A Multiple Content Variant Extension of the Multimedia Broadcast/Multicast Service

George Xylomenos and Konstantinos Katsaros
{xgeorge,ntinos}@aueb.gr
Mobile Multimedia Laboratory
Department of Informatics
Athens University of Economics and Business
Patision 76, Athens 104 34, Greece

***Abstract*— This paper describes an extension of the Multimedia Broadcast/Multicast Service supporting the distribution of multiple variants of the same content to heterogeneous receivers. We first outline the standard MBMS model, along with its state management and signaling procedures, and then proceed to describe our extended MBMS model, detailing the modifications that it imposes on MBMS state management and signaling procedures. Finally, we compare via analysis and simulation our extension against an approach for achieving the same goal by using standard MBMS mechanisms in terms of the signaling overhead incurred, showing that our proposal is clearly superior.**

I. Introduction

The high bandwidth of 3G systems is making feasible the provision of multimedia services over cellular links. While the amount of bandwidth consumed by services such as video distribution makes them too expensive for most users, costs can be dramatically reduced when many users desire to receive the same service, by transmitting the corresponding data only once per cell. This can be achieved either by *broadcast*, where all users receive the service, or by *multicast*, where only selected users receive the service. We focus below on multicasting, as it is more appropriate for targeting users that have subscribed to, and possibly paid for, a service.

Similar reasoning in the past led to the introduction of IP multicasting into the Internet. The *Universal Mobile Telecommunications Systems* (UMTS), specified by the *3rd Generation Partnership Project* (3GPP), added support for IP multicasting in the Release 99 specifications, but this simply consisted of sending each multicast packet over a separate unicast tunnel to each receiver. In contrast, the *Multimedia Broadcast/Multicast Service* (MBMS), introduced in the Release 6 specifications, allows resource sharing throughout the UMTS network, including over the air [1]. MBMS is suitable for services as diverse as media streaming and file downloads [3].

MBMS multicasting is compatible to some extent with IP multicasting, upon which it is based. Both IP and MBMS multicasting deliver the same content to all receivers, so as to transmit the data only once over each link of the multicast distribution tree. When the receivers are heterogeneous however, for example, terminals with different screen resolutions or users with different budgets, it is difficult to select a variant of the content that will satisfy everyone. If a high quality, expensive, variant is distributed, some terminals will not be able to receive it and some users will not want to pay for it; if a low quality, cheaper, variant is distributed, users that would pay more for better service will not be satisfied.

In our research for the IST B-Bone project we designed an extended MBMS model, supporting the distribution of multiple variants of the same content to different receivers; we refer to this model as *Multiple Content Variant* (MCV) MBMS. The desired variant can be dynamically and independently selected by each receiver based on terminal capabilities and/or user preferences. Our approach is combined with layered coding in order to transmit the minimum possible amount of data over each link between the sender and the receivers.

In Section II we introduce the standard MBMS model, along with its state management and signaling procedures. In Section III we describe our extended MBMS model along with its modifications to the standard state management and signaling procedures. In Section IV we compare our model in terms of the signaling overhead incurred against an alternative approach for achieving the same goal by using standard MBMS mechanisms. We present our conclusions in Section V.

II. Standard MBMS

Figure 1 depicts an example UMTS network. The *Broadcast/Multicast Service Centre* (BM-SC), is a new entity that controls the services provided by MBMS, while the *Gateway GPRS Support Node* (GGSN), the *Serving GPRS Support Node* (SGSN), the *Radio Network Controller* (RNC), the *Node-B* and the *User Equipment* (UE) are existing entities that must be modified to support MBMS; we commonly refer to the RNCs and Node-Bs as the *Radio Access Network* (RAN).

In the original IP multicasting model, each multicast group is identified by a class D IP address. Any host can *join* the multicast group to start receiving packets sent to it and later *leave* the group to stop receiving such packets. Any host can send packets to the group, even non group members, using the IP address of the group. This *open group* model is not very attractive for commercial content providers, as they cannot limit reception of their content to paid subscribers only.

The MBMS multicasting model departs from IP multicasting in that each multicast group is identified by a class D IP address *and* an *Access Point Name* (APN), which effectively

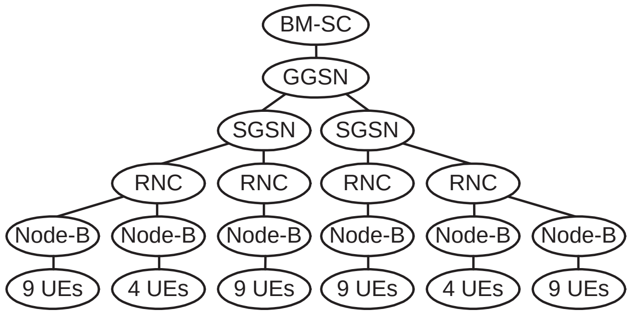

Fig. 1. Example UMTS network.

identifies the GGSN serving a specific UMTS network [2]. A more significant difference is that a *closed group* model is used. On the receiver side, a UE must first subscribe to a group in order to be allowed to join it. On the sender side, only the GGSN identified by the APN may send to the group; the content is first processed by the BM-SC and then delivered via the GGSN to the group. This model is quite suitable for the provision of commercial services over MBMS [1].

The IP multicasting implementation is split into *local* mechanisms, which track group membership and deliver multicast packets within a local network, and *global* mechanisms, which route multicast packets between local networks. The only local mechanism defined for IP multicasting is the *Internet Group Management Protocol* (IGMP) [4], a query/response protocol suitable for Ethernets. In contrast, many global mechanisms have been proposed for IP multicast routing. In MBMS multicasting, the local mechanisms cover the entire network served by a GGSN; the GGSN acts as the interface between the UMTS network, where MBMS multicasting is used, and the Internet, where IP multicasting is used. The GGSN may however receive data to forward to a group via either IP multicasting or IP unicasting. Inside the UMTS network, each UE that desires to join or leave an MBMS multicasting group sends IGMP messages to the GGSN, which are followed by MBMS signaling that creates a multicast distribution tree between the BM-SC and the UEs.

The 3GPP has defined a rich *Quality of Service* (QoS) model for UMTS networks. Each data stream may belong to the Conversational, Streaming, Interactive or Background class, and various parameters may be specified for a stream, depending on its service class. MBMS services may only belong to the Streaming or Background classes, since they are strictly unidirectional, from the BM-SC to the UE. An MBMS service must use the same QoS parameters for the entire multicast distribution tree, as defined by the BM-SC.

In order to support MBMS multicasting, each (internal) node in a UMTS network must maintain two types of state. First, packet forwarding state is required so that the node may determine which of its children should receive a multicast packet; this state is kept on a per group basis. Second, user accounting state is required so that the network may charge the receivers for participating in the group; this state is kept on a per group and a per UE basis. Each node maintains an *MBMS Bearer Context* (MBC) for each multicast group and an *MBMS UE Context* (MUEC) for each UE that is currently a member of the group [2], as shown in Figure 2 (fields in gray are used by our extended model).

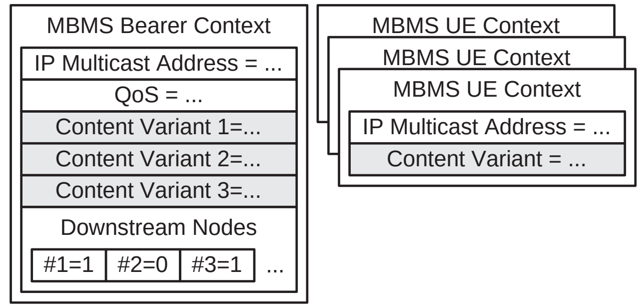

Fig. 2. MBMS contexts.

The MBC contains information for the entire group, such as its IP multicast address and its QoS parameters. The MBC includes a table indicating which downstream nodes (its children) 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 each child marked 1.

The MUEC contains information for a UE served by the node that is currently a member of the group. It is created (destroyed) when the UE joins (leaves) the group. Each MUEC is linked to a MBC via an IP multicast address. When the node forwards data to a multicast group, it uses the MUECs linked to the MBC to charge UEs. The MUEC does *not* indicate which child serves the UE, allowing the UE to move between children of the same node without notifying it.

When an MBMS multicast service is to be offered, the data describing it, such as IP multicast address and QoS parameters, are entered into a new MBC at the BM-SC. Additional MBCs and MUECs are dynamically created and destroyed at each node based on UE initiated signaling. When a UE desires to join (leave) a group, it sends an IGMP join (leave) message to its GGSN stating the corresponding IP multicast address [2]; this join/leave mode of IGMP is quite different from the query/response mode used with IP multicasting [6].

An IGMP join message triggers the multicast activation procedure shown in Figure 3. The GGSN asks the BM-SC if the UE has subscribed to the group and the BM-SC returns the APN of the GGSN that acts as the source. The GGSN then asks the SGSN if it can handle the MBMS multicasting group. The SGSN responds to the GGSN and notifies the UE that it can proceed. At this point (first dashed line) the UE knows the APN of the source, which may map to a different GGSN than the one initially contacted. The UE then requests the SGSN to start sending it multicast data. The SGSN creates the MUEC and notifies the GGSN corresponding to the APN. The GGSN notifies the BM-SC, the BM-SC creates the MUEC and responds to the GGSN. The GGSN creates the MUEC and responds to the SGSN, which responds to the UE.

When the first MUEC for a group is created at the GGSN or SGSN, the node initiates the registration procedure with its parent, indicating that it wants to start receiving data addressed

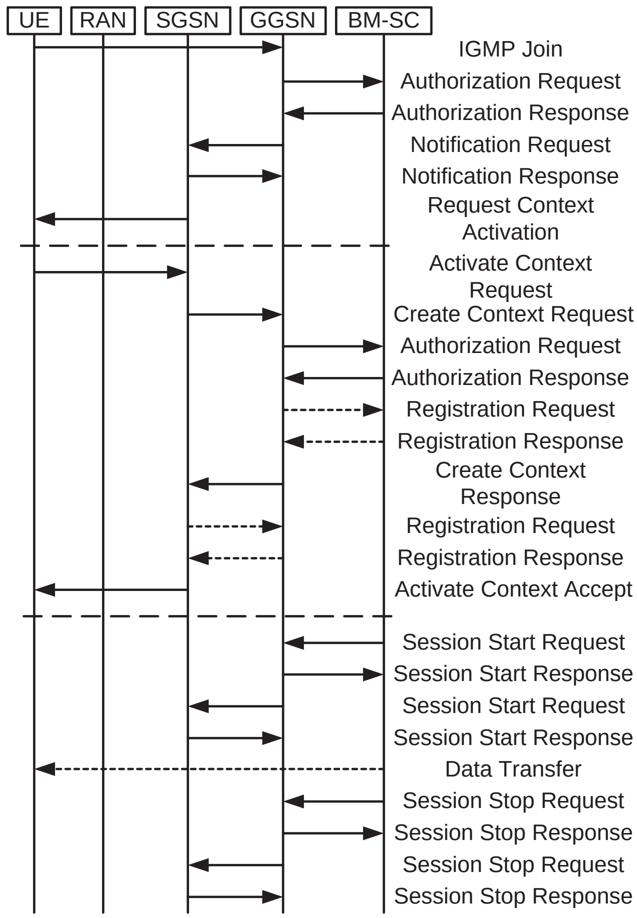

Fig. 3. MBMS multicast signaling.

to the group. The parent marks the corresponding entry in its MBC with 1 to start forwarding data to that child. Using the information provided in the response, the child creates the MBC for the group. These messages are shown with dotted lines, since they are not always exchanged after a join.

When the service is about to begin transmitting data (second dashed line), the BM-SC initiates the session start procedure towards each registered GGSN; similarly, these GGSNs initiate the session start procedure towards each registered SGSN. The SGSNs in turn instruct the RAN to establish appropriate radio bearers for data transmission, based on the QoS parameters of the service. After data transmission ends, the session stop procedure is initiated to release all reserved resources.

Similarly, an IGMP leave message triggers the multicast deactivation procedure (not shown) which reverses the activation procedure by removing the MUEC for that UE from each node. When the last MUEC for a group is destroyed at the GGSN or SGSN, the node initiates the deregistration procedure with its parent, indicating that it wants to stop receiving data addressed to the group. The parent marks the corresponding entry in its MBC with 0 to stop forwarding data to that child.

III. Multiple Content Variant MBMS

Our MCV MBMS model departs from standard MBMS by allowing a *single* service to offer different variants of the same content to different UEs, providing various tradeoffs between bandwidth and quality. This may increase the number of subscribers to a service by satisfying a wider range of heterogeneous receivers, such as terminals with different capabilities or users with different budgets. However, the QoS parameters for all variants of a service remain centrally determined by the BM-SC, as in the standard MBMS model, in order to prevent the degeneration of multicast groups to single receivers.

For MCV MBMS to be economical, the content variants must allow a node to generate all the variants requested by its children from the highest quality variant among them; this is the only variant that the node needs to receive from its parent. One method that may be used to generate such content variants is layered coding; another method is transcoding [5]. In the layered coding approach, the source encodes the lowest quality variant as the base layer and then encodes a series of successive enhancement layers. The next 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 network, and each node forwards to each child only the layers required to reconstruct the variant requested by that child, dropping the rest. We further discuss transcoding and layered coding elsewhere [7].

The variants must be chosen by the content provider to match common terminals and have sufficiently different costs. For example, for an audio service the variants could be CD, radio and telephone quality sound. The number of available variants must be small, to prevent the degeneration of multicast groups to single receivers. We have decided to support three variants, numbered 1 (low quality, LQ), 2 (medium quality, MQ) and 3 (high quality, HQ). The content variants available for each service are announced as in standard MBMS. A UE specifies the variant that it wishes to receive in its IGMP join message by including a variant number. The UE may later modify this request by sending a new IGMP join. For example, the user may request a higher quality audio variant to better hear a passage, or the terminal may request a lower quality audio variant when the bandwidth at its location is limited.

To support our MCV MBMS model, each (internal) node in a UMTS network must maintain additional information in the MBC and MUEC, as shown in Figure 2 (gray fields). The MUEC must be extended with the number of the requested variant, allowing the node to charge the UE accordingly. The downstream nodes table in the MBC must be extended with the number of the variant to forward to each child. For example, in Figure 2, a child marked 3 would receive variant 3.

In the MCV MBMS model, each node must inform its parent about the variant that it needs to receive, thus allowing the parent to maintain the entry for that child in the MBC. The node must thus determine the highest quality variant requested by *any* of its own children; from this variant it can produce any lower quality variants requested by its other children. We determine this information by counting the number of MUECs per variant and storing these counters in the MBC for the group, as shown in Figure 2. Each node requests from its parent the highest quality variant with a nonzero counter.

In order to implement this scheme, when a UE sends an IGMP join (leave) message, triggering the activation (deactivation) procedure, at the point where a node would create (destroy) a MUEC in standard MBMS, in MCV MBMS the node must instead do one of the following:

- If a MUEC was created (destroyed) at the node, the counter for the corresponding variant is incremented (decremented) by one.
- If a MUEC was modified at the node, the counter for its previous variant is decremented by one and the counter for its current variant is incremented by one.

The only modification needed in the MUEC management messages, shown with solid lines in Figure 3, is the addition of a variant number in the requests. Note that new IGMP join messages can be received from a UE that is already a member of a group and wishes to change its variant.

Furthermore, at the point where a node would create (destroy) a MBC in standard MBMS, in MCV MBMS the node must instead do one of the following:

- If the first (last) MUEC for a group was created (destroyed), the MBC is created (destroyed) and the parent is informed to start (stop) forwarding the proper variant.
- If the counter for a higher quality variant than the one currently requested became nonzero, the parent is informed to send the corresponding variant.
- If the counter for the currently requested variant became zero, the next nonzero counter is located and the parent is informed to send the corresponding variant.

The only modification needed in the MBC management messages, shown with dotted lines in Figure 3, is the addition of a variant number in the requests. Note that new registration messages can be received from a child that is already receiving a group and wishes to change its variant.

IV. Performance Evaluation

We will now assess the performance of MCV MBMS in terms of its signaling overhead, against the performance of a single MBMS service and a multiple MBMS service approach. The single MBMS service approach offers a single content variant only, therefore it will either transmit a content variant that is not requested by all UEs, or a content variant that does not satisfy all UEs. While it cannot be compared with MCV MBMS in terms of user plane traffic, it provides a performance baseline for control plane traffic.

The multiple MBMS service approach combines layered coding with a separate MBMS group per layer, with each UE joining the groups corresponding to the layers that it needs; the functionality provided is equivalent to MCV MBMS. This approach produces the same user plane traffic as MCV MBMS, as only the packets that are absolutely needed by each node are forwarded to it, but differs from it in control plane traffic: MCV MBMS slightly inflates the request packets, due to the inclusion of the content variant desired, but it does not require any UEs to join multiple MBMS groups.

As part of the IST B-Bone project we built an MBMS simulator based on the 3GPP Release 6 specifications. The simulator uses the Opnet Modeler 11.0 platform and, among other extensions to standard MBMS, it fully supports the MCV MBMS model described above. We used this simulator to compare MCV MBMS with the single and multiple MBMS approaches in terms of the signaling overhead incurred by the procedures shown in Figure 3.

We simulated the topology shown in Figure 1, consisting of a single GGSN, two SGSNs, four RNCs and six Node-Bs, where two of the Node-Bs control cells with 4 UEs and the other four Node-Bs control cells with 9 UEs. In each experiment some UEs are randomly chosen to join a multicast group. Each UE randomly selects a content variant with probability $p_L = 70\%$ for the LQ variant, $p_M = 20\%$ for the MQ variant and $p_H = 10\%$ for the HQ variant. The order in which UEs join is also random; depending on the actual order in an experiment, each node may receive 1 to 3 registration messages from each of its children. The number of UEs participating in the service was varied from 1 to 40, and each experiment was repeated 30 times with different seeds.

We have also performed an analytical evaluation of the alternative approaches discussed above under very general conditions [8]. By applying this analysis to our environment we find that the predicted signaling overhead of the multiple MBMS service approach in packets is

$$A_m(1 + p_M + 2p_H)N_u + 3S_n(R_m + S_m + SS_m)$$

where A_m, R_m, S_m and SS_m are the number of packets required for the activation, registration, session start and session stop procedures, respectively, N_u is the number of participating UEs and S_n is the number of SGSN nodes. By substituting values for these parameters from the standards, this expression reduces to $36 + 11.2N_u$. For a single MBMS service the predicted signaling overhead in packets is:

$$A_mN_u + S_n(R_m + S_m + SS_m)$$

or, after substituting values from the standards, $12 + 8N_u$; this is the minimum amount of signaling possible. Finally, the analysis for MCV MBMS indicates that the average predicted signaling overhead in terms of packets is

$$A_mN_u + S_n(2R_m + S_m + SS_m)$$

which reduces to $16 + 8N_u$. Figure 4 shows the simulation results and the analytical predictions for each approach; simulation results also indicate the 99% confidence intervals.

In the same manner, we find that the predicted signaling overhead of the multiple MBMS service approach in bits is

$$8[A_b(1 + p_M + 2p_H)N_u + 3S_n(R_b + S_b + SS_b)]$$

where A_b, R_b, S_b and SS_b are the number of bytes required for the activation, registration, session start and session stop procedures, respectively. By substituting values for these parameters from the standards, this expression reduces to $11184 + 3238.4N_u$. For the single MBMS service case the predicted signaling overhead in bits is

$$8[A_bN_u + S_n(R_b + S_b + SS_b)]$$

which reduces to $3728 + 2336N_u$. For MCV MBMS we assumed that 1 byte is added to each request packet to indicate the desired variant; the predicted signaling overhead in bits is

$$8[A_bN_u + A_mN_u/2 + S_n(2R_b + R_m + S_b + 16S_m + S_Sb)]$$

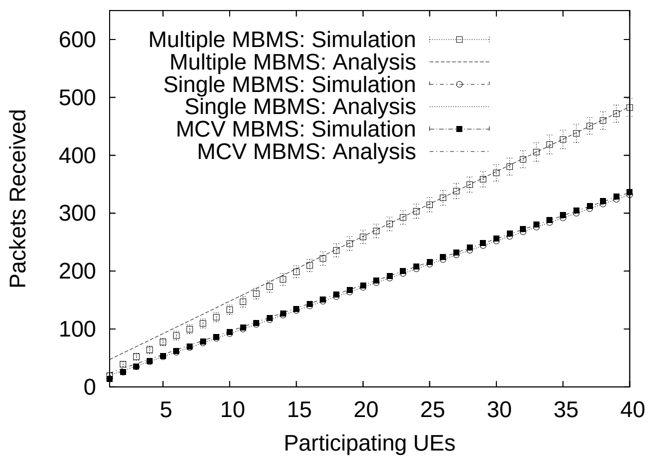

Fig. 4. Signaling overhead (packets).

which reduces to $5216 + 2368N_u$. Figure 5 shows the simulation results and the analytical predictions for each approach.

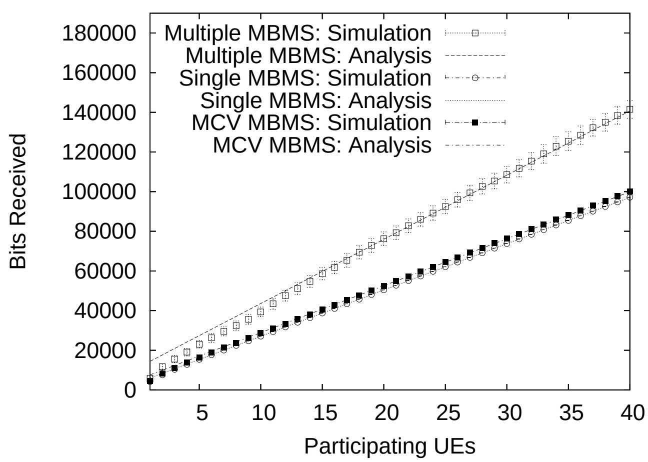

Fig. 5. Signaling overhead (bits).

Both figures show similar behavior for packets and bits and excellent agreement between the analytical predictions and the simulation results; the slight overestimation of the analysis for the multiple MBMS service approach with a few UEs is due to the simplifying assumption that at least one UE per SGSN has requested the HQ variant, which is obviously unlikely without a sufficient number of UEs for a probability of $p_H = 10\%$.

Comparing the MCV MBMS approach against the multiple MBMS service approach it is clear that there is a wide gap in the overhead in favor of our scheme, which grows with the number of participating UEs. For 40 UEs, the average number of control packets measured is 336 and 482, respectively, a difference of 43% in favor of MCV MBMS; the average number of control bits measured is 100,065 and 141,549, respectively, a difference of 41% in favor of MCV MBMS. The slightly smaller difference in bits is due to the inflated request messages with MCV MBMS. Therefore MCV MBMS provides a dramatic improvement over the functionally equivalent multiple MBMS service approach.

It is also very interesting to observe that MCV MBMS is only slightly more costly than the single MBMS service approach which incurs the minimum possible overhead. Indeed, the average number of control packets measured for 40 UEs with a single MBMS service is 332, while the figure for control bits is 97,168, therefore the extra cost of MCV MBMS is only 1.2% in terms of packets and 2.9% in terms of bits, again due to the slight inflation of the request messages by MCV MBMS. Therefore the benefits of MCV MBMS have a negligible cost compared to the single MBMS service approach.

V. Conclusion

We presented an MBMS extension supporting the distribution of multiple variants of the same content to heterogeneous receivers. We first described the standard MBMS model and its state management and signaling procedures, and then explained how our extended MBMS model can be derived from it. We compared our model in terms of the signaling overhead incurred against a multiple MBMS service approach, showing that MCV MBMS is clearly superior, as well as against a single MBMS service approach, showing that MCV MBMS provides enhanced functionality with negligible costs.

Acknowledgements

The work reported in this paper was supported by the IST B-Bone project under contract IST-2003-507607. The simulation platform was provided by the OPNET University Program.

References

[1] 3rd Generation Partnership Project (3GPP). Multimedia Broadcast / Multicast Service (MBMS); Stage 1. TS 22.146, V6.6.0, September 2004.
[2] 3rd Generation Partnership Project (3GPP). Multimedia Broadcast / Multicast Service (MBMS) user services; Architecture and functional description. TS 23.246, V6.4.0, September 2004.
[3] 3rd Generation Partnership Project (3GPP). Multimedia Broadcast / Multicast Service (MBMS) user services; Stage 1. TS 22.246, V6.2.0, September 2004.
[4] W. Fenner. Internet Group Management Protocol, Version 2. RFC 2236, November 1997.
[5] B. Li and J. Liu. Multirate Video Multicast over the Internet: An Overview. *IEEE Network*, 17(5):55–63, January/February 2003.
[6] G. Xylomenos. Group Management for the Multimedia Broadcast / Multicast Service. In *Proc. of the IST Mobile Summit*, June 2005.
[7] G. Xylomenos and K. Katsaros. Supporting Multiple Content Variants in the Multimedia Broadcast / Multicast Service. In *Proc. of the World Telecommunications Congress*, May 2006.
[8] G. Xylomenos, K. Katsaros, and G.C. Polyzos. Analysis of a Multiple Content Variant extension of the Multimedia Broadcast / Multicast Service. In *Proc. of the European Wireless Conference*, April 2006.