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

A.Garyfalos
BTexact Technologies
BT Adastral Park
Ipswich, UK
gary@jungle.bt.co.uk

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 compare Network (IP multicast) and Application Layer Multicast (ALM) under a specific assumption: end hosts are wireless devices using the Mobile IPv6 (MIPv6) protocol. This comparison has three main goals. First, we analyze the implications of running multicast in a mobile, wireless network using Mobile IPv6 (MIPv6). Second, we run a number of simulations to verify whether the network performance issues are different than in wired networks. Finally, using these results, we try to identify the factors that have the most significant impact on performance. Our results indicate that although ALM can be designed to work on top of a wireless network running MIPv6, there are a number of additional performance penalties beyond what occurs in wired networks. Essentially, the advantage of using IP multicast grows even stronger in mobile networks. Nevertheless, we recognize that there are significant barriers to ubiquitous network layer multicast and therefore believe that a more hybrid approach combining both IP multicast and ALM would offer the best performance.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous

General Terms

Measurement,Performance

Keywords

IP Multicast, Application Layer Multicast, Mobile IPv6

1. INTRODUCTION

Multicast for mobile hosts has recently been re-energized as an issue of major importance. 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 network nodes. By adding special functionality in the network, it allows packets to be routed to a specific set of end hosts using fewer network resources. As an alternative, Application Layer Multicast (ALM) follows a different paradigm. Operational responsibility and control is shifted to the application layer and the end hosts themselves. The advantage of this approach is the relative simplicity of deployment since no network modifications are required. In terms of network performance, although previous studies have shown IP multicast to outperform its counterpart[5, 6], ALM has progressed to offer comparable results.

In this paper we compare the two approaches under a specific assumption: end hosts are wireless devices and they operate over an IPv6 network. This comparison has three main goals. First, we analyze the implications of running multicast in a mobile, wireless network using Mobile IPv6 (MIPv6). Second, we run a number of simulations to verify whether the network performance issues are different than in wired networks. Finally, using these results, we try to identify the factors that have the most significant impact on performance.

Our results show that mobility introduces several new challenges for ALM that do not exist in wired networks. System stability is much more dynamic since instead of “node failure”, which is relatively uncommon in wired networks, nodes can move out of range in a mobile network. Furthermore, heterogeneity and capability become critical problems since mobile nodes may be less capable or more constrained in their ability to act as ALM end hosts. In terms of network performance, not only does the gap between IP multicast and ALM still holds, but the overall applicability of ALM is in serious doubt since ALM pays additional performance penalties for highly mobile nodes. However, given deployment concerns, a ubiquitous native multicast deployment might not be possible. Therefore, we advocate a hybrid system in which inter-domain multicast support is provided using ALM and intra-domain support is provided using native multicast.

The remainder of this paper is organized as follows. First we describe the motivation for our work. Section 3 gives an overview of MIPv6, how it supports IP multicast, and a review of current ALM protocols. Our analysis is described in Section 4, followed by the simulation results in Section 5. Section 6 concludes the paper.

MSWiM'03, September 19, 2003, San Diego, California, USA.

2. MOTIVATION

There are a wide range of applications that require wireless multicast support. These include military command and control, distance education and mobile commerce. Even interactive games and streaming video on mobile handsets can be considered. Multicast for mobile hosts is no longer regarded as only an optional service, but as the means to accelerate the transformation of mobile devices into a powerful marketing medium.

The impact of mobility on IP multicast has been studied by a considerable number of recent papers. Several specific problems have been identified and numerous solutions have been proposed [1, 2, 3]. 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, must decide on which modifications to apply.

At the same time Application Layer Multicast (ALM) has appeared as a promising alternative. Overall, ALM is an attempt to overcome the complexity of IP multicast by sacrificing a portion of the network efficiency gains. 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. One notable exception[11] shows how an overlay network could support mobile nodes. However, since the scheme does not use MIPv6, it is outside the scope of what we consider.

The main reason that ALM protocols disregard node movement is because they claim to be independent of underlying network topology characteristics. As the only concern is limited to detecting network failures, the potential for node mobility is neglected. Even if we assume that a protocol like MIPv6 handles all of the low-level mobility intricacies, the question becomes how we measure the effects of mobility. Equally important is the question on how ALM protocols compare to IP multicast. By providing the answers to these questions, we aim to understand how significant an impact mobility patterns have on performance and whether these differences are significant enough to influence (1) the choice between ALM and IP multicast deployment, and (2) the design of ALM for wireless, mobile networks.

3. BACKGROUND

This section presents an overview of the relevant protocols that we study in this paper. It consists of two parts: first we examine Mobile IPv6 (MIPv6) and how it supports IP multicast. Then we describe basic ALM protocol operation and identify the set of performance metrics.

3.1 Mobile IPv6

MIPv6 is a protocol which allows nodes to remain reachable while moving around in the IPv6 Internet [9]. Mobile nodes may receive packets in one of two ways. In a *reverse-tunneling* a router in the home domain, called the *Home Agent* (HA), 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 HA first) operation.

Multicast operation is influenced by these two modes. In a *remote-subscription* the node joins the 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 HA. Membership messages are tunneled to the HA 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 HA may be significant. In addition the delivery tree from the HA in such circumstances relies on unicast encapsulation from the HA to the mobile node and is therefore bandwidth inefficient compared to native multicast forwarding.

As MIPv6 provides only the basic mechanisms to enable multicast operation for mobile nodes, a set of remaining open issues had to be tackled by other research papers (e.g. the delay a MN may experience due to entering a domain that is not part of the multicast tree). Suggested solutions vary in approaches: 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. In this paper we do not need to take into account any of these ideas as they neither positively nor negatively affect the particular metrics we are using.

3.2 ALM protocols

The main difference between IP multicast and ALM is that group management and packet replication are shifted from IP routers (network layer) to end hosts (application layer). Packets are transmitted through standard unicast messages while replication takes place on the end hosts themselves. With most ALM protocols, the underlying physical topology is completely hidden from the tree creation algorithms. Though, one of the main goals of ALM is to discover as much useful information about the network as possible.

Although not as efficient as IP multicast, ALM aims to eliminate the need for additional support from network routers. Moreover it simplifies a number of other issues such as congestion control, pricing models and protocol interoperability. The main drawbacks of ALM are degradation of efficiency (one-to-many delivery function is achieved with multiple unicast calls) and robustness (due to the dependency of the distribution trees on end hosts). Existing work favors two metrics that help to quantify the network efficiency costs[5, 4, 10]:

Link Stress. This metric is defined per node/link and counts the number of identical packets sent by a node over a particular link. *For IP multicast this is equal to 1.*

Relative Delay Penalty (RDP). This is the path length of the overlay tree divided by the length of the direct path. *For IP multicast this is equal to 1.*

Various ALM algorithms exist with different approaches. However there is a set of metrics against which the all approaches can be compared. They are:

Stability is a measure of how quickly and how long overlay trees can be made to be "mature". When a node joins an overlay network, it is usually placed at a random location. From that point onward, there is an ongoing procedure which re-assigns overlay neighbors so that the overall cost of the tree is minimized. When no more significant gains can be achieved by shifting neighbors, the tree is considered to be mature.

Locality is a measure of a protocol's ability to minimize network attributes such as latency and cost. The goal is to minimize the tradeoff with IP multicast by considering the closeness of neighbors as a factor when constructing the routing table. A more detailed discussion is available in [5].

Network Performance is a direct measure of the inefficiency of using application layer multicast. The two metrics, just described, are Link Stress and Relative Delay Penalty (RDP). Previous efforts have shown that overlay routing based on locality characteristics have a maximum Link Stress of 2.5 to 3.5 and an RDP value of 1.5[5, 4, 10] (again, compared to a value of 1 for native multicast). The same experiments also show that ALM protocols that do not consider locality have RDP values of between 4 and 5. These results are important as they provide a reference point for our analysis.

Heterogeneity is a qualitative measure of the variations that exist in the capabilities of nodes–both in terms of what they can do and what they are expected to do. For example, many overlay algorithms assume that all participating nodes have the same capacity to process messages[6]. Furthermore, significant capability heterogeneity exists in many peer-to-peer populations. Even in wired networks this has an impact on performance since hot spots can cause significant performance degradation.

Robustness is a measure of the likelihood that key elements are likely to fail. Overlay networks are particularly prone to network failures due to their inherent dependency on end hosts acting as both clients and servers. It is widely assumed that end hosts are less robust than core network routers.

Our next objective is to investigate through analysis and simulation the impact of mobility on these metrics.

4. ANALYSIS

Having listed the most important aspects of overlay networks, we are now interested in examining how they may be affected by mobility. Our investigation is a two stage process. We first tackle each of the four issues from a theoretical level: first stability, then locality, network performance, heterogeneity, and finally robustness. In the next section we report on the second stage, a simulation study on whether the effects of mobility are more pronounced for ALM than for IP multicast and how specifically the mobility factors affect performance.

4.1 Stability

Stability is often an important assumption of overlay protocols. Current solutions expect changes in the underlying topology to be only because of network partitions or node failures. This assumption influences two important aspects of their operation: change detection and overlay construction. ALM protocols rely on periodic messages to perform reachability testing. Then based on the results, they continually update overlay routing tables until they reach a perceived optimum performance stage ("mature state"). Host mobility breaks this model since nodes move much faster than ALM designs anticipate. Latency metrics now become less robust since distance between two nodes can change dramatically in an instant if one of the nodes moves to a different branch of the network. Moreover, depending on the mobility pattern, a mature state may never be reached. This poses a significant problem for the use of ALM protocols for mobile nodes.

The only possible exception is if mobility can be restricted to certain domains. The effects would then not be as damaging. However, this would require a more hierarchical approach where each domain would have its own *stationary* overlay node. ALM would then be realized in two levels: first an overlay tree between domains and then a tree for each intra-domain with members. As a result, intra-domain movement would be hidden from the inter-domain overlay. This approach has been adopted by existing designs such as Overcast[8]. HMTP[7] goes even further by suggesting that IP multicast could be used to build the intra-domain tree.

In summary, node mobility breaks a fundamental assumption of ALM protocols with a specific impact on the formation and maintenance of the overlay network. Modifications, such as the 2-level approach, are required to limit the impact of such effects.

4.2 Locality and Network Performance

Because locality and network performance are closely related, we examine the two issues together. The main goal is to investigate how the link stress and RDP metrics are influenced by the movement of end hosts. Each of these two metrics are discussed below.

4.2.1 Relative Delay Penalty

We calculate the Relative Delay Penalty (RDP) as the ratio of the link costs for ALM compared to the link costs for IP multicast:

$$RDP = \frac{ALM\_link\_cost}{IPmulticast\_link\_cost}$$

The smaller the value, the better the ALM protocol since it means that it more closely matches the performance of IP multicast. As already stated, previous simulations have shown a mean RDP value of 1.5 for stationary nodes. The aim of this section is to analyze the effect of mobility.

Breaking RDP into its components results in a four-part equation. To be more specific, if we define **RDP_mob** as the RDP metric for mobile receivers, we have the following formula:

$$RDP\_mob = \begin{cases} \frac{ALM(reverse\_tunnelling)}{IPmulticast(home\_subscription)} \\ \frac{ALM(optimised\_routing)}{IPmulticast(remote\_subscription)} \\ \frac{ALM(reverse\_tunnelling)}{IPmulticast(remote\_subscription)} \\ \frac{ALM(optimised\_routing)}{IPmulticast(home\_subscription)} \end{cases}$$

The formula has four parts because IP multicast can be realized in two ways (home and remote subscription), and unicast (the ALM means of communication) can happen either through the home agent or directly between the participating nodes. We should note that although we consider all four combinations, we expect that only the first two to be of interest. The reason is that if handovers are frequent, it is more likely that reverse-tunneling (for ALM) and home-subscription (for IP multicast) will be used in order to keep signaling costs low.

In order to better understand the formula, we extend it to consider the following parameters:

- R, the number of receivers.
- D, the average path distance in the network. This is defined as the average number of links needed to traverse between two distinct nodes.
- $P(i)$, the probability that a receiver i is away from its home network.

We now estimate the four separate link costs as follows:

IP multicast (home_subscription). When all receivers are located in their home networks, the operation is identical to that of a stationary node. However, if we assume that a node has moved to a remote location (a distance D from its home network), the home agent will forward the multicast packets through tunneling. Since we have defined $P(i)$ as the probability that a node is away from home, and as we have N receivers in the group, we can define the total link cost as follows:

$$Multicast\_cost + (D * P(i) * N)$$

where Multicast_cost is the multicast cost associated with delivery the packet from the source to the home agent. We may note that as probability P gets closer to zero, the link cost approximates that for stationary nodes. This is expected since $P = 0$ implies every node remains in its home network and multicast follows the conventional procedure.

IP multicast (remote_subscription). With a remote subscription, mobility is perceived as nothing more than frequent leaves and joins and for the same group. Therefore, the total link cost is the same as that for standard multicast and can be defined as:

$$Multicast\_cost$$

Of course there is also the the increased signaling costs of this option. Receivers will have to re-join the multicast tree after every single handoff. However, the effect of control traffic on RDP is negligible. A further assumption of this scheme is that all foreign routers are both capable and willing to offer multicast to visiting nodes.

ALM (reverse_tuneling). In this option, all packets first go to the home address of each receiver. Consequently the link cost is initially equal to the standard ALM link cost (say ALM_cost). With a probability $P(i)$ that each receiver i is remotely located, the Home Agent will tunnel the packet to the new location. The main difference from IP multicast (home-subscription) comes after this first step. Since the receiver now has the obligation to forward the packets, it can only do so by reverse-tunneling them through the Home Agent. As a result the tunnel path is taken twice for each node whereas for IP multicast (home-subscription) it is taken only once. The total link cost can thus be described by the following formula:

$$ALM\_cost + 2(D * P(i) * N)$$

ALM (optimized_routing). We assume that bindings between mobile nodes are established as the overlay tree is constructed. This means that as an overlay node becomes aware of a new overlay neighbor, this information is passed down to the IP layer and MIPv6 nodes establish a binding association between them. The first difference from ALM (reversed tunneling) is that packets go directly to the remote location and not through the Home Agent. If the next neighbor is not connected to the same router, a new packet will traverse the network. In simple terms, ALM (optimized-routing) should be regarded as a standard ALM protocol with a loose sense of locality. Simply assuming that because two neighbors were originally identified as "close", will not necessarily hold since this proximity may break after a few moves. We therefore simply define this cost as:

$$ALM\_cost(loose\_proximity)$$

Summarizing our analysis, RDP for mobile nodes now looks like the following:

$$RDP\_mob = \begin{cases} \frac{ALM\_cost+2(D*P(i)*N)}{Multicast\_cost+(D*P(i)*N)} & (1) \\ \frac{ALM\_cost(loose\_proximity)}{Multicast\_cost} & (2) \\ \frac{ALM\_cost+2(D*P(i)*N)}{Multicast\_cost} & (3) \\ \frac{ALM\_cost(loose\_proximity)}{Multicast\_cost+(D*P(i)*N)} & (4) \end{cases}$$

In examining this formula, it is still the case that mobility still favors IP multicast over ALM. From (1) we see that in addition to the conventional performance gap between ALM and IP multicast, the overlay approach has an increased tunneling cost. However, mobility would not cause the gap to become much wider than that of a stationary nodes. The key reason being that there is an additional cost for multicast as well. We might therefore state that when movement is frequent (since this is when (1) is expected to occur), RDP performance remains relatively similar.

Part (2) of the formula will likely be used when movement is less frequent. What it shows is that all of the efforts to preserve locality will lose their effect. Previous papers (e.g. [4]) have shown that the RDP value for an ALM protocol with no locality measures approaches a mean value of 4.

As expected, (3) gives a large advantage to IP multicast. The remaining question is the significance of (4). This component is left to be evaluated by simulation and the results are described in the next section.

Overall, we would anticipate IP multicast to retain its performance advantage over ALM. For lower speeds this gain could reach a ratio of four to five. However, as handovers occur more often, this advantage is reduced.

4.2.2 Link Stress

Again this must be analyzed for all four operation types. Starting with multicast (remote-subscription) we would expect link stress to remain 1. This is because apart from the frequent leaves and joins, multicast distribution ensures that no duplicate packets traverse the same link. For multicast (home-subscription) though the situation is different. Even if link stress remains 1 until packets reach the home network, from that point onward tunneling to the Care-of-address implies that duplicate packets may traverse one or more of the same links. Consequently we anticipate link stress to be greater than 1.

Nevertheless, we still expect multicast link stress to be better than the ALM (reversed-tunneling) option. This is because of the similar problem of duplicate traffic flowing over the reverse tunnel plus the standard ALM link stress. As for ALM (optimal-routing), we leave for simulation a comparison with Multicast (home-subscription).

4.3 Heterogeneity

Mobile nodes vary not only in terms of their process capability but also in terms of their mobility pattern. Compared to stationary nodes, wireless handsets involve extra considerations in terms of battery capabilities and wireless link capacity. In addition, as the mobile node moves between cells or access points, packets are expected to be lost. The faster a node moves, the less able it will be to act as a stable forwarding node.

Although the importance of finding good peers has been recognized, the selection is based on metrics that neglect any of the characteristics of mobile nodes and lightweight devices. Accommodating mobility would require consideration for speed and hence handover rate. An additional consideration is "vertical handoffs". A vertical handoff is when a node changes its network connection even while remaining in the same cell. This might occur, for example, if a node wanted to switch from a WLAN to a GSM connection. This implies that overlay formation will have to handle rapid changes in topology and capability.

4.4 Robustness

Not only ALM, but IP multicast also suffers from packet losses due to mobility. However, there is a distinct difference. The effects are largely restricted to the moving node itself and do not influence other end hosts. On the other hand, as overlay networks form a forwarding chain between the participating nodes, it is possible that a fast moving node may adversely affect the robustness of a large portion of the overlay. This is because loss along the overlay path is additive. If node A has a $x\%$ packet loss rate due to handovers, this loss is propagated downstream to all other receivers. Node B would have at best $x\%$ loss rate (plus $y\%$ due to its own movement), and node C would have at least $(x+y)\%$. The impact of these fast moving nodes is therefore greater if they are closer to the source.

4.5 Conclusions from Analysis

For the sake of clarity we summarize the main points of our analysis. The impact of mobility on ALM protocols can be expressed through four main observations:

- Mobility breaks the fundamental assumption of relative stability in the underlying network. Consequently, formation of an overlay network is based on potentially often-changing metrics. Modifications, such as the 2-level approach would be required to limit the impact of such effects.

- The introduction of mobility makes the comparison of ALM and IP multicast performance less straightforward. Four different operation combinations have to be considered and evaluated. In a system where there is slow node movement, RDP can reach values of 4 and beyond. As movement becomes more frequent, the RDP value is expected to be similar.

- Heterogeneity is an inherent problem for any ALM protocol. Nodes selected to act as tree branches needs to be done with special consideration for limited capabilities such as processing power and bandwidth limitations. In addition, mobility makes the problem worse since it increases the rate at which nodes "fail", i.e. move out of range.

- Robustness has always been a concern for ALM. Mobility again exacerbates the problem because of more frequent node failure/movement.

Overall it is interesting to note that mobility has a significant impact on ALM. The fastest nodes move the worse for stability, robustness and heterogeneity. The only benefit of fast movement node comes in terms of the RDP metric. Due to the nature of MIPv6, RDP is much lower for fast than for slow nodes.

5. EVALUATION

In this section we use a simulator to further compare IP multicast and ALM. First we present the details of our simulation environment. Then, we present our results and analysis.

5.1 Simulation Configuration

This section gives an overview of our simulation environment. After describing the simulation environment, we provide a description of the simulation parameters, information about the protocol implementations, and finally the metrics used.

We have performed our simulations using a packet level discrete-event simulator written in Java. Our topologies form power-law graphs generated with 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 a cellular or router topology that is not mapped.

The parameters used in our simulations are as follows:

Parameter	Description	Value Range
N	Number of nodes (routers)	500
R	Number of receivers	10 ... 200
r	Ratio of mobile receivers	1
t	Experiment period	10000 (time units)
br	Packet transmission rate	1 over 10 time units
$link\_delay$	Link transmission delay	1 time unit
h	Handovers per experiment	0 ... 5
$Pattern$	Movement pattern	[Hop.Random.Trip]

Above, we describe three node movement patterns that we consider. In *Hop*, each time a node moves it connects to a random location in the graph. In *Random*, each node starts from its Home Agent and randomly chooses the next neighboring cell without any sense of direction. Finally, in *Trip*, while each node again starts from the Home Agent, it now moves towards a random router in the topology. *Hop* offers simplicity and generality while also serving as a "worst case" scenario (since the node will almost always by away from its Home Agent with a distance D). *Random* is more realistic with a strong sense of localized movement–trace files show that nodes to remain relatively close to their starting points. Finally, *Trip* contains what we believe is more realistic movement. Most of the tests were run with the *Hop* model. Although this choice may lack realism, it gives a feel for an abstract and extreme scenario. Nevertheless, where appropriate the other models were also used in order to strengthen the reality feature of our simulations.

In our simulator, we implemented each of the needed protocols but with certain simplifications. These simplifications include:

MIPv6. We implemented the basic functionality for Mobile Node, Home Agent, and Correspondent Node operation. The main exceptions are the absence of the Return Routability process (when sending Binding Updates) and the Duplicate Address Detection mechanism. Neither of these simplifications has any real impact on our results.

IP Multicast. We implemented a simple version of Source Specific Multicast. This is because we were interested in sparse-mode operation and wanted to avoid the complexity of the Rendezvous Point (RP) introduced by PIM-SM. This more straightforward approach captures the true essence of how multicast routing currently works.

ALM. A generic ALM protocol has been implemented. In order to capture the most important aspects, we compute a shortest path tree over the complete set of overlay nodes. This computation serves two important points. First, it is generic enough to capture the most relevant aspects of ALM routing. Second, it enforces locality in the strongest possible manner since it ensures that closely located receivers will be neighbors in the overlay topology. An important issue is that we did not examine is control overhead. The main reason is that it is difficult to implement, is ALM protocol specific, and does not have a particular impact on the metrics we study.

Based on this setup, we performed our simulations on reliability and performance using the following metrics:

Data throughput. This is the ratio of total received packets over those that should have been received assuming no losses.

RDP_mob. This is the four part equation presented in the previous section. The costs for IP Multicast and the ALM scheme have been calculated as a simple hop count.

Link stress. In our simulator each packet has an associated ID. Therefore we measured mean and maximum values for link stress by counting the number of identical packet IDs transmitted over each link.

Our simulation results and evaluation are presented in the next three sections. These are followed by a summary section that captures our main findings.

5.2 Robustness

In order to compare the robustness of the two systems, we measured their throughput. We performed a series of tests with a group size of 100 nodes and 0 to 5 handovers per session. Figure 3 shows our results with the x-axis displaying the frequency of handovers and the y-axis the percentage of lost throughput. This was calculated from the total received packets over those that should have been received. The Hop movement model was used.

The results show that although there are no or little packet loss for zero or one handover, there is loss for more handovers. The two IP multicast schemes, home and remote subscription, behave almost identically. Packet loss starts at 1.08% for 2 handovers and ranges to a maximum of 3.6% for 5 handovers per session. The two ALM protocols, reversed-tunneling and optimized-routing, both have much higher loss rates. The reversed-tunneling approach lost from 5.3% (2 handovers) to 19.48% (for 5 handovers). Optimized-routing demonstrated similar behavior with losses from 3.8% to around 20%.

There are three main conclusions from these results.

Slow movement. For up to 2 handovers we see that ALM manages to match IP multicast in low packet loss. This would imply that for slow moving devices (pedestrian model) the two schemes perform about the same.

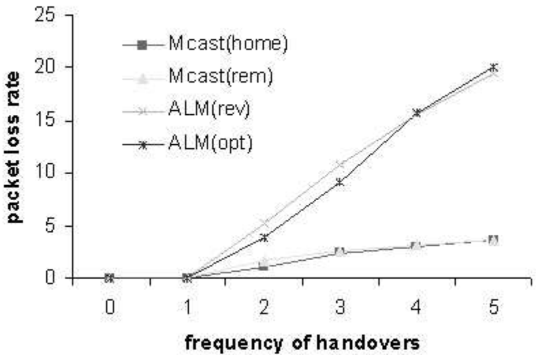

Figure 1: Packet drop rates for ALM and IP multicast.

Fast movement. For above 2 handovers though, ALM suffers greater losses. Extreme cases of 4 to 5 handovers indicate that ALM has around 4 times the drop rate of IP multicast. This suggests that for a fast moving node (e.g. car or train), ALM would not be usable.

Drop rate increase. The third point is that the drop rate is increasing much faster for ALM than for IP multicast. Furthermore, we noticed that as the group size grows larger in IP multicast, "saturation" starts to occur, i.e. there are group members in each cell. This implies that there are more branches of the multicast tree in the network. Consequently, as a mobile node moves to a new domain, there is a greater chance that either the new domain already has a member or a tree branch is not far away. ALM however cannot take advantage of saturation since there is no tree *in* the network. As a result, packet drop rates show a sharp increase.

5.3 RDP-mob

Since *RDP-mob* is a four-part equation, we simulated the four types of operation: IP multicast (home-subscription), IP multicast (remote-subscription), ALM (reversed-tunneling), and ALM (optimized-routing). Figure 4 shows the network efficiency of each operation measured as a total number of routing hops. In these results we used the Random movement model. The x-axis represents the different group sizes while the y-axis shows the RDP ratio, i.e. the number of hops traversed by ALM over those traversed by IP multicast. In order to more accurately measure the RDP ratio, we attempted to minimize any packet losses by considering the top 20% of counted hops. This is because in a real network Binding Updates would be lost and therefore counting the total hops would not be accurate.

There are five lines in the graph. The *stationary* line shows the standard RDP values assuming that nodes are stationary. This is used as a reference point in our discussion. Each of the other lines correspond to each of the parts in the RDP-mob equation. Each of these lines is discussed below.

(1) shows ALM (reverse-tunneling) over IP multicast (home-subscription). As expected, although RDP values are slightly increased, they follow quite closely to those of stationary RDP. Moreover, as the group size increases, RDP-mob grows to a maximum value of 2. Although

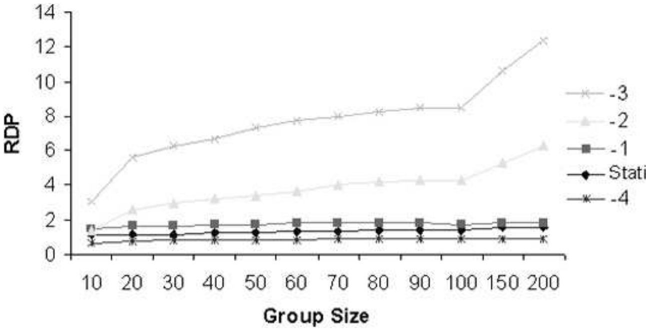

Figure 2: Effect of mobility on RDP.

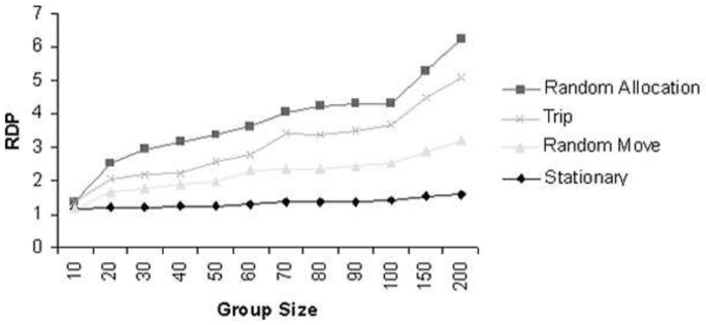

Figure 3: Effect of mobility patterns on RDP.

this is greater than the average of 1.5, the cost, relative to the other results, is not significantly greater.

(2) shows ALM (optimized-routing) over IP multicast (remote-subscription). The situation here is changed since RDP-mob may well reach a value of 6. This is considerable additional cost and is analogous to the link cost for an ALM operation with no consideration for locality information.

(3) shows ALM (reversed-tunneling) over IP multicast (optimized-routing). As expected the performance gains of using network layer multicast are significantly greater than ALM.

(4) shows ALM (optimized-routing) over IP multicast (home-subscription). This is the only scenario where mobility favors ALM over IP multicast. This is an interesting observation and could be an important issue for protocol designers and network administrators. However, this gain is marginalized since it neglects the increased control overhead since in ALM (optimized-routing) every mobile node has to send Binding Updates to their overlay neighbors. If we add the computational burden of the return routability procedure, we may end up heavily loading the network.

Since IP multicast (home-subscription) and ALM (reverse-tunneling) have smaller signaling costs they should be preferred for frequent handovers. Therefore, in an attempt to correlate these results with those for robustness, we associate (1) with fast movement and (2) with slow movement. The outcome is that for highly mobile nodes, the relative benefit of IP multicast is bound by a value of 2. For less mobile nodes, this benefit may reach values of 5 to 6, meaning that IP multicast performs much better for less mobile nodes.

In order to test the impact of different mobility patterns, we performed another series of tests that concentrate on ALM (optimized-routing) and IP multicast (remote-subscription). These tests varied the group size and the movement model. The results are shown in figure 5. We see that the Random model limits the efficiency gain of IP multicast by a considerable factor (RDP is around 3 for 200 nodes). This is almost half the gain of the Hop model (an RDP value of 5 to 6 for the same group size). The Trip model gives an intermediate value of 4 to 5. We can include that the actual movement pattern is an important factor in performance. Moving around a localized area would limit the performance gains of IP multicast by reducing the inefficiency of the ALM protocol.

5.4 Link Stress

Figures 6 and 7 show the maximum and mean values of link stress for remotely located nodes over different group sizes. In the same graphs we have also shown the corresponding values for stationary nodes in order to use them as a reference in our comparison. We again evaluate each of the four parts of the equation. The following are observations for each:

Multicast (remote-subscription). As expected, link stress is always 1.

Multicast (home-subscription). Tunneling from the Home Agent to the end node leads to values greater than 1. Mean values from Figure 7 show a slight increase over the group size: from 1.37 (group size=50) to 1.86 (group size=200). Maximum values for the range of group sizes vary from 5 to 11.

ALM (reversed-tunneling). This is the worst case scenario. Mean values start from 2.08 (group size=50) and range to 2.79 (group size=200). What this means is that for a group of 200 users (over a network of 500 routers), each packet will on average pass almost three times over each link. Such numbers quickly prove overwhelming for content-rich applications such as video streaming or interactive games.

ALM (optimized-routing). This offers better results than reversed-tunneling. Nevertheless, it still falls behind from IP multicast (home-subscription). Mean values range from 1.52 to 2.17.

An interesting observation is that compared to the ALM values for stationary nodes, mobility causes link stress to rise along with the group size. ALM for stationary nodes shows link stress to have a mean value of 1.08 (with a maximum of 3.22), while for mobile nodes the mean rises to the range of 2.08 to 2.79 (reversed-tunneling).

Overall, the ALM approach imposes considerable overhead on the network. With the obvious risk of over-simplifying our results, for an application of 150 to 200 users (over a network of 500 routers) each ALM packet would traverse each link in the distribution tree 1.7 times more than for IP multicast. For stationary nodes the ratio is below 1.1.

5.5 Conclusions from Simulations

The main lesson from our experiments is that, as expected, IP multicast performs better than ALM. However the extent of this advantage is dependent on issues such as speed and locality of movement. In addition, there are a

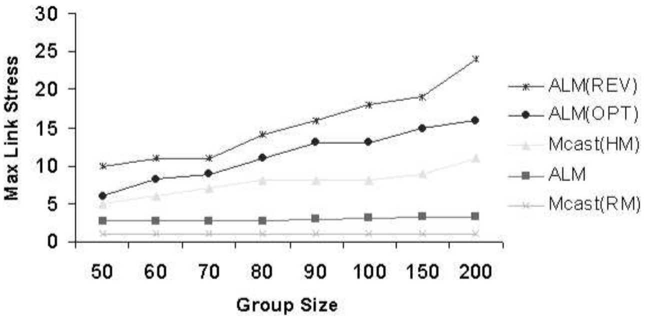

Figure 4: Maximum Link Stress for mobile hosts.

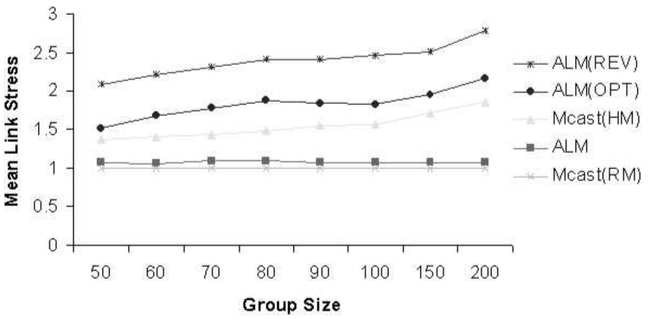

Figure 5: Mean Link Stress for mobile hosts.

number of additional important conclusions we can make. These are:

- In terms of robustness, low mobility gives no major advantage to IP multicast. However, as nodes start to increase their speed, ALM experiences additional packet loss of approximately 4 times that of IP multicast.
- In terms of Relative Delay Penalty, low mobile nodes 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.
- The RDP metric is also heavily dependent on the actual behavior of a user. The more "localized" the movement, the smaller the gain for IP multicast. In actuality, it is not whether a node is highly mobile but rather whether it moves frequently between cells or access points.
- ALM solutions have a considerably higher Link Stress when compared to IP multicast: around 1.7 times for 200 receivers on a 500 router network. And in general, Link Stress increases with group size.

Overall, our concerns about the suitability of ALM for mobile nodes has been confirmed. IP multicast outperforms ALM in all aspects. ALM suffers both when mobility is low and when it is high. High mobility gives better robustness but very high RDP. Low mobility gives better RDP values, but robustness is prohibitive. Link stress seems to be slightly larger when we compare it to a wired networks (1.7 versus 1.1) and essentially independent of mobility.

6. CONCLUSIONS

In this paper we have focused on the performance comparison of IP multicast and Application Layer Multicast when running in a mobile, wireless network. This comparison has been a two-stage process with three main goals. First we have analyzed the implications of MIPv6 on the two different multicast protocols. Second, we have run a number of simulations to analyze whether network performance is different in wired versus wireless networks. Finally, through simulations we investigated what factors most significantly impact the performance.

We believe that ALM can still be deployed in a wireless mobile network without any modifications to the underlying network. This is in contrast to IP multicast which still requires protocol support. Nevertheless, this advantage of ALM is now largely offset by new performance penalties. Our results show that not only does the gap between IP multicast and ALM still hold, but it is questionable whether ALM will even work. We observed that as node mobility grows, robustness decreases to the point where packet loss can be almost 25%. But even when mobility is low, ALM still has significant problems. IP multicast is a much better solution. Finally, with ALM, system stability is much worse and heterogeneity grows to be an even more significant problem.

Nevertheless, as we still recognize that global deployment of native multicast is a problem, a smarter alternative is needed. For example, a compromise solution could be developed which would apply a more hybrid approach. Under such a scheme, native multicast would be used for intra-domain operation while ALM would be used for inter-domain operation. The development and evaluation of such a scheme is left for future work.

7. REFERENCES

[1] C. Tan and S. Pink. Mobicast: A Multicast scheme for wireless networks. *Mobile Networks and Applications(5)*, pp. 259-271, 2000.

[2] T. Harisson, C. Williamson, W. Mackrell and R. Bunt. Mobile Multicast (MoM) Protocol: Multicast support for Mobile Hosts". *Mobicom*, Budapest, Hungary, September 1997.

[3] G. Xylomenos and C. Polyzos. IP Multicasting for Wireless Mobile Hosts. *MILCOM*, San Diego, California, October 1996.

[4] M. Castro, P. Druschel, Y. Hu, and A. Rowstron. Exploiting network proximity in peer-to-peer overlay networks. *Tech. Rep. MSR-TR-2002.*, Microsoft Research, June 2002.

[5] M. Castro, M. Jones, A. Kermarrec, A. Rowstron, M. Theimer, H. Wang, and A. Wolman. An evaluation of scalable application-level multicast built using peer-to-peer overlays. *IEEE Infocom*, San Francisco, April 2003.

[6] S. Jain, R. Mahajan, and D. Wetherall. A study of the performance potential of dht-based overlays. *USITS*, Seattle, Washington, March 2003.

[7] Z. Jamin and L. Zhang. Host multicast: a framework for delivering multicast to end users. *IEEE Infocom*, New York, June 2002.

[8] J. Jannotti, D. Gifford, and K. Johnson. Overcast: Reliable multicasting with an overlay network. *USENIX*, San Diego, California, October 2000.

[9] D. Johnson, C. Perkins, and J. Arkko. Mobility support in ipv6. *draft-ietf-mobileip-ipv6-21.txt*, February 2003.

[10] P. Radoslavov, H. Tangmunarunkit, H. Yu, R. Govindan, S. Shenker, and D. Estrin. On characterizing network topologies and analyzing their impact on protocol design. *Tech. Rep. USC-CS-TR-00-731*, USC, March 2000.

[11] B. Zhao, A. Joseph, and J. Kubiatowicz. Supporting rapid mobility via locality in an overlay network. *Tech.Rep. UCB/CSD-02-1216*, UC-Berkeley, November 2002.