

Lossless Multicast Handovers in Proxy Fast Mobile IPv6 Networks

Bernd Meijerink and Geert Heijenck

University of Twente, The Netherlands
{bernd.meijerink,geert.heijenck}@utwente.nl

Abstract. There is a demand in the Public Protection and Disaster Relief (PPDR) community for high bandwidth services on mobile devices. Group communication is an important aspect of PPDR networks. In IP based networks multicast is the preferred method to efficiently transmit data to more than one receiver simultaneously. It is important PPDR users can switch seamlessly between wireless networks. This paper describes improvements to multicast in Fast handovers for Proxy Mobile IPv6 (PFMIPv6) to provide seamless mobility to its users. We also identify and explore the specific problems stemming from difference in end-to-end delay between the old and new path during handovers for multicast traffic. A novel mechanism to determine the delay difference between two paths in a PFMIPv6 system is described and an implementation of this system is evaluated. It is shown the proposed approach can prevent multicast packet loss during a handover.

Keywords: Fast Handovers for Proxy Mobile IPv6, Multicast, Seamless mobility

1 Introduction

There is an increasing demand within the Public Protection and Disaster Relief (PPDR) community for high bandwidth mobile connections for services like streaming video and audio. PPDR systems are communication systems for public services like the police and fire department. Requirements for PPDR systems are different from the requirements for consumer systems in that they have to be highly resilient and provide strong security. Losing communication in a crisis situation can lead to dangerous situations and in the Netherlands prompted some fire departments to replace part of the TETRA system with analog equipment [1]. These systems differ from the more well-known mobile communications systems such as 3G and Long Term Evolution (LTE) by providing group call functionalities and operating (in a reduced manner) without support from a base station. While consumer demand for high bandwidth is driven by services like YouTube, Netflix and Spotify the demand in the PPDR community is based on live video from emergency sites and voice communication with groups of other users.

Current PPDR implementations used in Europe are mostly based on the TETRA system. While the system provides reliable voice communication among

users, there is very limited support for applications that require high bandwidth such as video streaming. For the future there is a desire from the PPDR community for more bandwidth and more services [2]. The logical step is to switch to newer technologies such as LTE and use existing infrastructure in the form of IEEE 802.11 networks (Wi-Fi) in buildings. To allow data services on the network, such as video and image sharing, a switch to a fully Internet Protocol (IP) based network is useful. This will allow better resilience due to packet switching and easier development of services and applications. Another benefit of a switch to IP is the general trend in mobile network technologies to switch to IP. The results are cost benefits and increasing availability in areas where it might not be efficient to deploy a private PPDR platform.

One of the problems for PPDR users on current networks is that seamless handovers between networks are difficult and lead to packet loss. Packet loss is problematic as it can lead to users missing important data, for example critical frames in a live video of a robbery. This problem exists when switching between similar network technologies (for example IEEE 802.11 to IEEE 802.11) and between different technologies (for example IEEE 802.11 and LTE). Another problem is that PPDR work is characterized by working in groups leading to a requirement for group communication on the network [2]. In traditional mobile networks this is problematic as a separate connection needs to be set up from some central point to all users in the group call. In IP based networks, multicast functionality allows data to flow from a sender to multiple receivers. This reduces network load and bypasses the need for a central server that connects all participants in a group. The lack of multicast support in wireless networks requires each user to send data separately to each other user in the group call, leading to inefficiencies.

Users of PPDR systems would like to be able to roam seamlessly between networks across borders [2]. Mobility in IP networks can be achieved with the Proxy Mobile IP (PMIP) standard [3]. There are also extensions to the standard such as Fast handovers for Proxy Mobile IPv6 (PFMIPv6) that allow seamless handovers between two different networks (technology independent) [4]. Other extensions cover multicast support for these network architectures through the use of multicast proxies at the PMIP network devices.

Currently there are no implementations of the PFMIPv6 protocol or Proxy Mobile IPv6 (PMIPv6) with multicast support integrated available. Supporting seamless handovers for multicast flows puts special demands on the handover processes, as these flows are typically destined to multiple users at the same time. The lack of available PFMIPv6 implementations with multicast support leads us to the following research question. How can we optimize multicast handovers in PFMIPv6 to allow minimal packet loss while minimizing the load on the wireless network?

The contributions of this paper are as follows. We present a method to determine the delay difference between two paths in a PFMIPv6 network. An implementation of the PFMIPv6 protocol supporting multicast listeners was made that makes use of the delay difference estimation. The implementation

is evaluated on performance aspects based on measurements taken in a test environment.

The following section will provide background information on the technologies related to PFMIPv6, multicast and related work. Its parts can be safely skipped by anyone who feels confident in their knowledge in these areas. Section 3 describes the timing problems surrounding multicast fast handovers and a solution to this problem. In Section 4 our implementation is validated by measurements in real wireless environment and finally we draw conclusion in the final section.

2 Background and Related Work

This section will explain the technologies used and built upon for this project and describe related work.

2.1 PFMIPv6

Fast handovers for Proxy Mobile IPv6 (PFMIPv6) is an IP mobility technology. It allows mobile devices to move through a network and switch attachment point without the device noticing this on the IP layer (Figure 1). To explain the way this system works we will start with explaining basic Mobile IP and build up to the more complex systems. It is interesting to note that Proxy Mobile IP is also an optional part of the 3GPP LTE specification, in which it replaces the GPRS Tunnelling Protocol (GTP) when used.

MIP Mobile IP (MIP) works by having a fixed point in the (mobile) network that acts as a mobility anchor for mobile devices. In MIP this entity is called the Home Agent and it will tunnel packets for a Mobile Node (MN) to the network that node is currently on. In MIP the MN knows about the mobility and actively helps maintain it.

PMIPv6 PMIPv6 works by adding an Local Mobility Anchor (LMA) to the provider core network [3]. This LMA is the anchor point for the MNs in the domain (All IP routes for the MNs lead to the LMA). The LMA tunnels the data for the MN to the Mobile Access Gateway (MAG) that serves as the MNs point of attachment. The main improvement over MIP is that the MN has no knowledge of the IP mobility, it always thinks it is connected to the same network and will have the same IP address.

When a handover is performed, the serving MAG lets the LMA know the MN has disconnected. The new MAG (nMAG) lets the LMA know the MN has connected to it and a tunnel is set up between the LMA and the nMAG. Because the nMAG will have the same link layer address as seen from the node it does not know that it is actually on a different link and keeps the same IP address (If the MN requests a new IP address it would also receive the same one it had before).

Fast handovers Proxy mobile IP does not reduce the handover times by itself. An extension called “Fast Handovers for Proxy Mobile IPv6” (PFMIPv6) enables faster handovers by changing the handover procedure in the MAG (The LMA is not changed) [4]. When a handover occurs, the previous MAG (pMAG) will set up a tunnel to the nMAG and forward all traffic it still receives for the MN to the nMAG. The nMAG will buffer all traffic that it receives for the MN until it connects. Once the MN is connected the nMAG will flush all buffered traffic towards it. The PFMIPv6 protocol covers two handover situations: (i) Predictive handover, in which the network knows about the movement of the MN beforehand. (ii) Reactive handover, in which the network does not know about the movement beforehand.

2.2 Related Work

The IETF has a working group on multicast mobility [5]. Several informational and experimental RFCs have been released that deal with multicast handovers. Implementing multicast functionality in Proxy Mobile IP networks is not straightforward. Simply enabling multicast routing on the MAGs and LMA would not lead to the desired results, as multicast traffic would be routed over the normal network instead of the MAG-LMA tunnels. To solve this problem the IETF multicast mobility working group has released several experimental status RFCs regarding this issue. RFC 6224 [6] explains the deployment of Multicast Listener Discovery (MLD) proxies at the MAG and LMA. These proxies query and listen for incoming MLD reports from mobile nodes. These reports are forwarded from MAG to the LMA, which in turn can subscribe to the multicast traffic on the greater network it is connected to. Multicast traffic is forwarded to the MN in a reverse fashion. When it arrives at the LMA it is forwarded to the MAGs which have subscribed MNs. The MAGs in turn forward it to the interfaces with these MNs.

For fast handovers, RFC 7411 [7] contains an extra option for the Handover Initiation and Handover Acknowledgement messages. This option contains all the multicast groups the MN is subscribed to. The nMAG can already subscribe to these groups while the handover is happening. As with unicast traffic it is also




Fig. 1: Proxy Mobile IPv6 architecture

possible to buffer multicast traffic being tunnelled from the pMAG preventing packet loss. This RFC does not deal with the problem of synchronizing multicast flows after the handover. There are also no reference implementations available of multicast fast handover systems.

A member of the IETF multicast mobility working group has published a multicast proxy implementation that can be used in PMIP networks [8]. The system is mainly built as a demonstration for peering between proxies to improve efficiency in situations where the MNs are also multicast senders. Enabling better support for multicast senders is not covered in this paper as the focus is multicast handovers. The application from [8] does however not integrate with fast handovers as is the goal in this paper. Integration with the PFMIPv6 system is vital as multicast subscriptions need to be transferred to the nMAG.

There is a large amount of work available on optimizing handovers for MIP. The main downside of MIP is the modified MN that needs to communicate with the network for mobility management. In [9] the performance of fast handovers in MIP enabled IEEE 802.11 networks is evaluated. It is shown that handover times can still be several seconds with up to 10 seconds in busy networks.

The authors of [10] show that PMIP provides significant improvements over MIP in that it improves handover performance, is link layer agnostic, and reduces the usage of wireless resources. While the paper does not look specifically at multicast all these improvements are also relevant to multicast traffic.

Work has been done on PMIP handovers that lead to the RFC 5949 standard for Fast Handovers for Proxy Mobile IPv6 [4]. This standard describes the handover procedure and how messages between the PFMIPv6 entities should be formatted. This standard will be used to implement fast handovers.

To make a handover appear seamless to a user it is very important that handovers occur as fast as possible. To reach this goal the handover delay needs to be minimized. In [11] it is shown through analytical modelling that using Media Independent Handover (MIH) can greatly improve the handover delay of PMIPv6 networks. The paper does however not cover fast handover situations with multicast. It is shown in [12] that multicast extensions for PFMIPv6 will greatly reduce handover delay for multicast traffic, but the authors did not test an actual implementation and all results were obtained analytically. In [13] it is shown that packets reordering during the buffer flush at the nMAG can be prevented by tagging the last packet sent from the LMA to the pMAG. While this method is shown to work well for unicast traffic, there is no potential end to multicast traffic streaming towards the pMAG, making this method unsuitable for multicast handovers.

3 Multicast Handover Timing

Handovers involving multicast traffic are different than those involving unicast traffic. With a unicast only handover the LMA can simply stop forwarding the traffic to the pMAG and start forwarding to the nMAG. With multicast traffic, and especially in a PPDR context, there is a high chance that there are still MNs

on the pMAG that want to receive this traffic. The result is that the LMA needs to keep forwarding this traffic to the pMAG, and there is no point at which the traffic forwarding can simply stop.

When performing a handover with multicast using PFMIIPv6 there are temporarily two interfaces at which multicast traffic arrives on the nMAG. Multicast traffic arrives at the nMAG from the pMAG-nMAG tunnel and the LMA-nMAG tunnel. Because the data on one tunnel takes a different route than traffic from the other tunnel, there can be a difference in arrival time of the same packet on both tunnels. This time difference can lead to problems during the handover. When the MN connects to the nMAG and the traffic buffer is flushed, the traffic should continue flowing from the LMA tunnel. When for example the LMA-nMAG tunnel has less delay than the pMAG-nMAG tunnel, this might lead to packet loss as these packets have not yet arrived on the LMA-nMAG tunnel. The nMAG needs to make a decision on when to start forwarding from the LMA-nMAG link. The packets themselves cannot be expected to contain information such as sequence numbers, so the nMAG needs a separate mechanism to make a forwarding decision.

Due to the possible differences in delay between the two tunnels there can be 3 different situations at the nMAG: (i) The delay is identical for both routes (Figure 2a), (ii) The LMA-nMAG route is faster than the pMAG-nMAG route (Figure 2b), (iii) The LMA-nMAG route is slower than the pMAG-nMAG route (Figure 2c). Below, we describe the consequences of the scenarios for multicast handovers and their possible resolutions.

No delay difference When the tunnels have the same delay, no special action needs to be taken. The application can simply flush the multicast traffic buffer that has built up during the handover, and switch to the data coming from




Fig. 2: Multicast synchronization

the LMA-nMAG tunnel once it has arrived there (see Figure 2a). In this figure the numbered boxes represent packets travelling to the nMAG. Packets coming from the pMAG are replaced by the packets coming from the LMA. In this case the switch can occur at any moment because there is no difference in the delay between the two tunnels. While this situation is ideal it is not very likely to occur during normal operations as it requires the delay on the direct LMA-nMAG tunnel to be equal to the delay on the LMA-pMAG-nMAG route.

LMA-nMAG tunnel is faster In this situation the multicast packets coming through the pMAG-nMAG tunnel arrive behind those from the LMA-nMAG tunnel. The result, as can be seen in Figure 2b, is that the MN is still receiving packets that would have already arrived over the LMA-nMAG tunnel. A useful solution in this scenario is to simply forward the packets when they arrive at the nMAG, and stop the pMAG-nMAG tunnel forwarding once all the data that was still missing from the LMA-nMAG tunnel has arrived. The duration that data still needs to be forwarded from the pMAG-nMAG tunnel is the delay difference between the two tunnels. This switching procedure ‘mixes’ packets from both tunnels until all packets that were missing on the LMA-nMAG tunnel have been transmitted to the MN.

pMAG-nMAG tunnel is faster In this situation the new LMA-nMAG tunnel contains older data than the LMA-pMAG-nMAG tunnel. It is possible to simply send all the packets from the LMA-nMAG tunnel onto the network, but this would waste wireless bandwidth. A solution in which the traffic is simply not forwarded until the LMA-nMAG tunnel is caught up with the LMA-pMAG-nMAG route is most efficient in terms of network resources. In the example (Figure 2c) the packets from the pMAG-nMAG tunnel are stopped forwarding once the LMA-nMAG tunnel is established and the packets from that tunnel start forwarding once there is new data for the MN on that tunnel.

3.1 Delay difference estimation

To make the correct decision on how to perform the multicast handover, the nMAG needs to know the difference in arrival time between the two tunnels. Traditional methods to determine one way delays are based on precise clock synchronization [14]. The delay estimation problem here is unique in that there is identical traffic arriving on two different interfaces and we are not interested in the delay of the link itself, but in the relative delay between the links.

We propose to send special timing packets from the LMA over all the LMA-MAG tunnels on a fixed interval. A timing packet consists of a UDP datagram containing a 16 bit sequence number and 32 bit timing interval. Timing packets are multicast down all tunnels. When the pMAG initiates the pMAG-nMAG tunnel for a handover, it will forward the incoming timing packets from the MN’s LMA over this tunnel. The nMAG will now receive timing packets on the LMA-nMAG and pMAG-nMAG links. It can use the difference between the

received time and the difference in sequence numbers on each link to calculate the difference in transmission times between both links. The delay difference between the two tunnels TD in ms is given by:

$$TD = (SeqN_{Tunnel1} - SeqN_{Tunnel2}) \times I - (T_{Tunnel1}^{arrival} - T_{Tunnel2}^{arrival})$$

In this formula $SeqN_{tunnel}$ is the sequence number of the last received packet on that tunnel. I is the timing interval the packets are sent on in ms (taken from the timing packet). $T_{tunnel}^{arrival}$ is the arrival time of the last packet for that tunnel in ms Unix time. Calculating the delay difference in this manner only requires us to store the last received sequence number and arrival time per link, greatly simplifying the system.

Once the nMAG knows the delay difference between the tunnels, decisions on how to switch the two multicast tunnels can be made based on the difference between them. The data on the difference is at most I ms old, the accuracy of the data is thus dependent on how often the timing packets are transmitted. In very congestion networks with a high variation in delay the value for I should be reduced to compensate possible sudden changes in the delay.

The benefit of this solution, over more traditional methods such as Internet Control Message Protocol (ICMP) echo requests, is that it does not take a round trip time to measure the delay. The timing packets are always arriving when a tunnel exists. Other active methods to find the one way delay require precise time synchronization between the sender and receiver. The maximum time that is required to measure the transmission delay between two tunnels is the transmission interval I , with an average of $\frac{1}{2}I$.

4 Validation

This section will present the validation of the PFMIPv6 and delay estimation implementation. The goal is to measure how effective the implemented system is in measuring the delay difference between two paths on the network and performing lossless multicast handovers. In order to validate the proposed approach, we extended an existing PMIPv6 implementation [15] with fast handover and multicast functionality. We also implemented the described time difference and handover optimization mechanism for multicast traffic.

4.1 Measurement network

The validation was done based on measurements taken from a PFMIPv6 network built for this purpose. Figure 3 shows the different devices in the network and how they are implemented on physical hardware. The LMA, MAGs and Corresponding Node (CN) are virtual machines running in VMWare (The box in Figure 3). The CN will multicast packets to the MN. The MAGs are each connected to physical IEEE 802.11 interfaces that are connected to the host computer using USB. The USB devices are shared with the MAG VMs to enable direct access to




Fig. 3: PFMIPv6 experiment network

the IEEE 802.11 hardware. The MN (A laptop with an IEEE 802.11 interface) will connect to the MAGs through the MAGs wireless interface. This laptop is running a recent Ubuntu Linux distribution with a kernel that was compiled with Optimistic Duplicate Address Detection (DAD). While Optimistic DAD decreases the time need for a layer 3 handover significantly, there is still a delay of above 2 seconds for the handover on the layer 2 level. It is outside the scope of this paper to attempt to further reduce this delay. It should be noted that our PFMIPv6 implementation transmits buffered multicast packets as unicast link layer packets to the MN. This is done to allow for higher data rates than would be possible using unicast, and to prevent other nodes from receiving packets they already have. This approach also reduces losses during the buffer flush.

4.2 Measurement method

Measurements were done using UDP multicast traffic flowing from the CN to the MN for 20 seconds. The handover trigger is sent exactly 5 seconds after the CN starts sending multicast traffic. The layer 2 handover takes 2.4 seconds on average after which the buffered traffic is flushed to the MN. Different packet loads will be used to measure the effect of the number of packets per second on the performance of multicast handovers.

In a larger network it would be common to find differences in the one way delay between the connections of the PFMIPv6 devices, due to delays in the routing and the distances involved. To simulate these delays on the measurement network, delay was added to specific tunnels. For this experiment several delay values will be used with or without an added variance in the delay. The delay between consecutive packets is distributed normally. It was shown in [16] that a normal distribution approximates real world packet delay. We will use 50 ms with 10 ms standard deviation, 100 ms with a 20 ms standard deviation and 200 ms with a 40 ms standard deviation as our different delay values. The delay

given to a packet will be correlated 50% with the previous packet. The delay is added using the NETEM functionality of the Linux kernel [17].

The timing packets were sent at a 500 ms interval. A single packet gives an overhead on the wired network of 94 bytes (2 IPv6 headers because of tunnelling, a UDP header, and 48 bits of data). In the test network the overhead amounts to 188 bytes per tunnel per second.

The measurements done can be divided into three groups: (i) Measurements with no added delay, (ii) Measurements with added delay on the pMAG-nMAG tunnel, (iii) Measurements with added delay on the LMA-nMAG tunnel. These groups cover the different situations that can occur on the network with respect to delay differences. If the delay is accurately measured at the nMAG it can take the appropriate forwarding action to ensure minimal packet loss and wireless network load.

When a handover occurs the nMAG will output its delay difference estimate. We will take all measurements of a single group (For example all measurements where the LMA-nMAG tunnel has 100 ms delay with 20 ms standard deviation) and calculate their average and standard deviation. We will compare these measured results to the delay settings the link was configured with. Each group of measurements will contain at least 30 measurement points.

For each measurement the packet loss will be recorded on the MN. Because multicast traffic is transmitted as broadcast packets on IEEE 802.11 networks, there is no retransmission mechanism as with unicast packets. The result of this unreliable delivery method is packet loss for multicast packets on the wireless interface. Because this background loss is not caused by the handover process itself, and we measure over a period longer than the handover, we try to compensate for this loss in the results. The packet loss measured surrounding the handover will be used to compensate for background losses during the handover. This compensation is done by subtracting the packet loss during the first 5 seconds of the measurement to compensate for losses before the handover. Compensating for losses after and during the handover is done by taking the number of packets lost from seconds 14 until 19 and multiplying them by 3 to compensate for losses that would otherwise have occurred during seconds 5 until 20. In an ideal scenario the packet loss will be close to 0. Each measurement will be run at least 10 times for each variable. The average of each measurement set will be reported with its 95% confidence interval.

4.3 Results

In Figure 4 we can see the delay that our measurement method reported at the nMAG during the handover. The measurement points are the mean calculated over all the measurements done at that specific delay. The bars around the measurement points represent the standard deviation of these results. We can see that the standard deviation correlates with the values used by NETEM to add delay to the links. A more detailed view of a single point in this graph (100 ms delay on the pMAG-nMAG link relative to the LMA-nMAG link) can be seen in Figure 5. In this graph we see the amount of delay measurements that fell within




Fig. 4: Measured delay




Fig. 5: Distribution of measured delay

a certain range. We can see this represents a normal distribution which can be attributed to the delay variation added in the experiment. Based on these results we can conclude that our system accurately finds the delay at the specific point in time the measurement is taken.

Figure 6 shows a graph of the throughput during the entire measurement process for PFMIPv6 handovers with different traffic loads with no added delay. The first 5 seconds the MN is connected to the pMAG and the throughput is around the value the data is sent at. At the 5 second mark the handover is initiated, and the throughput drops to 0 during the time the MN is disconnected from the network. After the handover is completed there is a visible spike in the throughput as the pMAG flushes the packet buffer to the MN.

The results for the packet loss during the measurements can be seen in Figure 7. In this graph the y-axis shows the packets lost during the handover and the x-axis shows the traffic load. The labels are marked with the component that added the delay with the delay distribution that was used (LMA (50,10) means that 50 ms delay with a 10 ms standard distribution as variance was added on the LMA-nMAG link). Compensating for losses that occur due to the unreliable




Fig. 6: Throughput during measurement

nature of multicast on IEEE 802.11 networks was not as successful as hoped, leading to negative values on the y-axis. The cause of this results has three components: (i) Compensating for loss due to link unreliability only works well when the loss occurs equally over the measurement period. We cannot control the loss during the measurement as it fully dependent on the wireless channel. (ii) More importantly, the delay difference estimation does sometimes not accurately represent the actual delay due to the delay distribution. The delay measurement in the current system only occurs every 0.5 seconds. If the delay of one or both of the timing packets is an outlier this will cause the system to think the delay is much higher or lower than is actually the case in this timeslot. We can see this effect on the results with 0 variance in the delay as they are closer to 0. Decreasing the interval on which the delay is estimated and taking into account previous measurements can potentially lead to a better estimate for the current situation with high delay variance. (iii) As the buffer is unicast towards the MN these packets are less likely to be lost. Packets arriving during the buffer flush are however transmitted using multicast during this period so we still have to compensate for this loss, leading to overcompensation in most cases. The effect is exaggerated at higher packet rates, because the loss surrounding the handover are higher than during the handover.

5 Conclusions and Future Work

This paper describes optimizations to the multicast handover process in PFMIPv6. We presented a mechanism to determine the delay difference between two paths in a PFMIPv6 network. The values found with this method are used to optimize the multicast handover process. The optimizations and delay finding mechanism were implemented and validated with measurements done in an actual deployment with IEEE 802.11 networks. It was shown that our delay estimation mechanism can accurately find the delay difference between two paths in a PFMIPv6 network.

We avoid transmission of duplicate packets on the wireless network during a handover by switching the interface the multicast packets are forwarded from at




Fig. 7: Compensated packet loss during handover

the right moment. The decision of when to switch between the pMAG-nMAG and LMA-nMAG tunnels is based on the difference in delay between the two tunnels. This prevents multicast packets from being transmitted twice, reducing the load on the wireless network.

The interval the timing packets are sent on needs to be tuned. Increasing the time between timing packets puts less load on the network, but makes the estimated delay between tunnels less accurate. In a network with high amounts of congestion or large variances in delay, it could be beneficial to have the timing packet interval much lower than on a low congestion network. Previous measurements will also need to be taken into account. A single timing packet can currently cause a handover to perform sub-optimally if it was delayed more or less than the average delay during the handover. Our current PFMIIPv6 implementation is still missing reactive handover support. Handovers that are made without initiating them first are performed as normal PMIPv6 handovers. Support for multicast senders needs to be added to the implementation. While it is already possible for a MN to act as a multicast sender improvements are possible to allow multicast traffic to flow from MAG to MAG directly, skipping the LMA and reducing latency.

Multiple MNs listening to the same group on a MAG present extra synchronization problems. These mainly occur on the air interface. If MN A performs a handover to a MAG that already has a MN B listening to the same multicast group, it becomes difficult to deliver the same data on the air interface twice. On IEEE 802.11 networks this is a problem as the multicast packets from both sources (pMAG-nMAG tunnel and LMA-nMAG tunnel) might be broadcast at the same time. The same problem occurs when two nodes from different LMAs are connected that listen to the same multicast group.

Acknowledgement The authors would like to thank the financial support provided by the SALUS project, co-funded by the EU under the 7th Framework Programme for research (grant agreement no. 313296).

References

1. Computable: Brandweer vervangt C2000-portofoons (2013), <http://www.computable.nl/artikel/nieuws/telecom/4923768/1276977/brandweer-vervangt-c2000portofoons.html>
2. Critical Communications Broadband Group: The Strategic Case for Mission Critical Mobile Broadband. Tech. rep., TCCA (2013)
3. Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., Patil, B.: Proxy mobile ipv6. Tech. rep., RFC 5213, August (2008)
4. Yokota, H., Chowdhury, K., Koodli, R., Patil, B., Xia, F.: Fast Handovers for Proxy Mobile IPv6. Tech. rep., RFC 5949, September (2010)
5. IETF Multicast Mobility Working Group: Multimob status pages (2014), <http://tools.ietf.org/wg/multimob>
6. Schmidt, T., Waehlich, M., Krishnan, S.: Base deployment for multicast listener support in Proxy Mobile IPv6 (PMIPv6) Domains. IETF RFC6224, April (2011)
7. Koodli, R., Waehlich, M., Fairhurst, G., Schmidt, T., Liu, D.: Multicast Listener Extensions for Mobile IPv6 (MIPv6) and Proxy Mobile IPv6 (PMIPv6) Fast Handovers (2014)
8. Schmidt, T.C., Wölke, S., Wählisch, M.: Peer my Proxy - A Performance Study of Peering Extensions for Multicast in Proxy Mobile IP Domains. In: Proc. of 7th IFIP Wireless and Mobile Networking Conference (WMNC 2014). pp. 1–8. IEEE Press, Piscataway, NJ, USA (May 2014)
9. Montavont, N., Noel, T.: Handover management for mobile nodes in ipv6 networks. Communications Magazine, IEEE 40(8), 38–43 (Aug 2002)
10. Kong, K.S., Lee, W., Han, Y.H., Shin, M.K., You, H.: Mobility management for all-ip mobile networks: mobile ipv6 vs. proxy mobile ipv6. Wireless Communications, IEEE 15(2), 36–45 (April 2008)
11. Magagula, L., Chan, H.: Ieee802.21 optimized handover delay for proxy mobile ipv6. In: Military Communications Conference, 2008. MILCOM 2008. IEEE. pp. 1–7 (Nov 2008)
12. Lee, J.H., Ernst, T.: Fast pmipv6 multicast handover procedure for mobility-unaware mobile nodes. In: Vehicular Technology Conference (VTC Spring), 2011 IEEE 73rd. pp. 1–5 (May 2011)
13. Quoc, A.K., Kim, D., Choo, H.: A novel scheme for preventing out-of-order packets in fast handover for proxy mobile ipv6. In: Information Networking (ICOIN), 2014 International Conference on. pp. 422–427 (Feb 2014)
14. De Vito, L., Rapuano, S., Tomaciello, L.: One-way delay measurement: State of the art. Instrumentation and Measurement, IEEE Transactions on 57(12), 2742–2750 (2008)
15. Open Air Interface: Open air interface proxy mobile ipv6 (OAI PMIPv6) (2011), <http://www.openairinterface.org/openairinterface-proxy-mobile-ipv6-oai-pmipv6>
16. Elteto, T., Molnar, S.: On the distribution of round-trip delays in tcp/ip networks. In: Local Computer Networks, 1999. LCN'99. Conference on. pp. 172–181. IEEE (1999)
17. Linux Foundation: netem (2009), <http://www.linuxfoundation.org/collaborate/workgroups/networking/netem>