On the efficiency of a dedicated LMA for multicast traffic distribution in PMIPv6 domains

Luis M. Contreras[1], Carlos J. Bernardos[1], and Ignacio Soto[2]

[1] Departamento de Ingeniería Telemática, Univ. Carlos III de Madrid, Spain
`contreras.uc3m@gmail.com, cjbc@it.uc3m.es`
[2] Departamento de Ingeniería de Sistemas Telemáticos, Univ. Politécnica de Madrid, Spain
`isoto@dit.upm.es`

Abstract. IP multicast allows the efficient support of group communication services by reducing the number of IP flows needed for such communication. Proxy Mobile IPv6 (PMIPv6) is a network-based mobility management solution, where the functionality to support the terminal movement resides in the network. Recently, a baseline solution has been adopted for multicast support in PMIPv6. Such base solution has inefficiencies in multicast routing because it may require multiple copies of a single stream to be received by the same access gateway. Nevertheless, an alternative solution to support multicast in PMIPv6 avoids this issue. This paper evaluates by simulation the scalability of both solutions under realistic conditions, and provides an analysis of the sensitivity of the two proposals against a number of parameters.

Keywords: Multicast, PMIPv6, Mobility.

1 Introduction

IP multicast allows the efficient support of group communication services (one-to-many or many-to-many) over IP networks. Applications like TV distribution take advantage of this extension of the IP protocol which facilitates the delivery of a single copy of a data stream to multiple listeners interested in receiving the same content simultaneously. The increasing generalization in the use of multicast has also triggered the need for supporting IP multicast in mobile environments.

Mobility management support in IP-based networks is a topic that has received considerable attention in recent years. A first approach to cover this issue was based on host-based solutions, mainly Mobile IPv6 (MIPv6) [1], where IP mobile terminals or Mobile Nodes (MNs) are aware of their IP mobility and have to perform operations in order to maintain their ongoing communication sessions. As a further step, a network-based solution has been standardized, called Proxy Mobile IPv6 (PMIPv6) [2], where the functionality to handle the movement of the MNs resides in the network.

Group communication is out of the scope of the PMIPv6 standard specification. Recently, a base solution has been adopted for multicast service delivery in PMIPv6 domains [3]. This baseline solution has been built on the capabilities of the existing multicast and mobility protocols, and considers the use of the PMIPv6 core entities to serve both unicast and multicast traffic, without any further improvement for a better adaptation to the multicast case. As consequence, the base solution has inefficiencies, specially the well known tunnel convergence problem, which implies that multiple copies of a single stream may be received by the same access gateway. To solve this issue, the architecture defined in [4] proposes a separate core entity being the unique responsible of serving multicast traffic to the access nodes, then guaranteeing that only one copy of the same stream is certainly received. This paper presents a comparative analysis of both solutions, providing some insights in the potential benefits of using dedicated infrastructure for multicast service.

2 Background on PMIPv6 and multicast

2.1 Proxy Mobile IPv6

PMIPv6 is based on MIPv6, reusing much of its concepts and packet formats. In PMIPv6, mobility support is provided by some specific network entities, namely Mobile Access Gateway (MAG) and Local Mobility Anchor (LMA).

The MAG takes care of the mobility signaling on behalf of the MNs attached to its links, tracking the MNs as they move, while the LMA stores all the routing information needed to reach the MNs in the PMIPv6 domain by associating each MN with the MAG that the MN is using. A tunnel between the LMA and the MAG allows the transfer of traffic from and to the MN. Using PMIPv6, the MN can move across a PMIPv6 domain changing its access link, while keeping its IP address.

The MN is registered by the MAG in the LMA by sending a Proxy Binding Update message (an extension of the MIPv6 Binding Update). The LMA then assigns one or more home network prefixes to the MN. The LMA acknowledges the registration process with a Proxy Binding Acknowledgment message that is sent from the LMA to the MAG, containing the home network prefixes allocated to the MN. The MAG completes the configuration to serve the MN traffic by setting up the appropriate forwarding rules for the downlink/uplink traffic to/from the MN.

Different LMAs can coexist in a certain PMIPv6 domain, for instance as a mechanism to perform load balancing. Figure 1 shows several MNs maintaining associations with distinct LMAs. The LMA forwards the traffic for a certain MN to the correct MAG using the configured tunnel. The MAG decapsulates the packets and forwards them to the MN transparently. In the opposite direction, traffic coming from the MN is encapsulated by the MAG (which is the MN's default router) and decapsulated by the LMA that routes it towards the final destination.

2.2 Multicast basics in access networks

By means of IP multicast, a number of receivers located anywhere in the network can subscribe to a content in the form of a multicast session group. The content is distributed using a particular data stream forming a multicast flow. A single copy of such flow is carried on every link in the network along the multicast path dynamically created to reach the interested receivers. The data stream is replicated on the routers where the multicast path topologically diverges.

The multicast source of the data stream does not maintain any subscription list of interested receivers. The source simply sends the data stream to an arbitrary group of hosts represented by an IP multicast address. The receivers indicate their interest in receiving certain content by explicitly joining the multicast group. The Multicast Listener Discovery (MLD) [5] defines the control messages for managing the group membership process in IPv6. Multicast protocols distinguish between multicast receiver (host part) and multicast router (network part) functionalities. Basically, the host part is devoted to the group subscription management, while the router part is focused on building and maintaining the multicast tree.

A multicast router in the receiver's sub-network will capture the control messages for joining or leaving a multicast group. In some cases, the router can act as a proxy [6] for the group membership indications of the receivers connected to it, instead of the multicast router role described above. This typically occurs in aggregation networks, where the first-hop router concentrates the traffic of a huge number of receivers. The proxy performs the router part of the group membership protocol on each downstream interface, while it plays the host role on the upstream interface towards the next multicast backbone router. The proxy is in charge of summarizing the subscription demand of the receivers.

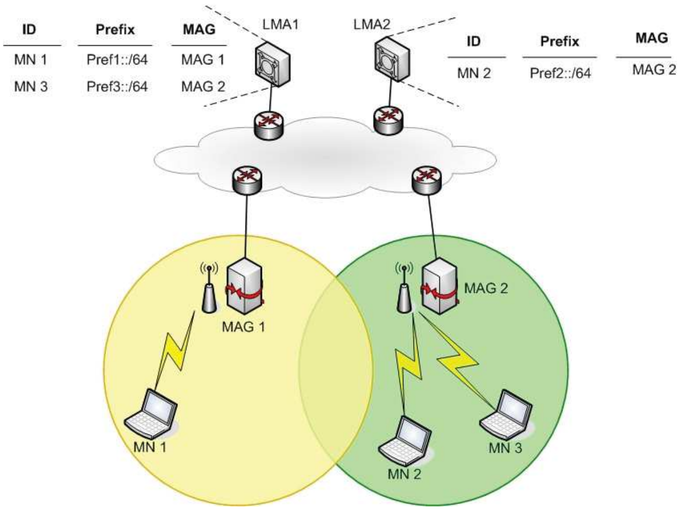

Fig. 1. Network entities in Proxy Mobile IPv6

2.3 Multicast in PMIPv6

Multicast is out of the scope of the PMIPv6 standard specification. This produces inefficiencies when distributing contents to multiple receivers (individual copies per MN). Several solutions have been proposed supporting multicast for Mobile IP networks [7]. However, they are not directly applicable to PMIPv6 because of the specificities of network-localized management environments, where the MN is not aware of network layer changes. A new approach is needed to provide multicast in PMIPv6 domains. With such aim, the MULTIMOB working group was chartered at the Internet Engineering Task Force (IETF) to specify a solution for multicast listener mobility compatible with the PMIPv6 and multicast standards.

Two kinds of solutions for multicast subscription in PMIPv6 can be differentiated: the remote subscription, where the MN gets the multicast data stream from the LMA, and the local subscription, where the MN directly obtains the multicast stream from the access router. According to the current MULTIMOB charter terms, only the remote subscription case is considered by now.

The base solution [3] provides a way to manage multicast traffic delivery to MNs unaware of their mobility. The MN expresses its interest in joining or leaving a multicast group by sending MLD control messages to the MAG, which acts as the first hop router for the MN. The MAG maintains the individual multicast status of the MN and handles the multicast traffic towards it accordingly to the MLD messages received. The MAG incorporates the functionality of MLD proxy, summarizing the group subscription requests of the MNs connected to it.

With the remote subscription model, the multicast traffic reaches the MNs after going through the corresponding LMA (note that there might be multiple LMAs in the same domain). A distinct MLD proxy instance is then defined per LMA connected to the MAG, in such a way that every MAG-LMA tunnel is part of a separate MLD proxy domain. For every proxy instance in the MAG, the tunnel interface pointing to the LMA becomes the proxy upstream interface, whereas the links towards the MNs are the corresponding downstream interfaces of each instance. The LMA is the entity in charge of interacting with the multicast infrastructure out of the PMIPv6 domain.

The summarization of control messages in upstream that the MAG performs is applied per set of MNs associated with a certain LMA, as the different proxy instances of the same MAG are isolated one from the other. The LMA maintains the multicast state of every tunnel interface. Such status reflects the summarized view offered by the MAG on behalf of the attached MNs bound to the LMA. A multicast stream will be delivered over the tunnel or removed from it according to the aggregated behavior of the MNs attached to the MAG.

The LMA, the MAG and the tunnel linking them are all part of the multicast tree. This branch will be common to every multicast tree providing any content subscribed by an MN in a MAG and associated to a particular LMA. It makes possible to send a single copy of a data stream per group of MNs demanding the same content.

3 Dedicated LMA for multicast traffic

In the base solution, the MAG can receive multiple copies of the same stream if several attached MNs associated to different LMAs demand simultaneously the same content. The Figure 2(a) shows this issue. Each LMA will forward the same stream to a certain MAG by using their corresponding LMA-MAG tunnel. Under these circumstances, multiple copies arrive to the MAG, creating tunnel convergence problem.

As an enhancement to this situation, a dedicated multicast LMA (M-LMA) [4] can be deployed to act as the unique mobility anchor for the multicast traffic to all the attached MNs in the PMIPv6 domain, whereas the unicast traffic is handled by the rest of LMAs, as in the PMIPv6 specification.

The M-LMA maintains a tunnel with every MAG in the domain for all multicast channels. This tunnel will be the path followed by the multicast traffic towards the MAG for all the MNs attached to it, thus avoiding the tunnel convergence problem, because only one stream copy arrives to the MAG. The Figure 2(b) represents graphically this solution.

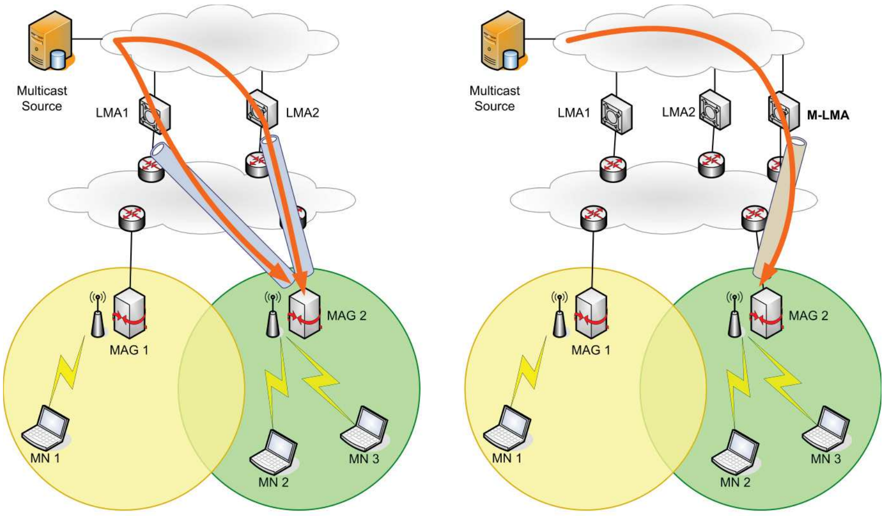

(a) (b)

Fig. 2. Tunnel convergence problem in base solution (a) vs. M-LMA proposal (b)

A number of advantages can be envisioned, such as the simplification of the multicast distribution tree topology or the significant reduction of the network bandwidth needed for providing multicast service within a PMIPv6 domain

The base solution specification includes a comparison with the M-LMA proposal which roughly evaluates the performance of each of them in a couple of extreme scenarios. The population of MNs in the comparison is set to 1 million which are uniformly distributed among 200 MAGs. These MAGs are then connected to 4 LMAs, for the base solution, or 1 LMA, for the M-LMA approach.

The comparison is based on three metrics: the number of simultaneous streams delivered per LMA, the total number of streams in the network for serving the connected MNs, and the number of repeated streams at MAGs to account for the tunnel convergence issue. The comparison is presented in Table 1.

This analysis is a theoretical exercise to study system scalability, but we argue that it does not represent a realistic scenario, so the conclusions on scalability obtained from it are not well supported. First, the model of the multicast channel demand (all the MNs subscribed to the same content in one case, all the MNs subscribed to a different content in the other) does not follow a realistic pattern, neither in terms of subscription behavior nor in terms of multicast flow volume. Second, the number of active multicast MNs and the MAGs supporting them seems to be unrealistically high.

In this paper we provide a more realistic comparison of both systems to better characterize the efficiency and scalability of them.

Table 1. System comparison with extremal settings in [3]

Case scenario	**Multicast scheme**	**Redundant streams per MAG**	**Simultaneous stream per LMA**	**Total streams in the domain**
Each MN subscribed to a different content	Base solution	0	250.000	1.000.000
	M-LMA	0	1.000.000	1.000.000
All MNs subscribed to the same content	Base solution	3	200	800
	M-LMA	0	200	200

4 Simulation framework

4.1 Channel popularity model

The channel preference in IPTV-like systems is commonly modeled [8, 9, 10] by a power-law distribution known as Zipf function. The Zipf function states that the occurrence of a certain event (here, the tuning of a multicast channel) is determined by:

$$k \times \left(1/r^{\alpha} \right), \tag{1}$$

being k a constant, r the rank or popularity of the event in the distribution, and α the factor which characterizes the skewness of the distribution. Then, the frequency or probability that predicts the eligibility of an event is provided by:

$$\frac{\left(1/r^{\alpha} \right)}{\sum_{i=1}^{R} \left(1/i^{\alpha} \right)}, \tag{2}$$

where R is the total number of ranked elements. As α increases, the popularity of the first ranked events increases, while the distribution tail concentrates less occurrences. We will consider also here a Zipf-type function to model the channel demand by the MNs in the domain.

4.2 Description of the simulation

We have used the numerical computation tool Octave [11] to simulate the multicast channel demand in a PMIPv6 domain, and to calculate the number of streams required per LMA-MAG tunnel defined in the system. The simulation focuses on the streams delivered in a certain time instant, thus not including the impact of user mobility during a longer observation period. The methodology followed here to simulate the system demand is basically the same approach followed in [10].

We firstly obtain a uniform random variable in the range (0,1) with as many samples as the number of MNs considered in the simulation (all are supposed to be multicast active and tuned to a channel). We then map the random variable to a Zipf function defined by both the skewness parameter α and the number of channels C, in order to simulate the MNs channel subscription. At this step we already have a picture of the channel demand in the PMIPv6, where each MN is associated with a tuned channel. Now, it is time to compare the behavior of the architectures under consideration.

To do that, we sequentially split twice the MNs in different groups, first accordingly to the number of LMAs, and then accordingly to the number of MAGs, in such a way that we form a matrix with the following elements: number of MNs, channel subscribed by the MNs, MAG where the MN is attached to, and LMA where the MN maintains a service association. In these conditions we are able to calculate the number of multicast streams needed per LMA-MAG tunnel in the domain, and to obtain the metrics previously used for scalability comparison.

The process is entirely repeated several times with random and independent sets of MNs subscriber choices (more than 100 times) guaranteeing convergent results in terms of mean and standard deviation values. The results obtained are the average values of the metrics under study.

4.3 Simulation parameterization

In order to develop a sensitivity analysis of both network solutions for the metrics defined above, we will consider different values for the parameters in the simulation.

In each iteration within the simulation routine, the same Zipf-like distribution representing the MN channel preference is used to coherently compare the impact of the distinct parameters against the considered network solution. The only case where different sets of subscriptions are compared is the case for the sensitivity on the number of MNs in the system, as it necessarily requires a distinct simulation universe. In our simulation we study the demand created for 6.000 and 12.000 MNs, respectively, all requiring multicast service.

The scenarios used for the analysis consists of a unique LMA, in the case of the M-LMA approach; and 2 LMAs or 4 LMAs, for the case of the base solution. In the later case, the variation in the number of LMAs provides an insight on the LMA scalability for multicast stream replication. Some other impacts can be evaluated, as follows:

- The impact of the user preferences on the system can be determined by the variation on the skewness factor of the Zipf distribution. The values considered in the simulation for α are 0.6 and 0.9 [8, 9, 10].
- The impact of the service provider content offer can be modeled by the variation on the number of eligible channels, that is, in the variation of C. In this simulation it takes values of 150 and 300 channels.
- The impact of the network access capillarity can be modeled by the number of MAGs in the domain. We study the impact of having 6 or 12 MAGs in the system.

We will confront every simulation scenario respect to the architecture defined by the number of LMAs taken into account. The different scenarios will be identified in the rest of the paper according to the rule defined by *{number of MNs, skewness factor, number of channels, number of MAGs}* (e.g., the scenario identified as *{6000, 0.6, 150, 6}* means that we are considering 6000 MNs which create a channel demand defined by a Zipf distribution with skewness factor of 0.6 over 150 channels, and that are evenly distributed in 6 MAGs).

5 System analysis

The following sub-sections show the sensitivity of the two solutions regarding to the parameters of the simulation. We first focus on the metrics of the number of streams per LMA, and the number of total streams in the PMIPv6 domain. The number of repeated streams per MAG is analyzed separately.

5.1 Impact of the content offer

Figure 3(a) depicts the impact of the content offer in the number of average streams per LMA. According to the figure, the number of streams per LMA grows with the number of channels accessible for the MNs. More channels in the system mean more distinct multicast streams needed to serve the MNs demand (all the channels are susceptible of having MNs tuned on them in our simulation).

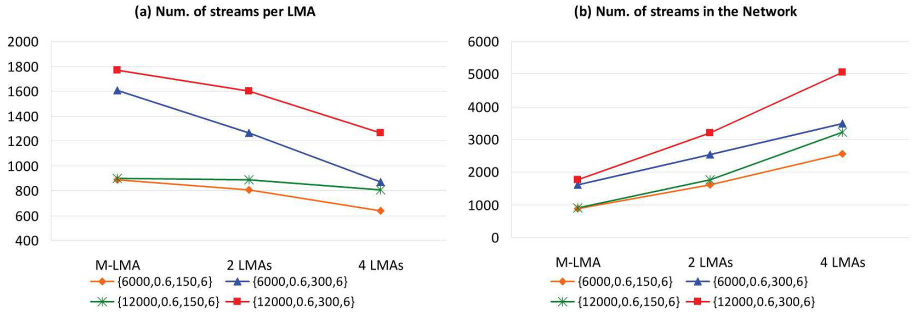

Fig. 3. Impact of multicast channel offer – (a) Number of streams per LMA; (b) Number of streams in the network

The requirements in terms of stream reception for the LMAs in the base solution are certainly more scalable than the M-LMA approach, improving the scalability as the number of served MNs grows. In the case of the base solution, the MNs are split among the LMAs, so if no MN associated with a particular LMA is subscribed to a particular channel, the LMA would not require receiving that channel. In the M-LMA approach, the unique LMA serves all the MNs, which means that it has to receive all the channels subscribed by any of the MNs in the PMIPv6 domain. However, the requirements for replicating multicast streams in the LMA in the base solution are not far from those needed by the M-LMA. This is because each LMA in the base solution has to send a copy of a channel to a MAG as soon as at least one MN associated with that LMA and attached to the MAG requires that channel. Popular channels will be replicated by every LMA and sent to each MAG that receives multiple copies of the channel, while in the M-LMA solution just one stream will be sent from the M-LMA to each MAG.

In fact, Figure 3(b) shows that the M-LMA scales better from the network perspective because the M-LMA architecture requires much less streams than the base solution to serve the same set of MNs. The increment in the number of streams as more channels are offered by the service provider is translated in a progressive increment of the average number of streams required by the base solution. The situation becomes worst as the number of LMAs in the domain increases.

5.2 Impact of the user preferences

Figure 4(a) presents the impact of α in the number of streams transmitted per LMA. As a general rule, greater values of α reduce the average number of streams per LMA. This is basically motivated by the shift in the Zipf distribution which accumulates the audience in the first ranks of the distribution. Some channels in the tail will not be necessarily tuned by any MN associated with an LMA and attached with some MAG, reducing the number of streams per LMA. It can also be observed that the streams per LMA increases with the number of MNs, because some of the previous channels not demanded before in some tunnels become now requested.

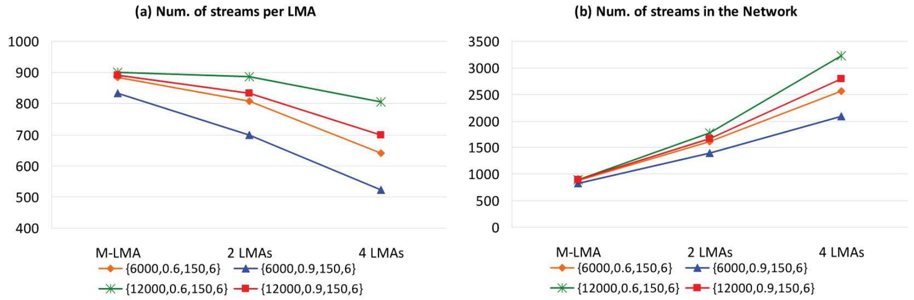

Fig. 4. Impact of the skewness factor α – (a) Number of streams per LMA; (b) Number of streams in the network

A relevant result is that the LMA scalability requirements are quite similar in both the base solution and the M-LMA approach. This is more notorious when the number of MNs increases. In the base solution, the MNs attached to a MAG are associated to different LMAs. Each one of these LMAs have to serve their respective MNs demand which will be composed of a wide set of common channels, thus driving to the tunnel convergence problem. This simulation confirms that when the number of MNs grows with respect to the number of channels, which is a reasonable configuration in multicast environments, the set of common channels that have to be managed by every LMA also grows, minimizing the scalability advantage of having several LMAs, while suffering the problem of the multiplication of stream copies in the domain.

Figure 4(b) shows the total average number of streams in the network for different values of α. It becomes also clear that the number of total streams decreases when α

increases due to the shift of the Zipf distribution which describes the MNs channel demand. The same arguments as before are applicable here. The total number of streams in the network grows significantly with the number of LMAs serving the multicast traffic. The volume of streams required for the service by the 4 LMAs case is close to 3 or 4 times (depending on α) larger than the resources needed by the M-LMA approach, where less tunnels are involved in the multicast service delivery.

5.3 Impact of the network access capillarity

The impact of the number of MAGs in the PMIPv6 domain can be observed in Figure 5. The number of streams per LMA increases if the MNs are distributed among more MAGs with the same channel demand. So, as the number of tunnels increases, the average number of streams per LMA also increases. In the M-LMA case, this implies more streams per (M-)LMA because only one device supports all the growth. For the base solution it can be highlighted that each LMA supports less streams, as the load is shared among the LMAs, but the difference in terms of number of streams supported per device (LMAs vs. M-LMA) are smaller than expected, because there is not a clean split of the load among the LMAs, much of it is replicated in every LMA.

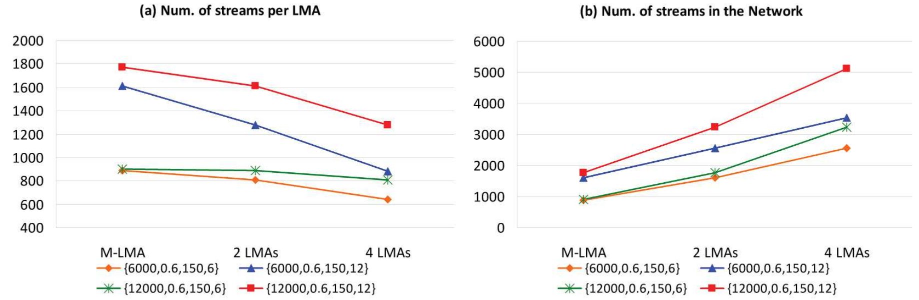

Fig. 5. Impact of the access capillarity – (a) Number of streams per LMA; (b) Number of streams in the network

On the other hand, the increment in the number of MAGs in the domain also impacts on the number of the total streams in the network. The number of streams in the network increases with the number of LMAs and MAGs.

5.4 Repeated streams per MAG

The M-LMA approach does not suffer from the tunnel convergence issue. There is only one LMA to forward multicast traffic towards the MAGs, thus avoiding stream repetition at the MAGs.

The situation is the opposite in the case of the base solution. Several tunnels, one per LMA, can feed the MAG with the same multicast channel. This happens when a MAG has attached MNs that are associated to different LMAs and subscribed to the same content.

Table 2. Average number of repeated streams per MAG

Sensitivity to	Scenario	2 LMAs	4 LMAs
-	{6000, 0.6, 150, 6}	121,24	133,20
Zipf factor α	{6000, 0.9, 150, 6}	94,08	107,60
Number of Channels	{6000, 0.6, 300, 6}	154,39	184,64
Number of MAGs	{6000, 0.6, 150, 12}	78,79	94,07
-	{12000, 0.6, 150, 6}	145,21	148,48
Zipf factor α	{12000, 0.9, 150, 6}	129,50	138,53
Number of Channels	{12000, 0.6, 300, 6}	239,34	263,95
Number of MAGs	{12000, 0.6, 150, 12}	121,19	133,07

Table 2 quantifies the tunnel convergence issue at the MAG, that is, the average number of simultaneously repeated streams in different scenarios. The trend in the number of repeated streams at a MAG is the same as the one followed by the total number of streams in the network. The number of repeated streams decreases with the increment of α, and with the increment in the number of MAGs. On the contrary, it increases when more channels are accessible in the system.

6 Conclusions

In this paper we have compared two approaches to provide multicast content to MNs within a PMIPv6 domain. Neither solution requires modifications to PMIPv6 or to multicast protocols. One is the base solution adopted by the IETF. In this base solution, an LMA serving the unicast traffic also acts as the first multicast router for the mobile nodes in the domain, and the MAGs act as MLD proxies. A PMIPv6 domain can use more than one LMA to serve the visiting MNs, splitting the load among the LMAs both for unicast and multicast traffic. The alternative solution is using a dedicated LMA for multicast traffic, different from the LMAs for unicast traffic, while the MAGs keep their role as MLD proxies.

An initial comparison of both solutions in [3] suggests that the base solution has better scalability properties. In this paper we show that in realistic scenarios this is not the case. The base solution tries to reduce the load in the LMAs, splitting the load by having several of them. However, for multicast traffic, the number of channels that an LMA has to serve does not decrease linearly with the reduction of the number of MNs associated to that LMA. The key factor is the set of channels subscribed by any MN associated to the LMA and attached to certain MAG. So reducing the number of MNs that an LMA has to attend does not result in the expected load reduction. In fact, as the number of MNs increases in the PMIPv6 domain, we have less advantage for having several LMAs, as each of them will probably manage all the multicast channels (or at least the popular ones) anyway. In fact we are just duplicating the work at the different LMAs. The base solution also increases the work load at the MAGs, because of the tunnel convergence issue, in which a MAG will receive several copies of the same flow for MNs subscribed to the same content but bound to different LMAs.

In contrast to this, the dedicated LMA solution for multicast traffic has several advantages: separation of the management of the multicast traffic and the unicast traffic in different LMAs, reduced load in the MAGs, and significant reduction in multicast traffic load within the PMIPv6 domain. This last advantage is due to the duplication that happens in the base solution with several LMAs forwarding the same multicast traffic to every MAG, a situation that does not happen in the solution with a dedicated multicast LMA.

Further work will focus on a more complete characterization of the performance of M-LMA proposal by including user mobility impact on the simulations.

Acknowledgements. The work of Luis M. Contreras and Ignacio Soto has been partially supported by the Ministry of Science and Innovation of Spain under the I-MOVING project (TEC2010-18907) and the EU Celtic UP-TO-US project (TSI-020400-2010-114).

7 References

1. D. Johnson, C. Perkins, and J. Arkko: Mobility Support in IPv6. IETF RFC 3775. (2004)
2. S. Gundavelli, K. Leung, V. Devarapalli, K. Chowdhury, and B. Patil: Proxy Mobile IPv6. IETF RFC 5213. (2008)
3. T.C. Schmidt, M. Waehlisch, and S. Krishnan: Base Deployment for Multicast Listener Support in Proxy Mobile IPv6 (PMIPv6) Domains. IETF RFC 6224. (2011)
4. J.C. Zuniga, A. Rahman, L.M. Contreras, C.J. Bernardos, and I. Soto: Support Multicast Services Using Proxy Mobile IPv6. IETF Internet draft, draft-zuniga-multimob-smspmip-05 (work-in-progress). (2011)
5. R. Vida and L. Costa: Multicast Listener Discovery Version 2 (MLDv2) for IPv6. IETF RFC 3810. (2004)
6. B. Fenner, H. He, B. Haberman, and H. Sandick: Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD) – Based Multicast Forwarding (“IGMP/MLD Proxying”). IETF RFC 4605 (2006)
7. I. Romdhani, M. Kellil, H.-Y. Lach, A. Bouabdallah, and H. Bettahar: IP Mobile Multicast: Challenges and Solutions. IEEE Communications Surveys & Tutorials, 6(1). (2004)
8. Z. Avramova, D. De Vleeschauwer, S. Wittevrongel, and H. Bruneel: Dimensioning Multicast-Enabled Networks for IP-Transported TV Channels. ITC 2007, LNCS, vol. 4516, pp. 6–17. (2007)
9. N. Sinha, R. Oz, and S. Vasudevan: The Statistics of Switched Broadcast. Proceedings of Conference on Emerging Technologies, Tampa, USA. (2005)
10. D. T. van Veen, M. K. Weldon, C. C. Bahr, and E. E. Harstead: An analysis of the technical and economic essentials for providing video over fiber-to-the-premises networks. Bell Labs Technical Journal, 10 (1), pp. 181-200. (2005)
11. GNU Octave, http://www.gnu.org/software/octave (access date: January 2011)