffic received by all RNCs is $\{l_1 + (1 - v_L^{N_R})l_2 + [1 - (1 - v_H)^{N_R}]l_3\}RB$ packets.

In Simulcast MBMS we can deal with each variant separately. Since at least one UE has joined the service in each cell, each SGSN will receive at least one variant. The LQ

variant will be received by an SGSN if at least one of the N_S UEs that it controls has asked for the LQ variant; this probability is the complement of the probability that none of these UEs have asked for the LQ variant, or $[1 - (1 - v_L)^{N_S}]$. We can treat the MQ and HQ variants in exactly the same manner for each SGSN, and then simply replace N_S with N_R in each formula to get the RNC requirements. Therefore, the user traffic received by all SGSNs is $\{[1 - (1 - v_L)^{N_S}]l_1 + [1 - (1 - v_M)^{N_S}](l_1 + l_2) + [1 - (1 - v_H)^{N_S}](l_1 + l_2 + l_3)\}SB$ packets, while the user traffic received by all RNCs is $\{[1 - (1 - v_L)^{N_R}]l_1 + [1 - (1 - v_M)^{N_R}](l_1 + l_2) + [1 - (1 - v_H)^{N_R}](l_1 + l_2 + l_3)\}RB$.

Turning to the RAN, the analysis of the user plane transmission power requirements must take into account the number of UEs that have requested a specific layer or variant in each cell: if it is less than the threshold T then multiple DCHs are used, otherwise a single FACH is used. In MCV/Multiple MBMS the transmission power required in cell i for the l_1B packets of the base layer depends on the fraction of the N_i participating users requesting it. Since $(v_L + v_M + v_H)N_i = N_i$, this power is:

$$P_{MCV1}(i) = \begin{cases} l_1BN_iP_D, & N_i < T \\ l_1BP_F, & N_i \geq T \end{cases}$$

For the l_2B packets of the first enhancement layer we note that $(v_M + v_H)N_i = (1 - v_L)N_i$, therefore:

$$P_{MCV2}(i) = \begin{cases} l_2B(1 - v_L)N_iP_D, & (1 - v_L)N_i < T \\ l_2BP_F, & (1 - v_L)N_i \geq T \end{cases}$$

Finally, for the l_3B packets of the second enhancement layer the fraction is v_HN_i , therefore:

$$P_{MCV3}(i) = \begin{cases} l_3Bv_HN_iP_D, & v_HN_i < T \\ l_3BP_F, & v_HN_i \geq T \end{cases}$$

Summing up over all C cells we find that the total transmission power consumed by MCV/Multiple MBMS is $\sum_{i=1}^C (P_{MCV1}(i) + P_{MCV2}(i) + P_{MCV3}(i))$. In Simulcast MBMS, we can again treat each variant separately. The transmission power required in cell i for the l_1B packets of the LQ variant depends on the fraction v_L of the N_i participating users requesting it:

$$P_{SC1}(i) = \begin{cases} l_1Bv_LN_iP_D, & v_LN_i < T \\ l_1BP_F, & v_LN_i \geq T \end{cases}$$

For the $(l_1 + l_2)B$ packets of the MQ variant we find similarly that:

$$P_{SC2}(i) = \begin{cases} (l_1 + l_2)Bv_MN_iP_D, & v_MN_i < T \\ (l_1 + l_2)BP_F, & v_MN_i \geq T \end{cases}$$

Finally, for the $(l_1 + l_2 + l_3)B = B$ packets of the HQ variant we find that:

$$P_{SC3}(i) = \begin{cases} Bv_HN_iP_D, & v_HN_i < T \\ BP_F, & v_HN_i \geq T \end{cases}$$

Summing up over all C cells we find that the total transmission power consumed by Simulcast MBMS is $\sum_{i=1}^C (P_{SC1}(i) + P_{SC2}(i) + P_{SC3}(i))$.




Fig. 5. Simulated topology.

Finally, in Single MBMS all layers are transmitted, therefore the total transmission power consumed is $\sum_{i=1}^C P_S(i)$ where:

$$P_S(i) = \begin{cases} BN_iP_D, & N_i < T \\ BP_F, & N_i \geq T \end{cases}$$

V. PERFORMANCE EVALUATION: SIMULATION

A. Simulation Setup

We have created a detailed MBMS simulator conforming to the common core of the 3GPP Release 6, 7 and 8 specifications, including full support for the MCV MBMS model described in this paper. The simulator uses the Opnet Modeler platform [13] and is available to the public as a contributed Opnet model [14]. In this section we compare MCV MBMS against the standard-based alternatives introduced in Section IV for a specific simulation scenario, applying in parallel the preceding analysis to the same scenario to compare our analytical predictions with the simulation results.

We simulated the topology shown in Figure 5, which consists of a single GGSN, two SGSNs ($S = 2$), four RNCs ($R = 4$) and six Node-Bs. Two Node-Bs control cells with 4 UEs (sparse cells) and four Node-Bs control cells with 9 UEs (dense cells). We varied the number N of participating UEs from 1 to 40 by randomly choosing UEs in each experiment and having them join in random order; due to occasional conflicts in the uplink, not all joins were successful, therefore this process continued until N UEs actually managed to join. We used two sets of parameters for the UE preferences: in scenario 1 each UE randomly selected a content variant with probabilities $v_L = 0.7$, $v_M = 0.2$ and $v_H = 0.1$, while in scenario 2 these probabilities were $v_L = 0.5$, $v_M = 0.25$ and $v_H = 0.25$. The UEs did not change their preferences over time, simplifying comparisons with the analysis. All simulation results shown indicate averages from 30 repetitions of each experiment with different seeds.

The MBMS service modeled was a stream of UDP packets with a payload of 968 bytes, generated every 0.125 s, or a data rate of roughly 62 Kbps, excluding the 28-byte UDP/IP header overhead. A total of $B = 1000$ user plane packets were generated in each experiment, randomly distributed to three layers with probabilities $l_1 = 0.5$, $l_2 = 0.25$ and $l_3 = 0.25$. At the radio link, each user plane packet was split into six segments of 1328 bits each; in the DCH an 8-bit header was added to each segment, while in the FACH a 32 bit header

was added. Each such segment was transmitted in a 20 ms *transmission time interval* (TTI). The resulting aggregate data rate was just below 64 Kbps, which is the highest FACH speed supported by the simulator. While a 64 Kbps data rate is low, the results shown below can be easily adapted to higher data rates: the control plane results remains exactly the same, while the user plane results for both the CN and the RAN scale linearly with the total number of packets. As shown in the preceding analysis, the signaling overhead does not depend on the data rate, while the user plane overhead only depends on the volume of data, not the data rate.

Each Node-B used a single sector antenna with a maximum transmission power of 20 W, covering a cell with a radius of 1 km. The path loss model used was the *Outdoor to Indoor and Pedestrian*, described by $L = 40\log_{10}d + 30\log_{10}f + 49$ dB, where d is the distance between the UE and the Node-B in km and f is the carrier frequency in MHz, which was 2110 MHz in our case. This model is valid for non-line of sight cases and describes the worst-case propagation [15]. The shadow fading loss was modeled as a log-normal random variable with zero mean and variance 10 dB, which is a valid assumption for outdoor users [15].

The UEs in each cell were manually placed so that half of them were close and half of them were far from the Node-B antenna. The transmission power of the FACH was set to $P_F = 0.4$ W, which is sufficient to cover UEs at a distance of *at least* $\frac{2}{3}$ of the cell radius, or 50% of the cell area; in practice, packet losses were negligible for all UEs. The transmission power of each DCH varied due to the outer loop power control modeled by the simulator. The average DCH transmission power in our experiments was measured to be $P_D = 0.08975$ W, therefore we set the threshold $T = 5$, that is, the FACH was used to serve 5 or more UEs, as $4P_D < P_F < 5P_D$. Note that this choice of T means that the sparse cells (with 4 UEs) never used the FACH, something quite reasonable for some cells in a real network.

B. Control Plane Results

Due to simulator limitations, the results presented below only cover the activation, registration, session start and session stop procedures shown in Figure 3 (signaling phases 1 to 4), omitting the deactivation and deregistration procedures (signaling phases 5 and 6). According to the relevant standards, for these procedures $A_p = 8$, $R_p = 2$, $SA_p = 2$ and $SE_p = 2$ packets, while $A_b = 2336$, $R_b = 472$, $SA_b = 992$ and $SE_b = 400$ bits [16], [17]. By applying the analysis of Section IV-B to our scenario and substituting values from the standards, we find that the predicted signaling overhead in terms of packets for Single MBMS is $12 + 8N$, while for MCV MBMS it is $16 + 8N$. For Multiple MBMS the predicted signaling overhead in terms of packets is $36 + 11.2N$ for scenario 1 and $36 + 14N$ for scenario 2, while for Simulcast MBMS it is $36 + 8N$. These predictions as well as the actual simulation results for each scenario are shown in Figures 6(a) and 6(b). Similarly, we find that the predicted signaling overhead in terms of bits for Single MBMS is $3728 + 2336N$, while for MCV MBMS it is $5216 + 2368N$. For Multiple MBMS the predicted signaling

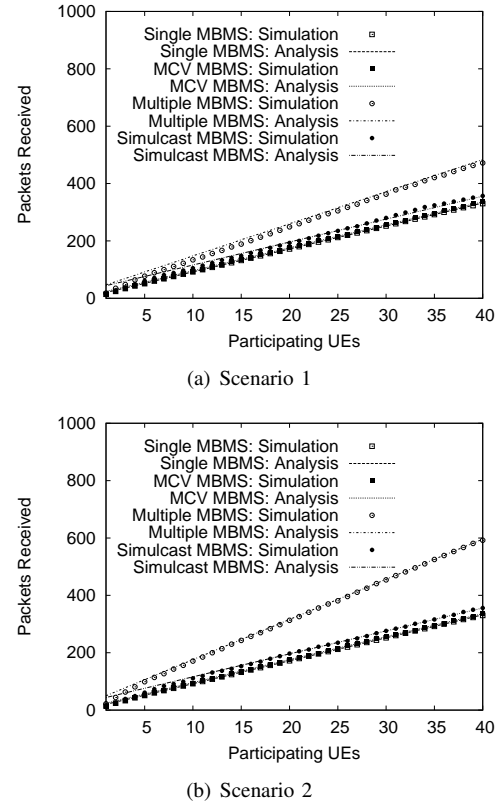



Fig. 6. Signaling overhead in terms of packets.


overhead in terms of bits is $11184 + 3270.4N$ for scenario 1 and $11184 + 4088N$ for scenario 2, while for Simulcast MBMS it is $11184 + 2336N$. These predictions as well as the actual simulation results for each scenario are shown in Figures 7(a) and 7(b).

The figures show similar trends for packets and bits and a nearly perfect match between the analytical predictions and simulation results. The slight overestimation in the analysis for Multiple MBMS and Simulcast MBMS is due to the assumption that at least one UE per SGSN has requested the HQ variant (Multiple MBMS) or each variant (Simulcast MBMS), which is unlikely with a few UEs. We observe that MCV MBMS is only marginally more costly than Single MBMS which incurs the minimum possible overhead. For 40 UEs, the simulation results indicate that the extra cost of MCV MBMS compared to Single MBMS is only 1.6% in terms of packets and 3.2% in terms of bits; the latter is relatively higher due to the inflated request packets of MCV MBMS. The additional signaling cost of MCV MBMS compared to Single MBMS is therefore negligible.

The figures also show that the signaling overhead of Multiple MBMS is sensitive to UE preferences: as the proportion of MQ and HQ UEs increases, so does the signaling overhead since more UEs need to join additional groups. We observe a wide gap in favor of MCV MBMS, which grows with the number of participating UEs. For 40 UEs the simulation results for scenario 1 indicate that the extra cost of Multiple MBMS compared to MCV MBMS is 40% in terms of packets and 38% in terms of bits; in scenario 2 the corresponding figures are




(a) Scenario 1




(b) Scenario 2

Fig. 7. Signaling overhead in terms of bits.



(a) Scenario 1



(b) Scenario 2

Fig. 8. User plane overhead in the CN in terms of packets.

77% and 74%. Therefore MCV MBMS provides a dramatic performance improvement over Multiple MBMS. Finally, the signaling overhead of Simulcast MBMS is slightly higher than MCV MBMS: even though each UE only joins a single group, multiple groups need to be maintained in the network. For 40 UEs the simulation results indicate that the extra cost of Simulcast MBMS compared to MCV MBMS is 6.2% in terms of packets and 5.0% in terms of bits, therefore MCV MBMS provides an improvement even over Simulcast MBMS.

C. User Plane Results

By applying the analysis of Section IV-C for the CN to our scenario, we find that the predicted number of received user plane packets with Single MBMS is simply 6000. With MCV/Multiple MBMS the expected number of received user plane packets is $6000 - 500(0.7^{N/2} + 0.9^{N/2}) - 1000(0.7^{N/4} + 0.9^{N/4})$ for scenario 1 and $6000 - 500(0.5^{N/2} + 0.75^{N/2}) - 1000(0.5^{N/4} + 0.75^{N/4})$ for scenario 2. Finally, with Simulcast MBMS the expected number of received user plane packets is $13500 - 1000(0.3^{N/2}) - 1500(0.8^{N/2}) - 2000(0.9^{N/2}) - 2000(0.3^{N/4}) - 3000(0.8^{N/4}) - 4000(0.9^{N/4})$ for scenario 1 and $13500 - 1000(0.5^{N/2}) - 3500(0.75^{N/2}) - 2000(0.5^{N/4}) - 7000(0.75^{N/4})$ for scenario 2. These predictions as well as the actual simulation results for each scenario are shown in Figures 8(a) and 8(b).

These analytical predictions closely match the simulation results for 15 UEs or more; with fewer UEs the assumption that at least one UE has joined the service in each cell is unrealistic, therefore the analysis overestimates reality. As

expected, Single MBMS requires nearly constant traffic, as all layers are sent to all nodes. With MCV/Multiple MBMS the user plane traffic depends on UE preferences: as the proportion of MQ and HQ UEs increases, so does the number of user plane packets that need to be forwarded. In both the scenarios, as the number of UEs grows MCV/Multiple MBMS traffic tends to reach that of Single MBMS; with a large enough number of UEs, all SGSNs and RNCs will serve at least one HQ UE, therefore they will receive all user plane packets as in Single MBMS. Simulcast MBMS on the other hand transmits each variant independently, without taking advantage of layered coding: with a large enough number of UEs, all SGSNs and RNCs will receive all variants. For 40 UEs the simulation results indicate that the extra cost of Simulcast MBMS compared with Single MBMS is 92% in scenario 1 and 118% in scenario 2; with more UEs the extra cost would reach 125%.

We next apply the analysis of Section IV-C for the RAN to our scenario, in order to estimate the energy consumption for user plane packets. Energy consumption can be calculated by multiplying the power required to transmit each packet by the time taken to transmit it; we plot the analytical predictions against the simulation results for both scenarios in Figures 9(a) and 9(b). Note that these simulation results also include the energy consumed for control plane packets, since their number, and therefore their energy consumption, is negligible compared to that of user plane packets. The match between analysis and simulation is excellent for up to 25 UEs. The gap that appears at this point is due to the assumption in the

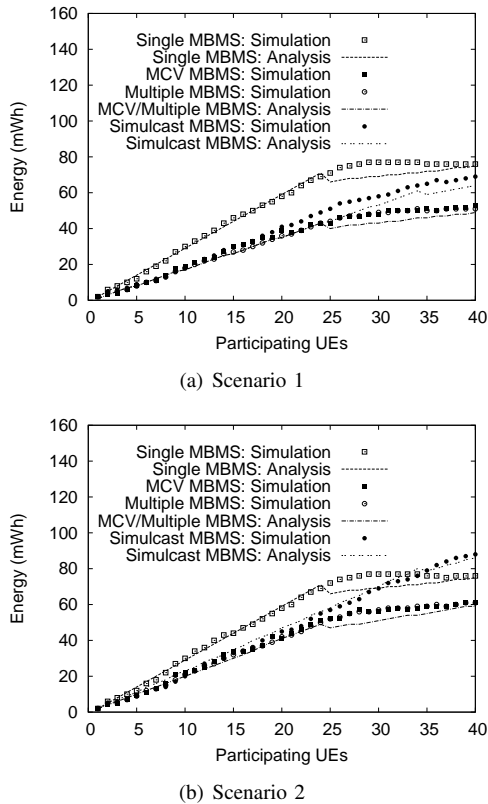


Fig. 9. Total energy consumed for transmissions.

analysis that all cells host the same (fractional) number of UEs, thus causing dense cells to switch from the DCH to the FACH at the exact same point. In the, more realistic, simulation, each cell hosts an integer number of UEs, therefore cells switch to the FACH at different points. As the number of UEs is increased, all dense cells switch to the FACH, therefore the gap closes again.

With a small number of UEs, MCV/Multiple MBMS will not transmit all packets belonging to layers not requested by any UEs in a cell, unlike Single MBMS which always transmits all layers in each cell. With a large number of UEs, when all layers are likely to be requested by some UEs in a cell, MCV/Multiple MBMS will not use the FACH to transmit layers requested by less than T UEs, while Single MBMS will use the FACH to transmit all layers even if only a single layer has been requested by T or more UEs. Even with a very large number of UEs in a cell, causing all layers to be requested by T or more UEs, MCV/Multiple MBMS would never lead to a higher energy consumption than Single MBMS. This means that MCV/Multiple MBMS can provide considerable energy savings: averaged over the entire range of 1 to 40 UEs, the simulation results show that MCV/Multiple MBMS consumes 35% less energy than Single MBMS in scenario 1 and 25% less energy in scenario 2.

Simulcast MBMS on the other hand may lead to lower or higher energy consumption than Single MBMS; Figure 9(b) shows that with a large enough number of UEs in scenario 2, Simulcast MBMS is more costly than Single MBMS. The reason is again that Simulcast MBMS does not take advantage

of layered coding: with a large enough number of UEs in a cell, all variants are separately transmitted in that cell, leading to duplicate transmissions of the same data. Since even the dense cells in our topology host relatively few UEs, this rarely occurs in our scenarios: averaged over the entire range of 1 to 40 UEs, the simulation results show that Simulcast MBMS consumes 25% less energy than Single MBMS in scenario 1 and 12% less energy in scenario 2, which is consistently worse than MCV/Multiple MBMS.

VI. RELATED WORK AND CONTRIBUTION

Addressing the issues raised by network and terminal heterogeneity for multimedia multicasting has been an important research topic for a long time, and the use of layered coding to handle these issues is a well-established idea. Early work in this area concentrated on selecting the optimal bit rates and routes for each content layer, assuming a static network topology and fixed user requirements [18]. After IP multicasting was proposed, most researchers assumed that a separate multicast group would be used per layer, with receivers joining the groups corresponding to the layers they could receive without excessive congestion losses; the underlying routing algorithm would take care of the rest.

In these schemes, a receiver can adapt to prevailing conditions by occasionally requesting an additional layer; when congestion appears, detected by excessive packet loss, layers are dropped until the network is stabilized. To prevent wasting bandwidth by constantly adding and dropping layers without coordination, one approach is for each receiver to monitor the attempts of its neighbors to add layers, thus automatically detecting the state of the network [19]. Another approach is to elect one receiver in each network area to coordinate these attempts and prevent conflicts [20]. The introduction of wireless links into the Internet renewed interest in such schemes, as wireless receivers also face unpredictable wireless losses and bandwidth variations induced by mobility [21].

Our approach departs from that work as we assume a different environment, thus avoiding many of the problems that plague IP multicast. First, multicast routing becomes trivial in MBMS, since a single distribution tree is used. Second, network resources can be reserved, therefore there is no inherent need to adapt to congestion. Third, due to the tree-based routing, our approach can easily handle layered coding, including both non-standard [22] and standard-derived schemes [23]; it can also be combined with hierarchical multicast error recovery [24].

Although MBMS avoids many of the problems that multimedia multicasting faces on the Internet, it introduces its own complications; the contributions of this paper lie exactly in addressing these complications. First, MBMS has already been standardized, thus changes must be kept to a minimum and they must be evaluated with respect to the additional load that they place on the network. Our proposal requires trivial changes to the standard MBMS: a single field is added to some messages, a few extra fields are added to the state maintained by each node and only constant cost state management processing is required.

Second, since a single distribution tree exists, it is critical to avoid overloading higher level nodes. Our proposal avoids the signaling overhead incurred by the management of multiple multicast groups, while allowing each node to easily select the content variant to forward toward the receivers. In addition, state and signaling are aggregated toward the root of the tree, exactly as in the standard MBMS. We quantify the gains of our single group approach in terms of control plane overhead in Section V-B and in terms of user plane overhead in the CN in Section V-C, showing that the cost of satisfying all users with MCV MBMS is minimal.

Third, since in UMTS networks the most important resource is the power required for transmissions over the air interface, our approach explicitly takes into account the issue of channel selection for each content layer transmitted, making the optimal choice between PtP and PtM transmissions separately for each layer. In Section V-C, we quantify the gains from our approach, showing that MCV MBMS leads to considerable energy savings.

Fourth, since analytical predictions often diverge from reality, we compare our predictions for control and user plane overhead with the results obtained from a detailed MBMS simulation model, showing a close match between analysis and simulation. This is especially important as energy consumption is heavily dependent on the actual number of UEs in each cell, exhibiting a discontinuous behavior at the switching points between PtP and PtM channels which is hard to capture by analytical modeling alone.

As the energy consumption of multimedia services is a critical factor for the success of MBMS [4], many researchers are trying to improve the energy efficiency of MBMS. One option is to choose between PtP and PtM channels based not only on the number of UEs, but also on their actual positions in the cell, so as to switch at the actual optimal point [25]. Another option is to employ a single PtM channel for UEs close to the center of the cell and separate PtP channels for the remaining UEs [25], [26]. Both options can be combined with MCV MBMS to optimize the energy consumption separately for each layer.

Another approach is to use layered coding to provide different levels of service to UEs depending on their location. One option is to transmit only the base layer with sufficient energy to cover the entire cell, using less energy for the enhancement layer(s), thus providing better service to UEs close to the center of the cell [25], [27]; remote UEs may receive the enhancement layer(s) via PtP channels [28]. Another option is to use hierarchical modulation within a single transmission stream, thus allowing UEs close to the center of the cell to receive more bits per modulation symbol [29]. These solutions are orthogonal to MCV MBMS, as they differentiate between UEs based on their reception conditions rather than their preferences.

Considerable research has also been motivated by the standardization of Raptor codes as an *Application Layer Forward Error Correction* (AL-FEC) scheme for MBMS [30]. Raptor codes provide efficient error correction without feedback from the, possibly huge, number of receivers of an MBMS service. There is a tradeoff between the amount of AL-FEC and

physical layer coding applied to each stream, which can be exploited to reduce the energy consumption [31]. This work can also be combined with ours to improve the decoding probability for each layer.

VII. CONCLUSIONS

We have presented an extension to MBMS that supports the distribution of multiple variants of the same content to heterogeneous receivers, aiming to maximize the number of UEs participating in an MBMS service while minimizing the amount of user and control traffic transmitted. We evaluated our MCV MBMS model against alternatives based on the standard MBMS via both analysis and simulation, showing a close match between the two approaches. We found that MCV MBMS requires less control plane overhead than the standard-based multiple service and simulcast approaches, being only slightly more expensive than the baseline single service approach. In the user plane, we found that MCV MBMS can satisfy all users with considerable savings in the energy consumption over the radio link compared with both the baseline single service approach and the standard-based simulcast approach.

ACKNOWLEDGEMENT

The work reported in this paper was partially supported by the European Union FP6/IST project B-Bone under contract IST-2003-507607. The simulation platform was provided by the OPNET University Program. The MBMS simulator was jointly produced by the partners in the B-Bone project.

REFERENCES

- [1] 3GPP. Multimedia broadcast/multicast service (MBMS); Stage 1. *TS 22.146; V8.4.0* June 2008.
- [2] 3GPP. Multimedia broadcast/multicast service (MBMS) user services; Stage 1. *TS 22.246; V8.5.0* March 2008.
- [3] Cheriton D, Deering S. Host groups: A multicast extension for datagram internetworks. *Proceedings of the Data Communications Symposium*, Whistler Mountain, BC, Canada, 1985.
- [4] Xylomenos G, Vogkas V, Thanos G. The multimedia broadcast / multicast service. *Wireless Communications and Mobile Computing*, 2008; **8**(2): 255–65.
- [5] 3GPP. Multimedia broadcast/multicast service (MBMS) user services; Architecture and functional description. *TS 23.246; V8.4.0* June 2009.
- [6] Fenner W. Internet Group Management Protocol, Version 2. RFC 2236, November 1997.
- [7] Xylomenos G. Group management for the multimedia broadcast/multicast service. *Proceedings of the IST Mobile Summit*, Dresden, Germany, June 2005.
- [8] 3GPP. Introduction of the multimedia broadcast multicast service (MBMS) in the radio access network (RAN); Stage 2. *TS 25.346; V8.4.0* December 2009.
- [9] Xylomenos G, Katsaros K. A multiple content variant extension of the multimedia broadcast / multicast service. *Proceedings of the IST Mobile & Wireless Communications Summit*, Myconos, Greece, June 2006.
- [10] 3GPP. Security of multimedia broadcast/multicast service. *TS 33.246; V8.4.0* June 2009.
- [11] Li B, Liu J. Multirate video multicast over the Internet: An overview. *IEEE Network*, 2003; **17**(5): 55–63.
- [12] Xylomenos G, Tsakanikas V, Polyzos GC. Reducing the transmission power requirements of the multimedia broadcast/multicast service. *Proceedings of the IST Mobile & Wireless Communications Summit*, Budapest, Hungary, July 2007.
- [13] Opnet Modeler Home. Available at http://www.opnet.com/solutions/network_rd/modeler.html.

- [14] Opnet Model Library Home. Available at http://www.opnet.com/support/des_model_library/index.html.
- [15] 3GPP. Selection procedures for the choice of radio transmission technologies of the UMTS. *TR 30.03U; V3.2.0* April 1998.
- [16] 3GPP. General packet radio service (GPRS); GPRS tunneling protocol (GTP) across the Gn and Gp interface. *TS 29.060; V8.11.0* March 2010.
- [17] 3GPP. Mobile radio interface Layer 3 specification; core network protocols; Stage 3. *TS 24.008; V8.9.0* March 2010.
- [18] Shacham N. Multipoint communication by hierarchically encoded data. *Proceedings of the IEEE INFOCOM*, Florence, Italy, May 1992.
- [19] Jacobson V, McCanne S, Vetterli M. Receiver-driven layered multicast. *Proceedings of the ACM SIGCOMM*, Palo Alto, CA, USA, August 1996.
- [20] Li X, Paul S, Ammar M. Layered Video Multicast with Retransmission (lvmr): Evaluation of hierarchical rate control. *Proceedings of the IEEE INFOCOM*, San Francisco, CA, USA, March 1998.
- [21] Wu D, Hou YT, Zhang Y-Q. Scalable video transport over wireless IP networks. *Proceedings of the IEEE Symposium on Personal, Indoor and Mobile Radio Communications*, London, UK, September 2000.
- [22] McCanne S, Vetterli M, Jacobson V. Low-complexity video coding for Receiver-driven Layered Multicast. *IEEE Journal on Selected Areas in Communications*, 1997; **15**(6): 983–1001.
- [23] Radha H. TransScaling: A video coding and multicasting framework for wireless IP multimedia services. *Proceedings of the ACM Workshop on Wireless Mobile Multimedia*, Rome, Italy, 2001.
- [24] Li X, Paul S, Pancha P, Ammar M. Layered Video Multicast with Retransmission (lvmr): evaluation of error recovery schemes. *Proceedings of the IEEE Workshop on Network and Operating System Support for Digital Audio and Video*, St. Louis, MO, USA, May 1997.
- [25] Alexiou A, Bouras C, Kokkinos V. Power saving techniques in MBMS multicast mode. *Proceedings of the ACM Workshop on Wireless Multimedia Networking and Performance Modeling*, Chania, Greece, 2007.
- [26] Alexiou A, Bouras C, Kokkinos V, Rekkas E. A Novel Mechanism for Radio Capacity Maximization during MBMS Transmissions in B3G Networks. *Proceedings of the ACM International Workshop on Modeling Analysis and Simulation of Wireless and Mobile Systems*, Vancouver, BC, Canada, 2008.
- [27] Chaudhry A, Jamil YK. An efficient MBMS content delivery scheme over the HSDPA network. *Proceedings of the International Symposium on Wireless Pervasive Computing*, Melbourne, VIC, Australia, 2009.
- [28] Pang A-C, Yang S-R, Yang C. Layer-encoded MBMS channel allocation for mobile networks. *Proceedings of the IEEE Conference on Local Computer Networks*, Dublin, Ireland, 2007.
- [29] Correia A, Silva J, Souto N, Silva L, Boal A, Soares A. Multi-resolution broadcast/multicast systems for MBMS. *IEEE Transactions on Broadcasting*, 2007; **53**(1): 224–34.
- [30] 3GPP. Multimedia broadcast/multicast service (MBMS); Protocols and codecs. *TS 26.346; V8.4.0* September 2009.
- [31] Luby M, Gasiba T, Stockhammer T, Watson M. Reliable multimedia download delivery in cellular broadcast networks. *IEEE Transactions on Broadcasting*, 2007; **53**(1) :235–4.

AUTHORS' BIOGRAPHIES

George Xylomenos is a member of the Mobile Multimedia Laboratory at the Athens University of Economics and Business, Greece, where he is an Assistant Professor of Computer Science. He received his Diploma in Informatics (1993) from the Athens University of Economics and Business, Greece, and M.S. (1996) and Ph.D. (1999) degrees in Computer Science from the Department of Computer Science and Engineering at the University of California, San Diego. His research interests include multicast based content distribution, information centric network architectures, provision of Quality of Service over wireless and mobile networks and supporting broadband services over next generation cellular networks. Email: xgeorge@aueb.gr.

Konstantinos Katsaros is a member of the Mobile Multimedia Laboratory at the Athens University of Economics and Business, Greece. He received his Diploma in Informatics (2003), and his M.Sc. (2005) and Ph.D. (2010) degrees in Computer Science from the Athens University of Economics

and Business, Greece. His research interests include future network architectures, content distribution and mobility support, as well as multicast and broadcast service provision over next generation cellular networks. He has also worked in the areas of mobile grid computing and cognitive radio. Email: ntinos@aueb.gr.

Vasilis Tsakanikas is an ex-member of the Mobile Multimedia Laboratory at the Athens University of Economics and Business, Greece. He received his Diploma (2005) in Electrical and Computer Engineering from the National Technical University of Athens, Greece. In his thesis, he performed an analysis of the established link between mobile terminals and users heads, utilizing Finite Difference Time Domain methods. He received his M.Sc. (2007) in Computer Science from the Athens University of Economics and Business, Greece. In his graduate thesis he performed an analysis of the MBMS power requirements. He is currently a technical manager for the Foundation for Research & Technology Hellas / Biomedical Research Institute. Email: vasilistsakanikas@yahoo.gr.