A Hybrid of Network and Application Layer Multicast for Mobile IPv6 Networks

A. Garyfalos
BTexact Technologies
BT Adastral Park
IPswich, UK
anargyros.garyfalos@bt.com

K. Almeroth
Computer Science
Univ of California
Santa Barbara, CA
almeroth@cs.ucsb.edu

J. Finney
Computing Department
Lancaster University
Lancaster, UK
joe@comp.lancs.ac.uk

***Abstract*— In this paper we evaluate a spectrum of various approaches to group communication for MIPv6 nodes: Network Layer Multicast (IP multicast), Application Layer Multicast(ALM) and a hybrid approach. Based on a previous comparison of the two extremes, we focus on the hybrid solution and we propose a number of optimizations. Our results show that although IP Multicast is still the best option in terms of network performance, due to deployment considerations, the hybrid approach is a more attractive option.**

I. Introduction

Network Layer (IP Multicast) and Application Layer Multicast (ALM) are two different approaches to support group communication in general. Focusing on Mobile IPv6 networks, our previous study[1] has shown that mobility introduces several new challenges for ALM that do not exist in wired networks. System stability, heterogeneity and node capability become critical problems since mobile nodes may be less capable or more constrained in their ability to act as ALM end hosts. Moreover, in terms of network performance IP Multicast outperforms ALM both for low and high speeds. However, given deployment concerns, a ubiquitous native multicast deployment might not be possible. Therefore a third alternative is required.

This paper investigates the potential of a hybrid system in which inter-domain multicast support is provided using ALM and intra-domain support is provided using native multicast. Although similar hybrid systems have been proposed in the past, our contribution is twofold. First, we examine the behavior of the system for mobile receivers. Secondly, we propose a set of modifications necessary to handle the mobility challenge. To evaluate the performance of our hybrid scheme we perform a set of simulations and show that it performs well.

The remainder of this paper is organized as follows. First we explain the motivation for a hybrid system by examining the impact of mobility on IP Multicast and ALM. Section 3 contains a more in-depth discussion of the architecture and deployment options of our hybrid system. Section 4 presents the simulation results and Section 5 concludes the paper.

II. Background

This section explains the motivation for a hybrid system by discussing the problems of pure IP Multicast and pure ALM. First we examine Mobile IPv6 (MIPv6) and describe how it affects IP Multicast and ALM. This is followed by a detailed performance comparison between the two schemes. We conclude with a summary of the main points.

A. Mobile IPv6

MIPv6 is a protocol which allows nodes to remain reachable while moving around in the IPv6 Internet[2]. Mobile nodes may receive packets in one of two ways. In *Reverse Tunneling* a router in the home domain (called the *Home Agent*) intercepts packets and tunnels them to the mobile node's location. In *Optimized Routing* packets are addressed directly to the new location. Transmission follows a similar pattern, with either direct (from the remote location) or indirect (tunneled through the Home Agent first) operation.

Multicast operation is influenced by these two modes. In *Remote Subscription* the node joins via a (local) multicast router on the foreign link being visited. Alternatively, in *Home Subscription* the mobile node joins the multicast group via a bi-directional tunnel to its home agent. Membership messages are tunneled to the Home Agent, which then forwards multicast packets down the other end of the tunnel. While the use of reverse tunneling can ensure that multicast trees are independent of the mobile node's movement, the round-trip time between the foreign subnet and the Home Agent may be significant. In addition, the delivery tree from the Home Agent in such circumstances relies on unicast encapsulation from the agent to the mobile node and is therefore bandwidth inefficient compared to native multicast forwarding.

B. The Case for IP Multicast

IP Multicast may be regarded as the traditional way of supporting group communication. Compared to the one-to-one operation of unicast and the one-to-all of broadcast, IP multicast is a more efficient way of addressing a *group* of receivers. By adding special functionality in the network, it allows packets to be routed to a specific set of end hosts using fewer network resources. However, due to both technical and commercial reasons, deployment of IP Multicast has not been fully deployed. The existence of numerous protocol specifications (PIM-SM, PIM-BIDIR, SSM, DVMRP, MOSPF, etc.) has resulted in “islands” of multicast capability. This, in conjunction with various technical complications, has made inter-domain deployment a major issue.

Concerning mobility, protocol operation becomes even more complicated. As MIPv6 provides only the basic mechanisms to enable multicast operation for mobile nodes, a set of remaining open issues has had to be addressed. Other efforts have proposed a variety of solutions. Some aim to dynamically change between home and remote subscription schemes, others deploy a hierarchical network infrastructure, while others focus on join delay issues by proactively joining to-be-visited networks. However, this plethora of potential modifications has only managed to increase the complexity and scepticism over actual deployment. System administrators are now even more confused since, in addition to deciding which IP multicast protocol to deploy, they must also decide on which modifications to apply.

C. The Case for ALM

In order to solve the deployment issues associated with IP Multicast, a number of ALM solutions have been proposed. Overall, ALM is an attempt to overcome the complexity of native multicast by sacrificing a portion of the network efficiency gains for increased deployability. Numerous efforts have been published with interesting, and at times, encouraging results. Nevertheless, to the best of our knowledge, the combination of ALM and mobility has not been extensively examined. Two notable exceptions[3], [4] show how an overlay network could support mobile nodes. However, since these scheme do not use MIPv6, they fall outside the scope of our focus.

The main reason that ALM protocols disregard node movement is because they claim to be independent of underlying network topology characteristics. When this assumption is removed, mobility becomes an interesting and critical factor. Even if we assume that a protocol like MIPv6 handles all of the low-level mobility intricacies, the question becomes whether ALM is still effective and at least modestly efficient. While some ALM protocols are capable of detecting network changes, almost all only focus on reorganizing the overlay when failures occur. They seem to neglect the potential for node mobility and dynamic network conditions.

D. IP Multicast vs ALM Performance

In a previous study[1] we have shown that, in terms of performance, IP Multicast has a clear advantage over ALM. We performed our simulations using the following metrics:

a) Data throughput.: This is the ratio of total received packets to those that should have been received assuming no loss.

b) Relative Delay Penalty (RDP).: This is the path length of the overlay tree divided by the length of the direct path. The smaller the value, the better the ALM protocol since it means that it more closely matches the performance of IP multicast.

A summary of our results is as follows:

- In terms of throughput, low mobility gives no major advantage to IP multicast. However, as nodes start to increase their speed, ALM experiences additional packet loss. At its worst, ALM suffers about 4 times the loss of IP multicast.
- In terms of RDP, low mobility causes IP multicast to perform much better: on the order of 4 to 5 times better than ALM. When mobility is high, IP multicast still performs better, but the improvement is less: an RDP ratio of 2 to 1.

Overall, ALM suffers both when mobility is low and when it is high. Low mobility gives better robustness but very high RDP. High mobility gives better RDP values, but robustness is poor.

E. Motivation for a Hybrid Scheme

In terms of network performance over MIPv6, IP Multicast holds a considerable advantage. However, given deployment concerns, a ubiquitous native multicast deployment might not be possible. Therefore, we need to use a solution that takes advantage of IP Multicast performance, but at the same time simplifies inter-domain deployment. Such a scheme is described in the next section.

III. Hybrid Schemes for MIPv6

Targeting group communication for a wide-scale application, implies that scalability is a primary requirement. Scalability can be expressed both as simplified inter-domain deployment and network efficiency.

As IP Multicast complexities discourage its wide-scale deployment and ALM suffers from mobility, we advocate a hybrid solution distinguishing between inter-domain and intra-domain operation. Given the inter-domain issues of IP Multicast, we expect ALM to handle this space and leave intra-domain to either a separate, smaller scale ALM protocol, or to locally supported IP Multicast. However, as support for IP Multicast may not be available, we also suggest a third alternative, called *DM Tunneling*.

In the remaining of this section we analyze the characteristics of each of the three proposals. We evaluate our suggestions through a number of simulations as these are shown in the next section.

A. ALM and ALM

In this scheme there is an overlay tree between domains and then a different overlay tree for each intra-domain communication. A number of members are expected to act as the anchor points between the two overlays. Althoug this approach has been suggested before (OMNI[5] and Overcast[6]), mobility has never been considered.

Being independent of IP Multicast is the key advantage of these solutions. They can be applied to any domain without assuming network multicast support. Moreover, as the mobility of nodes is abstracted from inter-domain operation, it is expected that performance concerns will be reduced.

Nevertheless, as our previous evaluation has shown, problems may well arise even in the smaller-scale intra-domain environments. This is because even a small percentage of mobile nodes may have implications for the whole overlay tree. Figure 1 gives such an example. The topology on the left shows an overlay tree consisting of the stationary nodes *a*,*b*, *c* and *d*. We may assume that this is an efficient overlay as it is constructed on latency metrics to reflect the underlying network topology. If at this stage a fast moving node *e* simply passes through the domain, then the ALM protocol will reshape the overlay tree based on the latency metrics at a given instant in time. The right topology in Figure 1 shows such a possible reformation. As *e* moves rapidly out of the domain, the overlay will eventually reach its initial state. The outcome is that node *e* did not receive any

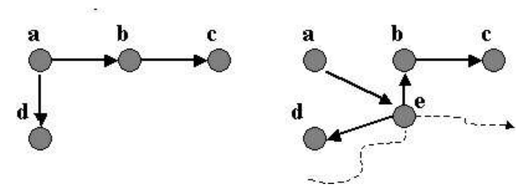

Fig. 1. Effect of a moving node on an overlay tree.

special treatment and that the operation of the whole delivery tree has been considerably disrupted.

We would like to take advantage of native multicast support wherever available. The details of such a scheme are discussed below.

B. ALM and IP Multicast

Assuming native multicast support is present, packets reaching a domain through ALM operation can then be relayed through IP Multicast. A similar approach has already been proposed by the Host Multicast Protocol (HMTP)[7]. It treats domains as *IP Multicast islands* that communicate through a simple overlay tree. The anchor point between ALM and IP Multicast is called a *Designated Member* (DM) and is selected among the local participating members. However, as this work was designed without consideration for mobile nodes, two points need to be clarified.

The first clarification relates to the selection of the Designated Member. Simply selecting one group member from the domain may have serious implications if that node is moving rapidly around or through the domain. Consequently DMs should be selected from a set of stationary or slowest moving nodes. We would also argue that it may be preferable to select the DM from a set of dedicated network nodes. This is the case because a DM has to perform additional functions such as session advertisement (within the domain) and address allocation. Mobile nodes will likely not want this burden. Of course, deploying this solution would now require infrastructure modifications. In our simulations we select the DM by randomly picking a router from each island.

The second clarification is the implementation of IP Multicast itself. As mentioned in 2.1, multicast can be implemented in one of two modes. With Home Subscription, nodes join the multicast group through their Home Agents (which then tunnels packets to the remote location). With Remote Subscription, nodes join the multicast group from their current local router. The benefit of the first option is reduced tree maintenance overhead while the second option results in more efficient data distribution. Because IP Multicast will run only in the “islands”, we would expect the Remote

Subscription to be the preferable choice. Not only is data distribution likely to be more efficient but tree maintenance will be reduced. Assuming our simulations validate our hypothesis, our first choice would be to use Remote Subscription.

C. ALM and DM Tunneling

Finally we want to handle for the case when IP Multicast is not available locally. Although ALM is the obvious alternative, we have seen that even a few mobile nodes may raise serious concerns. By the term *DM Tunneling* we imply a scheme where as mobile nodes enter a new domain, instead of joining the local ALM, they *register* with the local DM. The DM then has the responsibility to act as a Home Agent for each of the registered nodes and tunnel packets (as these are always received through the inter-domain ALM) to each node. In effect, this scheme may be viewed as an intermediate solution between IP Multicast Home and Remote Subscription. In fact, we would expect it to give better performance results (in terms of data delivery hop count) to Home Subscription since the tunneling is now restricted within single domains. However, the choice between DM Tunneling and ALM may depend on the specific characteristics of mobility patterns. The faster the nodes move, the better DM Tunneling will perform. Assuming the simulations prove our hypothesis, we could propose a more adaptive and multi-dimensional system where nodes receive different treatment based on their mobility pattern. Using DM Tunneling for fast nodes and ALM for the stationary/slow ones would limit the impact of fast moving nodes on the overlay tree.

D. Summary

Based on the schemes we have proposed so far, we believe the best alternative will be for a system to use ALM for inter-domain communication and IP Multicast (with Remote Subscription) in the intra-domain. If IP Multicast is not supported, the alternative is to use DM Tunneling. This solution could lead us toward an adaptive system that when IP Multicast is not supported, highly mobile nodes are supported by DM Tunneling and the rest organize themselves into an ALM overlay. The next section evaluates whether our suggested choice does indeed perform best.

IV. Evaluation

In this section we use a simulator to evaluate our proposed set of solutions. We have three specific objectives. First, to demonstrate the potential of a hybrid solution compared to pure ALM and IP Multicast solutions for MIPv6 nodes. Second, because we have shown the problems of ALM with mobile nodes in a previous paper[1], we focus on comparing Home Subscription, Remote Subscription, and DM Tunneling for the intra-domain case. Third, we test the conditions which favor DM Tunneling over pure ALM for the intra-domain operation. We start by presenting the details of our simulation environment.

A. Simulation Configuration

We have performed our simulations using a packet level discrete-event simulator written in Java. Our topologies form power-law graphs generated using Brite. Each of these nodes is mapped to a different radio cell forming a simple one-dimensional radio cell topology. Although we recognize that this cell topology is a potential weakness, we argue that it actually models an unmapped cellular or router-based topology.

For our simulations, the set of parameters, and then the value or range of values used are as follows:

Parameter	Description	Values
N	Number of nodes (routers)	500
R	Number of receivers	10...300
D	Number of domains	20
r	Ratio of mobile receivers	1
t	Experiment time units	10000
br	Transmission bit rate	1/10
link_delay	Link transmission delay	1
h	Handovers per experiment	0...5
Pattern	Node movement	Random

Random is a movement pattern where each node starts from a place different than its Home Agent and randomly chooses the next neighboring cell without any sense of direction.

In terms of protocols, we implemented basic versions of MIPv6 and the proposed hybrid approaches. For network layer multicast, we implemented Source Specific Multicast (SSM). This is because we were interested in sparse-mode operation and we wanted to avoid the complexity of PIM-SM's Rendezvous Point (RP) and shared/shortest path tree switching. This more straightforward approach captures the essence of how multicast routing currently works. Finally, a generic ALM protocol has been implemented by computing a shortest path tree over the complete set of overlay nodes.

Based on this simulation environment, we performed our evaluation on reliability and performance using the following metrics:

c) Data throughput.: This is the ratio of total received packets to those that should have been received assuming no loss.

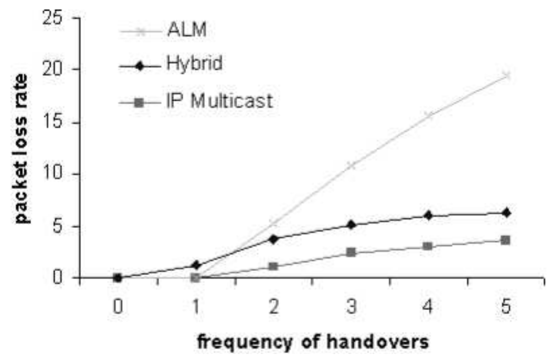

Fig. 2. Loss rate in Hybrid, ALM, and IP Multicast.

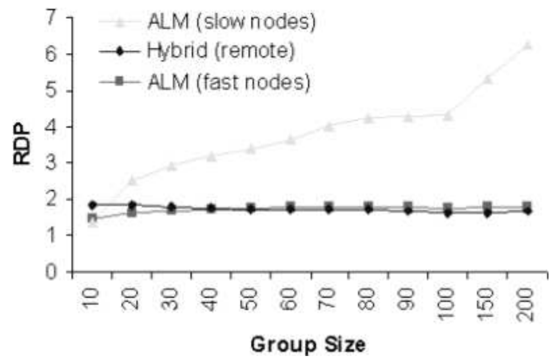

Fig. 3. RDP for Hybrid and ALM schemes.

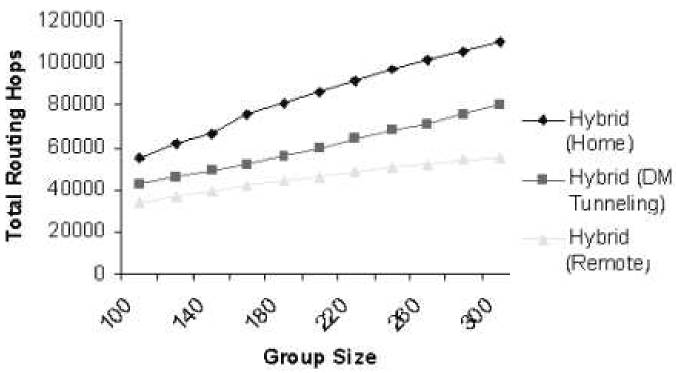

Fig. 4. Hop count for three hybrid schemes.

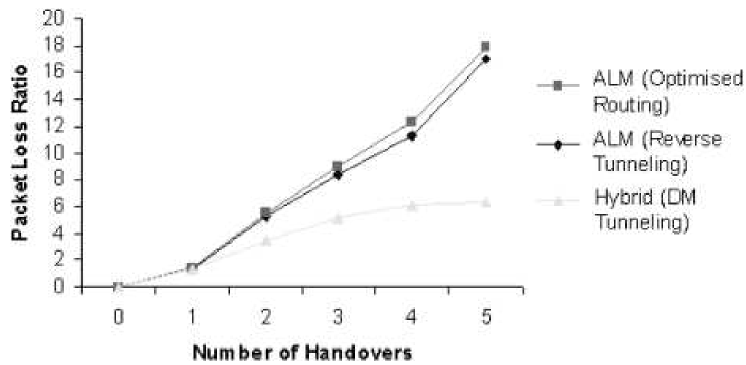

Fig. 5. Loss rate for ALM and DM Tunneling.

d) RDP.: This is the ratio of ALM or hybrid hop count to that of IP Multicast.

Our results are presented in the next two sections. These are followed by a summary section that describes our main findings.

B. Hybrid versus Non-Hybrid

Our results show that both in terms of data throughput and RDP, our hybrid scheme is a better alternative than ALM and compares closely to the performance of IP Multicast.

e) Data Throughput.: We performed a series of tests with a group size of 100 nodes and 0 to 5 handovers per session. Figure 2 shows our results with the x-axis displaying the frequency of handovers and the y-axis the percentage of lost throughput.

There are three main conclusions from these results. First, for slow movement (up to 2 handovers) we see that the hybrid system has similar and in times slightly higher loss rate then ALM. Second, for fast movement (2 handovers and above) ALM suffers greater losses. Extreme cases of 4 to 5 handovers indicate that ALM has around 4 times the drop rate of IP multicast whereas the hybrid approach follows more closely that of pure IP Multicast. The third point is that for the hybrid approach, the drop rate remains relatively stable. This is in contrast to ALM which increases much faster.

f) RDP.: As our previous study showed, the RDP ratio for ALM is different for fast and slow node movement. Figure 3 again shows both cases. However, these results also now show that for the hybrid approach (with Remote Subscription) there is no notable gap. Therefore only a single line is used to represent the hybrid solution. What our results show is that the hybrid solution significantly outperforms ALM for slow nodes. For fast moving nodes the difference is insignificant. However, as previous stated, even though ALM has low RDP for fast moving nodes, it has very high loss. Therefore, we believe that the hybrid approach is preferable to a pure ALM scheme.

C. Comparing Hybrid Schemes

The aim of this section is to verify the preference of IP Multicast Remote Subscription over Home Subscription. We also compare DM Tunneling in order to test its potential. Our tests have measured data throughput and simply the number of routing hops instead of RDP. For data throughput we found that all three schemes give very similar results. Therefore, the line in Figure 2 for the hybrid approach essentially represents the performance for all three schemes. In terms of hop count however, we see important differences. Figure 4 shows that the hybrid implementation with Remote Subscription outperforms the other two options. Moreover, as expected DM Tunneling gives better results than hybrid with Home Subscription. We therefore believe that DM Tunneling is a valid option for a hybrid system. Its comparison to ALM is discussed below.

D. DM Tunneling vs ALM for Intra-Domain

We believe there is a clear advantage to using ALM compared to unicast tunneling when there are stationary nodes. We now try to verify under what conditions, if any, DM Tunneling becomes advantageous. For this reason we performed our tests on single domains with 100 group members. The parameter varied is the number of handovers, and the metric used for comparison is the rate of dropped packets. Figure 5 shows our results.

Figure 5 has three lines. The first two represent the two different approaches in which ALM can be implemented for MIPv6: reversed tunneling and optimized routing. We can see that in terms of data throughput there is no real difference between them. The third line represents DM Tunneling. According to our results, until there are at least two handovers, DM Tunneling performs about the same. Above 2 handovers, DM Tunneling significantly outperforms ALM. This leads us to conclude that treating nodes differently based on their mobility patterns can lead to improved performance.

V. CONCLUSIONS

In this paper we have investigated the potential of a hybrid system in which inter-domain multicast support is provided using ALM and intra-domain support is provided using native multicast. Although similar hybrid systems have been proposed in the past, our contribution is twofold. First, we have examined the impact of mobility on these hybrid protocols, and second, we have proposed a set of necessary modifications. Simulations have shown that a hybrid scheme achieves good performance.

In terms of future work we need to examine how to implement a function that classifies nodes based on their mobility patterns. The use of predictive algorithms is one option. Another interesting issue arises when we consider mobile networks and multicast support. Some islands can be regarded as complete islands that move as a single entity (cars, trains, etc.). Such a characteristic implies that ALM for inter-domain communication will suffer from the same symptoms as any ALM does for mobile nodes. The development and evaluation of a scheme to accommodate these issues is potentially very interesting but left for future work.

REFERENCES

[1] K.Almeroth A.Garyfalos and J.Finney, "A comparison of network and application layer multicast for mipv6 networks," in *Proceedings of MSWIM*, San Diego, CA, September 2003.

[2] C. Perkins D. Johnson and J. Arkk, "Mobility support in ipv6," February 2003, draft-ietf-mobileip-ipv6-21.txt.

[3] A. Joseph B. Zhao and J. Kubiatowicz, "Supporting rapid mobility via locality in an overlay network," Tech. Rep. UCB/CSD-02-1216, University of California, Berkeley, November 2002.

[4] H. Siao and C. King, "Bristle: A mobile structured peer-to-peer architecture," in *Proceedings of IPDPS*, Nice, France, April 2003.

[5] S. Shi and J. Turner, "Routing in overlay multicast networks," in *Proceedings of IEEE Infocom*, New York, June 2002.

[6] D. Gifford J. Jannotti and K. Johnson, "Overcast: Reliable multicasting with an overlay network," in *Proceedings of the USENIX*, San Diego, CA, October 2000.

[7] Z. Jamin and L. Zhang, "Host multicast: A framework for delivering multicast to end users," in *Proceedings of the IEEE Infocom*, New York, June 2002.