Reducing the Transmission Power Requirements of the Multimedia Broadcast/Multicast Service

George Xylomenos, Vasilis Tsakanikas and George C. Polyzos
Mobile Multimedia Laboratory
Department of Informatics
Athens University of Economics and Business
Patision 76, Athens 104 34, Greece
Email: xgeorge@aueb.gr, btsakan05@aueb.gr and polyzos@aueb.gr

***Abstract*— The Multimedia Broadcast/Multicast Service (MBMS) was designed to support the economical distribution of multimedia content to large numbers of receivers in 3rd generation cellular networks. In this paper we present and evaluate an MBMS extension that reduces the transmission power requirements while increasing the number of potential users of such services, by supporting the distribution of multiple variants of the same content to heterogeneous receivers. We first describe the standard MBMS model, along with its state management and signaling procedures, as well as our extended MBMS model, in terms of the modifications that it imposes on the standard. We then present an analytical and simulation evaluation of the transmission power requirements of our approach against alternatives based on standard MBMS, showing that our approach maximizes the number of potential users, without excessive transmission power requirements.**

I. Introduction

The increased bandwidth available in 3rd generation cellular networks makes them an attractive platform for multimedia services, such as video distribution. While the resource requirements of such services make them expensive for most individual users, these costs can be dramatically reduced by sharing them among many users receiving the same service. The *Universal Mobile Telecommunications System* (UMTS), specified by the *3rd Generation Partnership Project* (3GPP), has introduced to this end the *Multimedia Broadcast/Multicast Service* (MBMS) [1], which enables resource sharing throughout the network, including over the air. In this paper we focus on MBMS multicasting, which is more suitable for users that have subscribed to, and possibly paid for, a service.

The MBMS multicasting mode, similarly to IP multicasting, delivers the exact same content to all receivers. However, when these receivers are heterogeneous, for example, terminals with different screens or users with different budgets, a content provider faces a dilemma when choosing an appropriate variant of its content for transmission: a low quality variant will not satisfy users prepared to pay more for better service, while a high quality variant will not be received by users with simple terminals or limited budgets. In both cases, potential users, and the corresponding revenues, are lost. This problem is exacerbated by the fact that MBMS multicasting only provides transmission power savings compared to unicasting when a sufficient number of receivers exists in a cell. Therefore, increasing the number of potential users is critical for MBMS.

In order to maximize the number of potential users, without consuming an excessive amount of transmission power, as part of our research in the IST B-Bone project we have designed an MBMS extension supporting the distribution of multiple variants of the same content to different receivers, *Multiple Content Variant* (MCV) MBMS [2]. In this paper we compare our approach against some standards based alternatives in terms of their transmission power requirements. In Section II we describe the standard MBMS model, while in Section III we describe our extended MCV MBMS model. In Section IV we present an analytical comparison of MCV MBMS against two standards based alternatives, and in Section V we present a corresponding simulation based comparison.

II. The standard MBMS model

An example UMTS network supporting MBMS is shown in Figure 1. A new functional entity, the *Broadcast/Multicast Service Centre* (BM-SC), is added to control the provision of MBMS services. 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 the existing UMTS network elements that need to be modified so as to handle the establishment of multicast distribution trees and the transmission of MBMS data over these trees.

Even though MBMS is based on IP multicasting, it departs from it in many ways. While in IP multicasting groups are identified by a class D IP address, in MBMS a group is identified both by a class D IP address and by an *Access Point Name* (APN), which resolves to the GGSN serving a UMTS network. Therefore, MBMS services are defined with respect to a specific UMTS network and their scope is limited within that network. Furthermore, while in IP multicasting anyone can send to and receive from a group, in MBMS a UE must first *subscribe* to a group in order to be later allowed to *join* it so as to *receive* data, and only the GGSN identified by the APN may *transmit* data to a group. These are clearly ideal properties for commercial, subscription based, services [1].

Each network node supporting MBMS must maintain two types of state. First, packet forwarding state is required, allowing the node to determine which of its children should receive a packet. Second, user state is required, allowing the network to charge the participating receivers. Each such

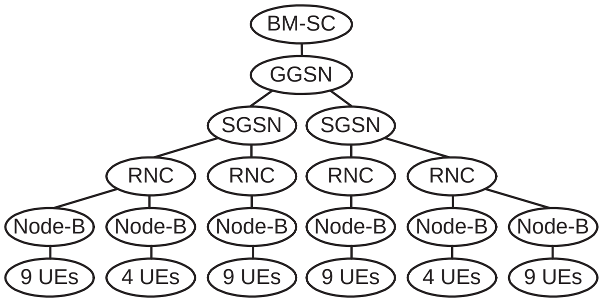

Fig. 1. An example UMTS network.

node therefore maintains an *MBMS Bearer Context* (MBC) for each multicast group present in its area and an *MBMS UE Context* (MUEC) for each UE served by the node that is currently a member of such a group [3], as shown in Figure 2 (fields in gray will be discussed in Section III). The MBC contains information pertaining to the entire group, including its forwarding state, which consists of a table indicating the downstream nodes that 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). The MUEC contains information pertaining to a specific UE; it is linked via its IP address to an MBC. When forwarding data to a group, the node uses the MUECs to charge the UEs.

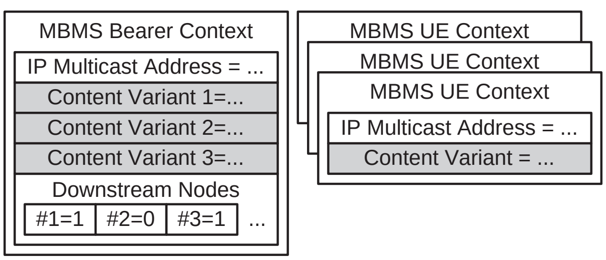

Fig. 2. The MBMS bearer and user contexts.

When an MBMS service is to be offered, its attributes are administratively entered into an MBC at the BM-SC. Additional MBCs and MUECs are dynamically created at nodes belonging to the multicast distribution tree, based on UE initiated signaling. Each UE desiring to join an MBMS group sends an *Internet Group Management Protocol* (IGMP) join message to the GGSN. This message triggers the *multicast activation* procedure, in which, after the BM-SC verifies that the UE is subscribed to the group and returns to it the APN of the group, a MUEC is created at the BM-SC, the GGSN and the SGSN serving the UE, thus establishing the user state. To establish the forwarding state, when the first MUEC for a group is created at the GGSN or SGSN, that node sends a *registration* to its parent, and the parent marks the proper entry in the forwarding table of the MBC with 1 so as to later forward packets to that child. Using the information provided in the registration response, the child creates its own MBC.

The multicast activation and registration procedures lead to the establishment of a multicast distribution tree from the BM-SC towards all UEs participating in an MBMS service, but they do not reserve any transmission resources: a separate *session start* procedure is used to establish radio bearers when the service is ready to start transmitting data. At the end of data transmission, a *session stop* procedure is used to release the radio bearers. A *multicast deactivation* procedure can be triggered by a UE desiring to leave a group by sending an IGMP leave message to the GGSN, thus releasing the user state from all nodes between the BM-SC and the UE. Finally, when the last MUEC for a group is destroyed at a GGSN or SGSN, a *deregistration* message is sent by that node to its parent, which marks the corresponding entry in the forwarding table of the MBC with 0.

When the session start procedure indicates to an RNC that it should establish radio bearers for a multicast service, the RNC must first decide what type of channel it should employ in each cell, and then instruct the Node-B controlling the cell accordingly. Two options exist: either establish separate *Point to Point* (PtP) links towards each participating UE using the *Dedicated Transport Channel* (DCH), or establish a common *Point to Multipoint* (PtM) link towards all participating UEs using the *Forward Access Channel* (FACH) [4]. With the FACH, a single transmission reaches all UEs in the cell, regardless of their number and position. This means however that the FACH must always transmit at a high enough power level so as to reach UEs even at the edge of a cell. With the DCH, a separate transmission on a separate DCH is required for each participating UE. Each DCH however employs power control, that is, each transmission is performed with the minimum power required to reach the target UE. As a result, for a few UEs it is normally more economical to employ multiple DCH links rather than a single FACH link.

In order for the RNC to select the appropriate type of radio bearer for each cell, it must estimate how many UEs participate in the group in that cell. This is achieved via a procedure known as *UE Counting* [4], whereby the RNC asks the UEs participating in a multicast group to establish a signaling connection with the network 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 the number of UEs in the cell is lower than an administratively defined *threshold* T, individual DCH links are established, while if it is larger than or equal to T, a common FACH link is established. Ideally, T should be selected so that the average transmission power required to send T copies of a packet over DCH links is slightly higher than the transmission power required to send that packet once over a FACH link.

III. The Extended MBMS Model

Our *Multiple Content Variant* (MCV) model extends the standard MBMS model by allowing a single MBMS service to offer different variants of the same content to heterogeneous receivers. This is achieved by using *layered coding* [5] to create the variants, and by extending the standard MBMS signaling and state management procedures so as to distribute these variants in an economical manner. In layered coding, the source encodes the content as a *base layer* and a series of successive *enhancement layers*. In MCV MBMS each content variant consists of the base layer and a set of successive enhancement layers. Since the number of available variants must be kept small to prevent the degeneration of multicast groups to single receivers, in our simulator implementation we support up to three variants, *low quality* (LQ), *medium quality* (MQ) and *high quality* (HQ). A UE specifies the desired variant by including it in its IGMP join message; it may later modify this request with a new IGMP join. For example, the user may request higher quality audio to better hear a passage, or lower quality audio when the bandwidth at its cell is limited.

The traditional way to combine multicasting with layered coding is to transmit each layer via a separate group, with each receiver joining the groups corresponding to the layers required to reconstruct its desired content variant [6]. The use of multiple groups leads to an increase of the signaling load placed on the network, thus limiting MBMS scalability. In our approach we exploit the tree topology of UMTS networks, which implies that all multicast distribution trees will overlap, in order to merge all groups into one. In order to achieve this, the source injects all layers into the multicast distribution tree, but 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 2, if child #3 was marked 2, that node would need 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 the base layer plus the first enhancement layer to child #3.

In our extended model each node must maintain additional user and forwarding state in each MUEC and MBC, shown in Figure 2 with a gray background. The MUEC includes a number indicating the requested variant, so as to allow the node to charge the UE accordingly, while the MBC stores a count of the number of MUECs for each variant, so as to allow the node to determine the content variant to request from its parent. Finally, the forwarding table in the MBC for the group is also modified to hold the number of the variant to forward to each child, that is, 0 (none) to 3 (HQ).

The state management procedures of the standard MBMS model must also be modified as follows [2]. When a UE sends an IGMP join, at the point where in standard MBMS a node would create or destroy a MUEC, in MCV MBMS the node must instead do one of the following: a) if a MUEC was just created (destroyed), the counter for its variant is incremented (decremented), or, b) if a MUEC just modified its variant, the counter for its previous variant is decremented and the counter for its current variant is incremented. This ensures that the counters in the MBC are kept up to date. In addition, at the point where in standard MBMS a node would create or destroy an MBC, in our model it must instead do one of the following: a) if the first (last) MUEC was created (destroyed), the MBC is created (destroyed) and the parent is informed to start (stop) forwarding data, b) if the counter for a higher quality variant than the current one became nonzero the parent is informed, c) if the counter for the current variant became zero, the next nonzero counter is found and the parent is informed. This ensures that the parent's forwarding state is kept up to date.

When the session start procedure indicates to an RNC that it should establish radio bearers for a multicast service, the RNC in our model must make a separate decision about the type of channel it should establish for each layer. Therefore, the RNC will 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, depending on whether the number of users requesting it is less than or greater than the threshold T. In our simulator implementation, all the UEs (LQ, MQ and HQ) need to receive the base layer, the MQ and HQ UEs also need to receive the first enhancement layer and the HQ UEs also need to receive the second enhancement layer. As a result, the lower layers may be transmitted over a common FACH link, while the higher layers may be transmitted over separate DCH links.

IV. Performance Evaluation: Analysis

In this section we first review our past work on the evaluation of our MCV MBMS model against some standards based alternative approaches and then present an analytical evaluation of the transmission power requirements of each option. The first alternative, *Base MBMS*, uses a single group to distribute only the LQ variant to all receivers; it incurs the lowest overhead, but does not satisfy MQ and HQ users. The second alternative, *Single MBMS*, also uses a single group but distributes the HQ variant (all layers) everywhere: it satisfies all users, but incurs the highest overhead. The third alternative, *Multiple MBMS*, uses an independent group for each layer: it satisfies all users without incurring high user plane overhead, but increases control plane overhead, since UEs need to join multiple groups. Note that Base MBMS and Single MBMS incur the same control plane overhead, while MCV MBMS and Multiple MBMS incur the same user plane overhead.

Regarding control plane overhead, in terms of packets and bits received, both analysis [2] and simulation [7] indicate that MCV MBMS has a negligible cost compared to Base MBMS and Single MBMS, unlike Multiple MBMS which incurs considerable additional overhead. While the amount of control plane traffic is small compared to user plane traffic, it is concentrated at the highest levels of the hierarchy, thus limiting MBMS scalability. Regarding user plane overhead, in terms of packets and bits received, both analysis and simulation [8] indicate that MCV MBMS provides dramatic savings in the radio access network over Single MBMS; its user plane overhead is actually closer to that of Base MBMS.

The limitation of our past work is that it focuses on the amount of packets and bits received, which do not directly reflect the transmission power consumed in each cell: when PtM channels are used, each transmission costs more but reaches many UEs; when PtP channels are used, transmissions cost less but only reach a single UE. In this paper we therefore focus on comparing the transmission power requirements of MCV MBMS against Base MBMS and Single MBMS. Note that Multiple MBMS is nearly identical to MCV MBMS in this respect, since they only differ in the control plane, which has a negligible effect on the total power consumption.

To estimate the transmission power requirements of each option, we assume that N_u users are interested in an MBMS service offered in three variants, LQ, MQ and HQ, comprising three layers, also called LQ, MQ and HQ. The probability that a user will request each variant is p_L, p_M or p_H, respectively, while the probability that each of the B packets generated by the source are part of each layer is q_L, q_M or q_H, respectively. If C cells exist, the expected number of users N_i in a cell i where U_i potential users exist is $N_i = (U_i / \sum_{i=1}^{C} U_i) N_u$.

In MCV MBMS, the transmission power required for the Bq_L packets of the LQ layer depends on whether the $(p_L + p_M + p_H)N_i = N_i$ users requesting the LQ layer are less or more than the threshold T. If we denote the power required to transmit a packet via the DCH by P_D and the power required to transmit that packet via the FACH by P_F, then the total expected transmission power in cell i for the LQ layer is:

$$P(i)_L = \begin{cases} Bq_L N_i P_D & : \quad N_i < T \\ Bq_L P_F & : \quad N_i \geq T \end{cases}$$

Similarly, for the MQ and HQ layers, we find that:

$$P(i)_M = \begin{cases} Bq_M (p_M + p_H) N_i P_D & : \quad (p_M + p_H) N_i < T \\ Bq_M P_F & : \quad (p_M + p_H) N_i \geq T \end{cases}$$

$$P(i)_H = \begin{cases} Bq_H p_H N_i P_D & : \quad p_H N_i < T \\ Bq_H P_F & : \quad p_H N_i \geq T \end{cases}$$

Therefore, the total transmission power consumed by MCV MBMS is $\sum_{i=1}^{C} [P(i)_L + P(i)_M + P(i)_H]$. For Base MBMS, only the LQ layer is transmitted, therefore the total transmission power consumed is simply $\sum_{i=1}^{C} P(i)_L$. For Single MBMS, all layers are transmitted to all UEs, therefore the total expected transmission power in cell i is:

$$P(i)_S = \begin{cases} BN_i P_D & : \quad N_i < T \\ BP_F & : \quad N_i \geq T \end{cases}$$

Hence, the total transmission power consumed is $\sum_{i=1}^{C} P(i)_S$.

V. Performance Evaluation: Simulation

As part of the IST B-Bone project, we have implemented a detailed MBMS simulator based on the 3GPP Release 6 specifications, using the Opnet Modeler 11.0 platform, which, among other extensions, fully implements MCV MBMS. In order to evaluate our approach, we used the topology shown in Figure 1 with the parameter set discussed below. For comparison purposes, we applied the same topology and parameters to the analytical model presented in Section IV. In the simulated topology we have two types of cells: the four *dense* cells host nine UEs, while the two *sparse* cells host only four UEs. The UEs did not move or change their content variant preferences over time, to allow comparisons with the analytical model.

In the radio network, each Node-B used a single sector antenna with a maximum transmission power of 20 W, to cover 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; d is the distance between the UE and the Node-B antenna and f is the carrier frequency in MHz. This model is valid for Non Line-of-Sight cases and describes the worst case propagation. The shadow fading loss was modeled as a log-normal random variable with zero mean and variance 10 dB, a common assumption for outdoor users.

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 $2/3$ of the cell radius, again a common assumption for this environment. The transmission power of each DCH depended on both distance and time, due to the outer loop power control modeled by the simulator. The average DCH transmission power in our experiments was P_D = 0.08975 W, therefore we set the threshold $T = 5$, that is, the FACH was used to serve 5 or more UEs, implying that the sparse cells never used the FACH, something reasonable in a real network.

The MBMS service modeled was a stream of IP packets with a payload of 968 bytes, generated every 0.125 s, that is, a bit rate of roughly 62 Kbps, excluding the UDP/IP header overhead of 28 bytes per packet. At the radio link these packets were 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 segment was transmitted at the power level indicated above in a 20 ms interval. While the simulation results include *all* downlink packets, in practice the power spent for signaling packets was negligible.

In each experiment, N_u UEs were randomly chosen to join a multicast group; we varied N_u from 1 to 40 and repeated each experiment 30 times. Each UE randomly selected a content variant with probabilities $p_L = 0.7$, $p_M = 0.2$ and $p_H = 0.1$. The source generated 1000 packets, which were transmitted as $B = 6000$ segments, distributed to the three content variants with probabilities of $q_L = 0.5$, $q_M = 0.25$ and $q_H = 0.25$.

Figure 3 shows the total transmission *energy* consumed (in mWh), calculated by multiplying for each packet its transmission power by its transmission time and then summing up over all transmissions. For each option, we show both the analytical predictions (see Section IV) and the simulation results averaged over the 30 runs. The agreement between analysis and simulation is nearly perfect for up to 25 UEs. The gap that appears at this point is due to the assumption of the analytical model that all dense cells host the same (fractional) number of UEs, causing them to simultaneously switch from the DCH to the FACH, unlike in the simulator where each cell hosts an integer number of UEs. At 30 UEs all dense cells have switched to the FACH and the gap begins to close.

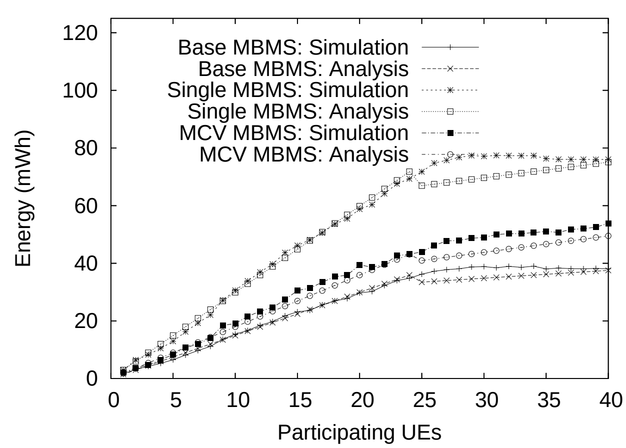

Fig. 3. Total Transmission Energy.

Averaged over the entire range of 1 to 40 UEs, the simulation results show that MCV MBMS consumes 29% more energy than Base MBMS, while Single MBMS consumes 99% more energy than Base MBMS. Therefore, while both Single MBMS and MCV MBMS satisfy all users, the extra energy required by MCV MBMS to achieve this is less than one third of that required by Single MBMS, hence the number of users is maximized without an excessive amount of energy.

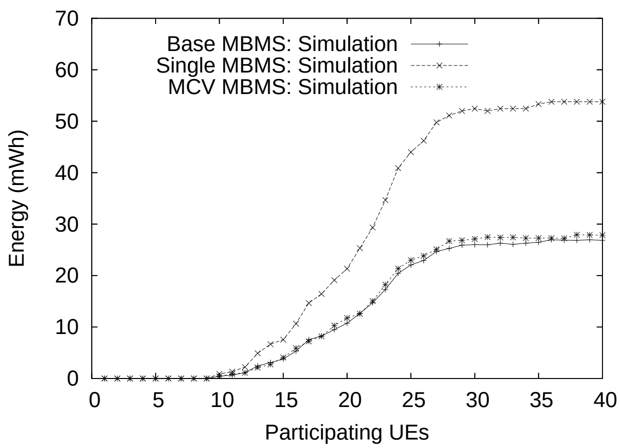

Fig. 4. FACH Transmission Energy.

To gain a better insight into these results, Figure 4 and Figure 5 show how much of the total transmission energy is due to the FACH and the DCH, respectively. From Figure 4 we see that at 10 UEs the dense cells start switching to the FACH. The energy consumption increases until 30 UEs and then remains constant: at this point all dense cells have switched to the FACH. Note that MCV MBMS uses the same energy over the FACH as Base MBMS, since MCV MBMS only uses the FACH for the LQ layer; the MQ and HQ users are too few to warrant use of the FACH. In contrast, in Single MBMS all layers are transmitted to all users, thus doubling the energy consumption in the FACH compared to Base MBMS.

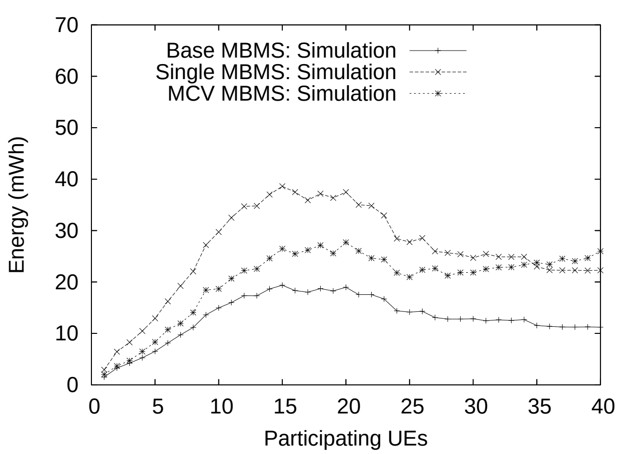

Fig. 5. DCH Transmission Energy.

Finally, from Figure 5 we see that the energy consumption of the DCH follows similar curves for Base MBMS and Single MBMS. The peak is reached at 15 UEs, where most cells still use the DCH. From this point on, DCH energy consumption starts dropping as the dense cells switch to the FACH; at 30 UEs they have all switched to the FACH and DCH energy consumption stabilizes. Single MBMS again spends twice as much energy as Base MBMS. MCV MBMS behaves similarly until 15 UEs, but after that point its DCH energy consumption remains at the same level. The reason for this is that while the LQ layer switches to the FACH in the dense cells, decreasing the DCH energy consumption, the MQ and HQ layers are still sent via the DCH, so, as the number of UEs increases, the DCH energy consumption due to these layers is also increased.

VI. Conclusion

We presented an extended MBMS model that transmits different variants of the same content to each UE in order to maximize the number of UEs participating in an MBMS service while minimizing the transmission power required to serve them. We explained how our extended MBMS model can be derived from the standard MBMS model, and evaluated our extensions against two alternatives based on standard MBMS via analysis and simulation, showing that our approach can satisfy all users without consuming excessive power.

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] 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.
[3] 3rd Generation Partnership Project (3GPP), "Multimedia Broadcast / Multicast Service (MBMS) user services; Architecture and functional description," TS 23.246, V6.4.0, September 2004.
[4] ——, "Introduction of the Multimedia Broadcast Multicast Service (MBMS) in the Radio Access Network (RAN); Stage 2," TS 25.346, V6.6.0, September 2004.
[5] B. Li and J. Liu, "Multirate video multicast over the Internet: An overview," *IEEE Network*, vol. 17, no. 5, pp. 55–63, 2003.
[6] N. Shacham, "Multipoint communication by hierarchically encoded data," in *Proc. of the IEEE INFOCOM*, vol. 3, May 1992, pp. 2107–2114.
[7] G. Xylomenos, "A Multiple Content Variant of the Multimedia Broadcast / Multicast Service," in *Proc. of the IST Mobile Summit*, June 2006.
[8] G. Xylomenos and K. Katsaros, "Supporting Multiple Content Variants in the Multimedia Broadcast / Multicast Service," in *Proc. of the World Telecommunications Congress*, May 2006.