A Flexible Overlay Architecture for Mobile IPv6 Multicast

Anargyros Garyfalos
British Telecom
BT Adastral Park
IPswich, UK
anargyros.garyfalos@bt.com

Kevin C. Almeroth
Department of Computer Science
University of California
Santa Barbara, CA
almeroth@cs.ucsb.edu

***Abstract*— The Internet has evolved from a wired infrastructure to a hybrid of wired and wireless domains. As network access is now provided with much of the last mile being a wireless mobile environment, delivering rich multimedia to users is now a necessity. However, despite the advent of new technology and standards such as Mobile IPv6 (MIPv6), there is still an important dilemma over the choice of systems that either achieve high levels of performance or offer easier deployment. The very deployment of IPv6 is delayed for this reason; network providers continue to use legacy systems. The goal of this paper then is to offer insight into this issue by examining the case of data streaming to MIPv6 users through the use of multicast. By specifically considering the debate over network and application layer multicast, we examine a spectrum of possible alternatives and evaluate the potential of enhancing the functionality of access routers. The result is an overlay architecture that can bring the desired balance between deployment complexity and performance.**

***Index Terms*— Mobile IP, MIPv6, multicast, ALM.**

I. INTRODUCTION

The advent of new technology, such as Mobile IPv6 (MIPv6), often imposes a dilemma to service and network providers: either achieve full compatibility by extending the currently deployed network to meet the new requirements, or follow a more cost-efficient paradigm where performance and efficiency are compromised in order to protect previous investments. Another alternative is sometimes possible: add hardware in a few selected places that operates as a service gateway and performs application layer functionality, e.g. multicast. Frequently related to the principles of the *end-to-end* argument [1], this paper provides insight into this dilemma by focusing on the example of deploying multicast service for users operating in a mobile IPv6 network.

Compared to the one-to-one operation of unicast and the one-to-all of broadcast, *multicast* is a more efficient way of reaching a specific set of network nodes [2]. Implementations of multicast can be broadly realized in two prominent layers: either the network or the application. The first case is most commonly referred to as “IP multicast” and operates by adding special functionality in core Internet routers [2]. The second approach is referred to as Application Layer Multicast (ALM) and follows a paradigm of shifting control to the end hosts themselves [3]. The advantage of IP multicast is its more efficient use of network resources, whereas ALM offers relative deployment simplicity since few or no network modifications are required. For this reason, and in combination with recent advances, ALM has recently attracted much attention in the research community and a number of examples have been the focus of commercial development.

However, the introduction of mobility, and specifically MIPv6, poses new challenges that shift the balance between the two options. In a previous study [4], we have shown that when mobility is introduced, the performance gap between IP multicast and ALM widens considerably. Moreover, the effect of numerous other issues such as network control/management, trust, and node cooperation have two important implications. First, there is a need to re-evaluate the tradeoffs offered by both ALM and IP multicast. Second, as numerous extensions have been proposed to the basic schemes, there is now a wide and complicated spectrum of alternate deployment options. Choosing the right balance of complexity and efficiency is a challenging and multi-dimensional problem. To this end, we have proposed the *Intelligent Gateway Multicast (IGM)* as an overlay architecture aiming to bring a compromise between the two extremes [5]. However, this idea has only been contemplated and has not been thoroughly evaluated.

The contribution of this paper is therefore twofold. First, we describe our investigation of the deployment-versus-performance issue [4], [5]. This summary includes both a theoretical evaluation of the available choices and a representative set of simulation results. Second, based on the resulting evaluation framework, we fully define IGM and provide an extensive evaluation of its operation and deployment suitability.

The remainder of this paper is organized as follows. In Section II we explain the motivation and background for our study and explore related work. Section III analyzes the range of existing alternatives to IP multicast and ALM. In Section IV we describe our proposed IGM solution and evaluate it. The paper is concluded in Section V.

II. BACKGROUND AND RELATED WORK

Current solutions for providing multicast to mobile nodes can be categorized into two groups: native IP multicast and Application Layer Multicast (ALM). This section offers a basic description of each of the two categories and provides a basis for our comparison in the following section.

A. IP Multicast and MIPv6

MIPv6 can accommodate IP multicast operation in two ways. In *home subscription*, the *mobile node* joins the group via a bi-directional tunnel to its *home agent*. Membership messages are tunneled to the home agent, which then joins the group. Data packets that are then received by the home agent are forwarded to the mobile agent via the tunnel. Alternatively, in *remote subscription* the mobile node joins the group directly–via the multicast router to which the mobile node is connected. This second solution may suffer poor performance when the mobile node is moving rapidly and changing *access routers* frequently. In particular, there may be significant *control traffic* because the multicast tree must be frequently changed. This problem can be solved by using the *home subscription* solution. Home subscription can ensure that multicast trees are independent of the mobile node's movement. However, the delay between the multicast group source and then through the home agent and finally the mobile node may be significantly worse than a direct path. Other problems include *tunnel convergence*, which occurs when members of the same group reside in the same domain but are served through different home agents, and *handover performance*, which refers to the packet loss and jitter experienced while the mobile node to re-join the multicast tree after changing access points [6].

In order to overcome the various performance issues, a plethora of routing and/or host extensions have been proposed. Solutions range from *hybrid solutions* [7] to *hierarchical schemes* [8], [9]. Overall, experimental evaluations have shown that each proposal can address a specific set of problems, and yet, every solution is associated with high complexity due to the required configuration of core routers. Therefore, applying any of these techniques in the global Internet is a challenging task.

B. Application Layer Multicast (ALM) and MIPv6

ALM is an attempt to overcome the complexity of native multicast by sacrificing a portion of the network efficiency gains for increased deployability. Compared to IP multicast, ALM has two important differences. First, packets are transmitted through standard unicast messages and replication takes place, not on the routers, but on the end hosts themselves. Second, group membership knowledge is no longer distributed within core routers, as in IP multicast, but is either stored at the end hosts or at other special servers. These two features give ALM two major advantages. First, operation is controlled by the end devices so there is no need for additional support from network routers. Second, it simplifies a number of issues such as congestion control, pricing models and, protocol interoperability. In general, ALM protocols take advantage of the combination of protocol flexibility, application compute power, and the relative simplicity and maturity of unicast technology.

Similar to IP multicast, ALM protocols can accommodate mobility simply based on the operation of MIPv6. Nodes can move while receiving packets through the standard MIPv6 unicast techniques of *reverse tunneling* or *optimized routing*. However, as ALM protocols claim to be independent of underlying network characteristics, their only concern is quickly detecting network failures. The result is that the potential for node mobility is ignored since frequent user movement is not anticipated. Even if we assume that a protocol like MIPv6 handles all low-level mobility intricacies, the question becomes how we measure the impact of mobility on ALM protocol performance.

III. ALM VS IP MULTICAST MIPV6 EFFICIENCY

In this section, we investigate, through analysis, how mobility can affect the network efficiency of both ALM and IP multicast. Our analysis is the first step of a two-stage process. In the next section, we report on a simulation study to evaluate the effects of mobility on ALM protocols and IP multicast. In the remainder of this section, we tackle a number of important issues from a theoretical perspective: stability, heterogeneity, robustness, and network performance.

A. Stability

Stability is a measure of how quickly and for 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 on, there is an ongoing procedure that reassigns 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.

Stability is often an important assumption of overlay protocols. Current solutions expect changes in the underlying topology to only occur 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. Based on the results, they continually update overlay routing tables until they reach a perceived optimum performance stage, called a "mature state". Host mobility breaks this model since nodes move much faster than ALM designs anticipate. ALM protocols now become less robust since the distance between two nodes can change dramatically in a short period of time. 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.

B. Heterogeneity

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 [10]. However, because significant capability heterogeneity exists, even in wired networks, the result is the potential creation of hot spots and significant performance degradation. And while the importance of finding good peers has been recognized, selection criteria often neglect the characteristics of mobile nodes and lightweight devices. This poses another significant problem for the use of ALM protocols among mobile nodes.

C. Robustness

Robustness is a measure of the likelihood that key tree elements will fail. Overlay networks are particularly prone to failures. The greater potential for failure is due to the overlay's inherent dependency on end hosts acting both as clients and as servers. End hosts are thought to be less robust than core network routers.

Both ALM protocols and IP multicast suffer from packet losses due to mobility. However, there is a distinct difference. The effects when running IP multicast are largely restricted to the moving node itself and do not affect other end hosts connected through the mobile node. On the other hand, because overlay networks form a forwarding chain between participating nodes, a fast moving node may adversely affect the robustness of a large portion of the nodes in the overlay. This is because loss along the overlay path is additive. If node A experiences $x\%$ packet loss due to handovers, all downstream receivers will also see $x\%$ packet loss. The performance impact of fast moving nodes grows as the number of downstream receivers increases. This poses another significant problem for the use of ALM protocols for mobile nodes.

D. Network Performance

Network performance is a direct measure of the inefficiency of using application layer multicast. Existing work favors two metrics that help to quantify the costs [11], [12]. The first is *link stress,* defined as the number of identical packets sent by a node over a particular link. *For IP multicast, link stress is always equal to 1.* The second metric is *Relative Delay Penalty (RDP)*. It measures the path length of the overlay tree divided by the length of the direct path. *For IP multicast, RDP is always equal to 1.* Each of these metrics is further described below.

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. Previous evaluations of ALM protocols have shown that overlay routing based on locality characteristics have a maximum RDP value of 1.5 [11], [12]. The same experiments also show that ALM protocols which do not consider locality have RDP values of between 4 and 5.

When we start to consider mobility, four cases are created, each representing a combination of options for ALM and IP multicast. RDP now consists of four parts defined as follows:

$$RDP = \begin{cases} \frac{ALM(reverse\_tunneling)}{IPmulticast(home\_subscription)} \\ \\ \frac{ALM(optimized\_routing)}{IPmulticast(remote\_subscription)} \\ \\ \frac{ALM(reverse\_tunneling)}{IPmulticast(remote\_subscription)} \\ \\ \frac{ALM(optimized\_routing)}{IPmulticast(home\_subscription)} \end{cases}$$

The formula has four parts because IP multicast can be realized in two ways (home and remote subscription), and ALM protocols can operate either through the home agent or directly between the participating nodes. Even though we consider all four combinations, we expect that only the first two will be of interest. The reason is that if handovers are frequent, reverse-tunneling (for ALM) and home-subscription (for IP multicast) will more likely 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, i.e. 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 delivering a packet from the source to the home agent. We note that as $P(i)$ approaches zero, the link cost approximates that of stationary nodes. This result 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 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 cost 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_tunneling). In this option, all packets first go to the home agent of each receiver. Consequently the link cost is initially equal to the standard ALM link cost, (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 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 as:

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

ALM (optimized_routing). We assume that bindings between mobile nodes are established as the overlay tree is constructed. This assumption means that as an overlay node becomes aware of a new overlay neighbor, this information is passed 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 even a short period of movement. We therefore simply define this cost as:

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

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

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

An examination of this formula reveals that mobility still favors IP multicast compared to 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 is that there is an additional cost for multicast as well. We therefore believe 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 (2) shows is that all of the efforts to preserve locality will lose their effect. Previous papers have shown that the RDP value for an ALM protocol with no locality estimate approaches a mean value of 4 [12].

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 the RDP gain could reach a ratio of four to five. However, as handovers occur more often, this advantage is likely to become smaller.

2) Link Stress: As with RDP, link stress 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).

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

IV. ALM VERSUS IP MULTICAST PERFORMANCE

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

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 with Brite[13]. 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 range of parameters used in our simulations are as follows:

Parameter	Description	Value Range
N	Number of nodes (routers)	500
R	Number of receivers	10. . . 200
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	RandomWaypoint

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 (SSM)[14]. This is because we were interested in sparse-mode operation and wanted to avoid the complexity of the Rendezvous Point (RP) introduced by the Protocol Independent Multicast (PIM) protocol[15]. 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.

Based on this setup, we performed our simulations on reliability and performance using three metrics. These metrics are:

1) **Data Throughput.** This is the ratio of total packets received over total packets sent.
2) **RDP.** This is the four part equation presented in Section III.D. The costs for IP multicast and the ALM scheme have been calculated as a simple hop count.
3) **Link Stress.** In our simulator each packet has an associated ID. Therefore, we measure mean and maximum values for link stress by counting the number of identical packet IDs transmitted over each link.

B. Simulation Results

In terms of throughput, low mobility gives no major additional advantage to IP multicast compared to ALM and a no-movement scenario. However, as nodes start to increase their speed, ALM experiences additional packet loss. At its worst, ALM suffers about four times the loss of IP multicast. This result is shown in Figure 1 where the x-axis displays how many handovers occurred on average for each node during the simulation period. The y-axis shows the packet loss percentage. This value was calculated by dividing the total received packets by those that should have been received. Experiments were run for 100 mobile nodes in a network consisting of 200 routers. Both IP multicast schemes considerably outperformed the two ALM schemes.

Turning to RDP, results are straightforward and so graphs are not presented. However, we summarize the results, starting with the low mobility case. When nodes do not frequently move, IP multicast performs much better: on the order of four to five times better than ALM. When mobility is high, IP multicast still performs better but the improvement is less: an RDP ratio of two to one. The reason is that IP multicast suffers a greater performance hit than does ALM. Overall RDP performance is still better for IP multicast, but the gap shrinks as mobility increases.

Finally, results for the link stress metric are shown in Figure 2. The results show that link stress is considerably higher for ALM when compared to IP multicast (around 1.7 times) and increases as group size grows. This result is a direct consequence of the scalability provided by IP multicast and lack of scalability in ALM protocols.

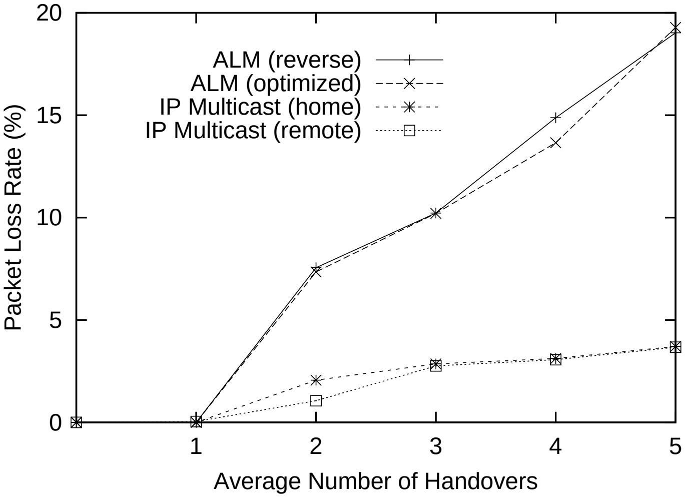

Fig. 1. Packet loss rates for ALM and IP multicast.

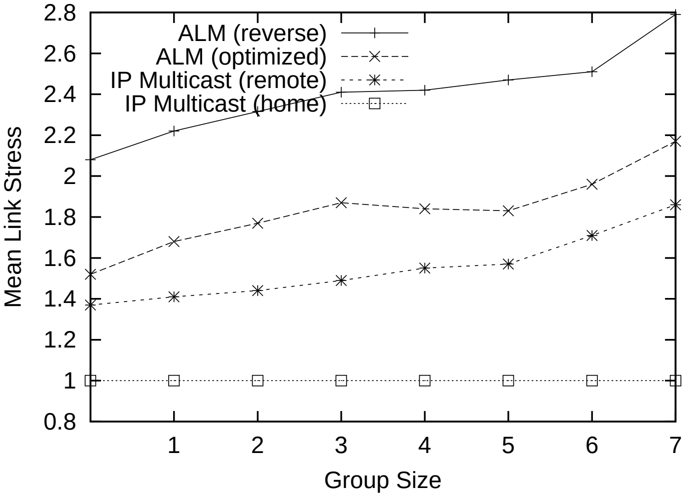

Fig. 2. Mean link stress for ALM and IP multicast.

Overall, ALM suffers both when mobility is low and when it is high. Low mobility gives better robustness but very high RDP. High mobility has poor robustness and also has RDP values that are worse, but comparatively closer to the performance of IP multicast.

Clearly, the different ways of providing ALM-based and IP-multicast-based support for MIPv6 perform differently. In addition, there is also a different level of deployment complexity required. Figure 3 shows how the possible set of services fits into a spectrum. At the one extreme, IP multicast offers the best performance but faces the most significant deployment requirements. On the opposite side, a standard ALM solution is a relatively simpler alternative, but at the same time, the least efficient. The remainder of this paper looks at a hybrid solution that attempts to find a compromise between performance and deployment complexity.

V. INTELLIGENT GATEWAY MULTICAST (IGM)

In this section we present our Intelligent Gateway Multicast (IGM) protocol, an overlay scheme that provides a compromise between the inefficiency of ALM and the deployment complexities of IP multicast. Our description consists of three parts. First, we describe a set of principal design considerations. Second, we describe the operation of IGM, and third, we compare our work to other proposed solutions.

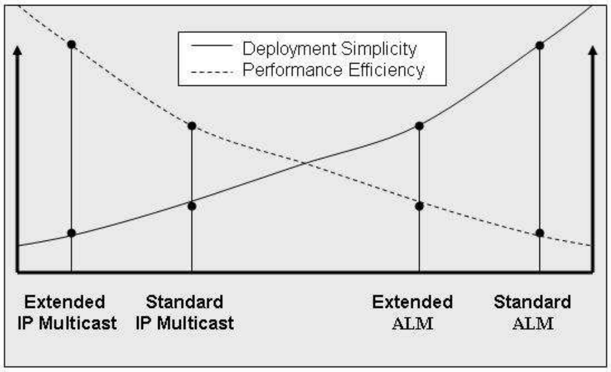

Fig. 3. Spectrum of tradeoffs for ALM and IP multicast.

A. Design Considerations for IGM

There are three main goals for IGM. These include:

1) **Minimize the impact of mobility.** User mobility, as already mentioned, can place significant performance overhead on core and edge network elements. Apart from the performance issues [4], the penalty of expecting routers to frequently modify a multicast tree in response to mobility is a very expensive proposition for both ALM and IP multicast. Therefore, there is a need to hide as much of the effects of node movement as possible. In order to achieve this goal, IGM aims to extend the notion of "device heterogeneity". IGM extends this concept by treating nodes differently based not only on their hardware capabilities, but also on their mobility pattern. In this paper, we have considered speed and movement as relevant factors.
2) **Reduce deployment complexity**. This can largely be realized by avoiding complexity in the core of the network. Configuring core routers is not only a complicated process, but also requires effort by network providers. In addition, as backbone routers are usually heavily loaded, particular emphasis needs to be given to minimizing operating requirements. With IP multicast and the associated costs of mobility being a perfect example, we will avoid the same problem with MIPv6 multicast.
3) **Need for operational control.** One general disadvantage of ALM protocols, is that participating end nodes must be trusted to behave correctly. We believe this assumption should be avoided as much as possible. Because our perspective is one of actually trying to deploy a real-world system, we believe that an Internet Service Provider (ISP) will want some level of operational control [16]. For example, an ISP might not want to provide a streaming data service to a rapidly moving node. Minimizing end-node decision-making is an important requirement.

Considering these three issues, we propose a solution that is based on multicast support in *intelligent gateways*. These devices are not new routers, but devices that are co-located with radio access stations. The complexity of providing a service like multicast is concentrated in these devices. Although the installment of such machines has a certain investment cost, we expect the cost to be considerably lower compared to new functionality in core routers.

B. Overview of IGM

The core element of our IGM architecture is the *gateway*, an intelligent node co-located with the radio access station of any network. Each *gateway* is responsible for a number of operations:

- Advertise their presence and set of services to local mobile nodes [17]. On reception of such advertisements, nodes can then direct *join* messages to the gateway.
- Keep a record of the currently served nodes. Such records will be a simple mapping of the mobile node ID and the corresponding home agent address.
- Relay packets destined to the mobile nodes as they arrive from the home agent or the corresponding node.
- If requested, relay incoming packets to neighboring gateways. This is an attempt to hide the mobility of the node from the home agent by relaying packets instead of reconfiguring the path.
- Exchange information with neighboring gateways in order to acquire node mobility details.
- Act as a firewall and access control point for authorizing user join requests and filtering unwanted traffic [16].

In order to take advantage of the features offered by IGM, mobile nodes must modify their operation in two ways. First, nodes must refrain from sending *binding update* messages when their gateway is IGM compatible. The update will instead be handled by the gateway itself. Second, nodes must inform a new gateway of the following information: the address of the previous gateway and addresses of the node(s) that need to be notified of the mobile node's new address (the home agent and the set of corresponding nodes).

Figure 4 displays the message exchange for an example scenario of IGM. In this example, Gateway 2 is the first gateway visited by the mobile node. The mobile node then moves to Gateway 1. Each of the steps are:

1) When a node enters a new area, it first discovers the associated gateway. It can do this by (1) including the information as part of the handover, (2) broadcasting a service advertisement, or (3) waiting for a periodic broadcast from the gateway. The most efficient of these choices is for the gateway information to be part of the handover process. In Figure 4, we show an example of a mobile node first associating with *Gateway 2*. As part of this communication, the mobile node will discover the gateway's IP address and will receive a list of available services (e.g. QoS provisions). More generally, we believe there will exist a need for intelligent handovers when possibly many wireless companies are offering varied and competing services [18].

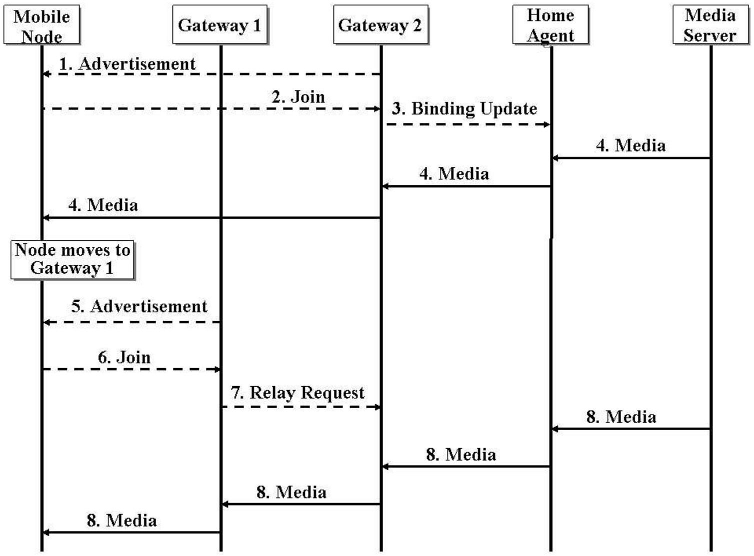

Fig. 4. IGM protocol message exchange.

2) At some time after connecting to a gateway, the mobile node joins a multicast group by sending a *join* message to that gateway. This *join* must contain two IP addresses: one for the *last relay* and one for the *last gateway*. The last relay is valid if there was a previous gateway delivering packets to the mobile node. This address could be the mobile node's home agent or any other correspondent node (if ALM is in operation). The last gateway address is needed so that the current gateway can establish a forwarding tunnel. If this is the node's first gateway (as in this example), the field for this information will be empty.
3) The new gateway enters a mapping of the mobile node ID to the corresponding *last relay* and the *last gateway*. It then sends one of two messages. The default option is an *information request* sent to the *last gateway*. The purpose is twofold: to establish a forwarding tunnel between the two gateways, and to request information concerning the recent mobility pattern of the mobile node (recorded at the last gateway). The alternative is for MIPv6 *binding update* messages to be sent to the home agent and the correspondent nodes (one of which will be the last relay). This message contains the new mobile node's address (care-of address) to which the last relay should be transmitted. Although an *information request* is the default option, since in our example we assume that this is the first gateway the mobile node visits, the *last gateway* option will be empty, and therefore, Gateway 2 sends a binding update to the *last relay*.
4) The outcome of the previous step is for Gateway 2 to indirectly join the multicast group. It then relays incoming packets to the mobile node. The mechanism by which the gateway receives packets can be either ALM or IP multicast.
5) If at a later stage the node moves to a new gateway (e.g. Gateway 1), the process is repeated.
6) Similar to before, after associating with the new gateway, the mobile node will send a *join* message. This time, the *last gateway* address will be Gateway 2.
7) As before, the new gateway enters a record for the mobile node's ID. This time however, since the *last gateway* field is non-empty, it sends an *information request* to the specified address (Gateway 2). Again, this message is to request information about the mobile node's recent mobility pattern and to establish a forwarding tunnel between the two gateways.
8) Upon reception of this message, Gateway 2 must reply with an *information reply*. As before, this message has two purposes. First, t provides locally recorded information to the new gateway. And second, it responds if the gateway is willing to set up a forwarding tunnel. A reason for refusing this setup could be heavy load or support for too many existing tunnels.
9) Assuming a tunnel can be set up, Gateway 1 does not inform the *last relay* of the mobile node's new care-of address. Multicast content still arrives at the *last gateway* and is then tunneled to the recorded care-of address. As a result, the encapsulated packets are first sent to Gateway 2 and then tunneled to Gateway 1. Finally, the content is relayed to the mobile node. Tunnels are kept alive based on a soft state mechanism where links are dropped if a keep-alive message is not received for a specified time period.

When the mobile node visits an area that is not supported by an IGM gateway, it simply switches back to normal operation by sending the binding update messages itself. The operation is then re-initialized as the mobile node sets the *last gateway* field to be empty and the existing forwarding tunnel expires.

Finally, in order to avoid long forwarding chains between the involved gateways, we use a "path check" mechanism. After processing a join message and starting the process of forwarding data packets, the gateway sends a ping message to the packet source. By comparing the hop count (or other quality metrics) of the ping response to the hop count of relayed data packets, the gateway can decide whether it would be more efficient to join the group directly. The threshold for switching to a direct join can be determined in a variety of ways. Here, we have simply described it as a hop count computation.

C. IGM Related Work

Through the use of gateway tunnels, IGM effectively aims to minimize the impact of mobility on multicast operation by hiding part of node movement from the rest of the network. This abstraction is often referred to as *localization*, and is the focus of a set of protocols, called *Localized Mobility Management* (LMM) protocols. LMM protocols address issues such as performance and reliability by minimizing MIPv6 signaling caused by frequent changes in care-of addresses. This is achieved by establishing, "an entity instance of the home domain, similar to the home agent, into a visited network domain hosting the mobile node" [19]. The two main LMM protocols are *Hierarchical MIPv6* and *Fast Handovers for MIPv6*.

Hierarchical Mobile IPv6 (HMIPv6) [20] relies on a wide deployment of specially configured core routers, called *Mobility Anchor Points* (MAPs). MAP nodes divide the global network into specific domains, or *regions*. Each MAP router is then responsible for serving the nodes that reside within its region. Intra-region movement is then hidden from the rest of the network and signaling overhead is reduced. By forming a hierarchy of HMIPv6 routers, the effect of *localization* takes place in different levels of the network depending on the location of the mobile node.

Although very efficient in terms of reducing signaling load [21], [22], the major disadvantage of HMIPv6 is the associated deployment complexity. The requirement for a wide-scale deployment of MAP nodes contradicts our design consideration for reducing deployment complexity.

Fast Handovers for Mobile IPv6 (FMIPv6) [23] aims to minimize handover latency through two mechanisms. First, mobile nodes acquire a new care-of address prior to the completion of movement. This way, they will be able to use the new care-of address right after connecting to the new access router. Second, forwarding tunnels are established between the old access router and the new one during the transition process. As a result packet loss is considerably reduced and operation becomes much smoother.

Although deployment complexity is still of concern for FMIPv6, compared to HMIPv6 it poses fewer concerns. This difference is because configuration does not take place at the core of the network, but only at the edges. Nevertheless, since FMIPv6 aims at a much lower level of localization, it hides less of nodes' mobility.

Based on the specific advantages of each of the two schemes, a combination of the two schemes is possible and at times encouraged [24]. However, as indicated by other studies [19], maintaining compatibility between the two schemes is not a straightforward issue. This is because of two main reasons: the care needed when ordering messages, and the careful configuration required. The outcome of this process not only increases deployment concerns, but also has questionable performance benefits. In certain scenarios, e.g. low packet rate or specific saturation conditions, the generated overhead makes standard MIPv6 the preferred option [24].

IGM overcomes these concerns by placing its intelligence at the edge of the core network. Although we acknowledge that a real life deployment will face initial complexity issues, in the long run, IGM avoids the high cost of HMIPv6 and the compatibility issues of two independent schemes. As a final note, IGM extends the basic hybrid scheme of HMIPv6 and FMIPv6 in the following two ways:

- IGM advocates longer forwarding chains between access routers/gateways. While FMIPv6 restricts handover issues to the mobile node, IGM uses gateway tunnels to save core network resources.
- IGM also aims to take advantage of the intelligence that resides at the edges of the network. As smart antennas and radio transmission technologies advance, gateways will be able to offer new services to mobile nodes. For example, slow moving nodes may receive better service (e.g. QoS) compared to fast nodes because of the smaller impact on the network.

Overall, we believe that given deployment considerations, a solution that exists between the two extreme points of efficiency (IP multicast) and deployment simplicity (ALM) is best. For mobile communication, this means a compromise between the efficient but expensive core-based approach (IP multicast), and the less efficient but simpler alternative (ALM).

VI. EVALUATION OF IGM

This section presents an evaluation of our IGM protocol. First, we analyze how IGM addresses the main design considerations and then we proceed with a simulation of the system. We conclude with a summary of the main points.

A. Analysis of IGM Performance

In our analysis, we follow a two-step process. We begin with a more abstract analysis by investigating how IGM addresses the main design objectives. We then proceed by focusing on network performance issues. From an abstract level, IGM achieves its original three goals as follows:

- **Minimize the impact of mobility.** The impact of mobility is minimized in two ways. First, IGM localizes node movement using tunnels between gateways. Second, IGM can treat nodes differently based on their mobility pattern. Traditionally, functionality on end devices was based on knowledge collected in the core of the network (IP multicast). There was then a trend to use information that is known at the very edge (ALM). With IGM, we create a compromise which takes advantage of both worlds. By using intelligence located at the edge of the core network combined with information acquired from end devices (e.g. node speed), we can now categorize nodes according to parameters deemed most important.
- **Reduce deployment complexity.** Deployment complexity is reduced in two ways. First, IGM is not dependent on the mechanism by which data is delivered to gateways, whether the gateway receives data via IP multicast or some ALM protocol. Second, although IGM uses gateways, no changes to the core of the network are required. In general, we anticipate that use intelligent gateways as a growing trend [17], [18].
- **Need for operational control.** Since we believe that an ISP will want some level of operational control, gateways are designed to provide this functionality. Even if an ALM-style solution is widely deployed, ISPs will still have some level of control since tunnel setup will be processed by their gateways. Filtering and throughput control can be used to allow or impede data delivery.

In Section III, we measured network performance separately for each of the four multicast options: IP multicast with either home or remote subscription and ALM with either optimized routing or reverse tunneling. At this point we analyze the impact of IGM on link cost by again considering these four options:

- **IP multicast (home_subscription).** We have already specified the total link cost to be:

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

where the additional cost (when compared to stationary nodes) is given by the distance, D, from the home agent, $P(i)$ is the probability that the node is away from home, and N is the number of receivers in the group. The difference for IGM is the calculation of D. To be more specific, for non-IGM IP multicast, D is equivalent to the distance from the home agent to the new gateway (nGW). In Figure 5, this path is indicated as (b).

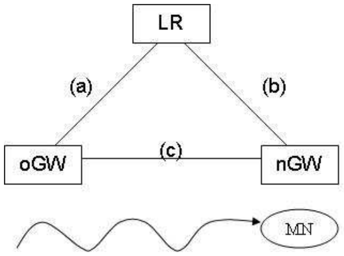

Fig. 5. Tunneled traffic (Paths (a) & (c)) and direct traffic (Path (b)).

In IGM, D has a different value and is also based on time. Initially, when a the mobile node moves to the new gateway, a tunnel is established and D is equivalent to the distance (a) from the home agent to the *old gateway* (oGW) plus the tunnel cost, (c). However, after a specific time interval, *t*, if the new gateway decides that a direct path to the home agent would be more efficient, it switches and the link cost becomes (b). Overall, D for IGM IP multicast with home subscription is:

$$D = \begin{cases} \sum_{i=0}^{\infty}(a+c) & \text{if } b \geq (a+c) \\ \\ \sum_{i=0}^{t}(a+c) + \sum_{i=t+1}^{\infty}(b) & \text{if } b < (a+c) \end{cases} \qquad (1)$$

When $b = (a+c)$, IGM does not provide any gains in link cost. The advantage of IGM occurs when $b > (a+c)$, while the negative effects are restricted to the time when the gateway switches to the direct mode of operation.

- **IP multicast (remote_subscription).** With a remote subscription, mobility is perceived as nothing more than frequent leaves and joins. Therefore, estimating the total link cost is a more complicated procedure since there are no standard reference points, such as the home agent. This complexity occurs because the new gateway that the mobile node joins may already be a member of the multicast distribution tree or just closer to it than the old gateway. Because the opposite is equally possible, we once again simply denote the link cost as follows:

  $$Multicast\_cost$$

  IGM makes no difference in this scenario because the multicast data packets will necessarily flow through the new gateway.
- **ALM (reverse_tunneling).** We have already specified the total link cost to be:

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

  where ALM_cost is the standard ALM cost to the home agent. The difference for IGM is identical to the case for IP Multicast with home subscription. D can therefore be calculated as in Equation (1) with the same implications.
- **ALM (optimized_routing).** In Section III we regarded this case as a standard ALM protocol with a loose sense of locality. We will follow the same approach here:

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

  As in the IP multicast remote subscription case, the use of IGM makes little difference.

Overall, the advantages of IGM are simplified deployment compared to that of IP multicast, and the potential exploitation of information concerning the mobility patterns of the served nodes. In terms of network performance, apart from the expected reduction in packet losses and signaling overhead, link costs may also be reduced. This is particularly true for IP multicast with home subscription and ALM with reverse tunneling. In the next section, we describes the simulation results that validate these hypotheses.

B. Simulation Results

Using the same evaluation methodology described in Section IV, we implemented the IGM protocol by adding IGM intelligence into the access routers. As this functionality conformed to the defined IGM specification, no other change was required for the ALM and IP multicast implementations. Management of gateway tunnels was based on simple ping messages while speed of end users was measured as the rate of handovers over a specific time interval.

In terms of other parameters and metrics, in addition to packet loss, RDP, and link stress, we add as a fourth metric called *control load*. Control load is measured as the total number of routing hops traversed by all *control* (binding update) messages. Related to this metric, we have also evaluated how information about node mobility patterns affects performance. The results for each of the three metrics are described below.

Robustness. In this set of experiments we investigate the extent to which IGM can protect the system from rapidly moving nodes. Depending on whether IP multicast or ALM is deployed, the negative effects are of a different nature. In the IP multicast case, packet losses are restricted to the node that is actually moving. In the ALM scenario however, packet losses are propagated down the ALM forwarding chain. Therefore, we expect IGM to be perform considerably better than ALM.

In order to evaluate our claims we measured the recorded packet loss. We performed a series of tests for all the multicast options, both IGM-enhanced and not. This means that we tested IP multicast (with home and remote subscription) and ALM (with optimized and reverse tunneling) with and without IGM enhancement. With a group size of 100 nodes over a network domain of 500 routers, each mobile node initiated an average of 0 to 5 handovers per session (5 corresponding to vehicular movement). Finally, we used the random waypoint mobility model for all our nodes, but with an important distinction; half of the nodes were originally located in their home network and half at a random remote location. The requirement for this distinction came from our analysis which showed that the distance D can have an important impact on system performance.

Figure 6 shows our results. The x-axis displays the frequency of handovers, and the y-axis shows the percentage of lost packets. In this graph, we have not shown the IP multicast with remote subscription and ALM with optimized routing since they perform similar to IP multicast with home subscription and ALM with reverse tunneling respectively.

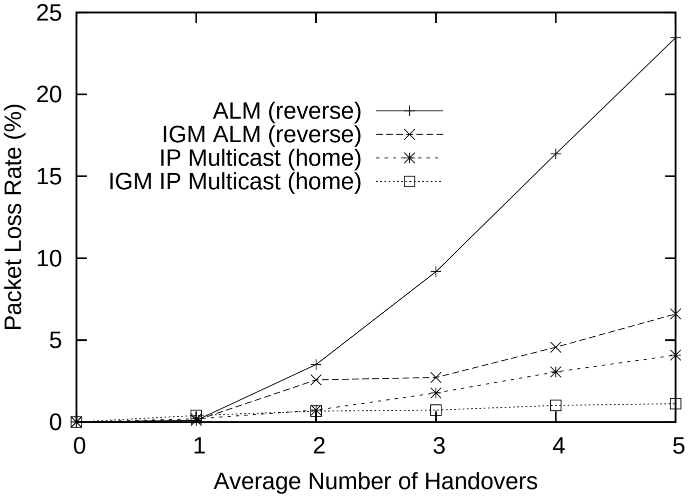

Fig. 6. Packet drop rates for IGM-enhanced ALM and IP multicast as the speed of movement varies.

From the results in Figure 6, we make three observations. The first is that, as expected, low mobility gives no major advantage to IGM for either IP multicast or ALM. Second, IGM gives considerably better values for high mobility. This is more evident in the ALM case with high mobility. IGM reduces the packet loss rate by approximately 70%, from 23% to 6.5%. The gains are of similar magnitude for IP multicast, from 4% to 1%, but should be less noticeable because of the already good performance of IP multicast. The final observation relates to the comparison of ALM and IP multicast. According to Figure 6, a possible deployment of IGM would enhance the operation of an ALM protocol such that network layer multicast would not be required. This is an important observation since multicast support at the network layer has a high deployment cost.

In order to better understand IGM performance, we also examined system throughput over a varying percentage of nodes located close to their home agent. In the previous graph, the ratio was 1:1. Figure 7 has an x-axis that corresponds to the percentage of nodes that move close to their home agent. A standard handover frequency of 5 has been chosen since this imposes the highest demands on the system.

The main observation from Figure 7 is that IGM-enhanced protocols are only slightly affected by the proximity of mobile nodes to their home agents. This result is true for both the IP multicast and the ALM case. On the contrary, the non-IGM options performs poorly when initially there are few modes moving close to their home agent. Only when the percentage of such nodes starts to decrease does the packet loss rate improve, from 26% to approximately 19%.

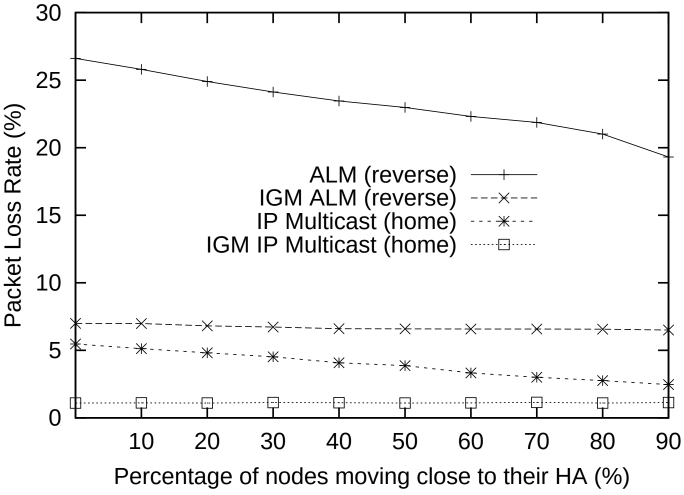

Fig. 7. Packet drop rates for IGM-enhanced ALM and IP multicast as the percentage of nodes that are close to their home agent varies.

Data Load Results. In this set of experiments, we investigate the impact that tunnels impose on the distribution of data packets. More specifically, we focus on evaluating the impact of Equation (1) on the system and how the results differ under varying conditions. To this end, we used two metrics: *total link cost* and *mean link stress*.

Total link cost is measured as the total number of routing hops traversed by all data packets. Using the same parameters as before, we ran a number of tests varying both the handover rates and the percentages of users moving close to their home network. Figure 8 shows the results with the x-axis displaying the percentage of nodes that move close to their home agent and the y-axis showing the total link cost, i.e. the total number of routing hops divided by 10,000.

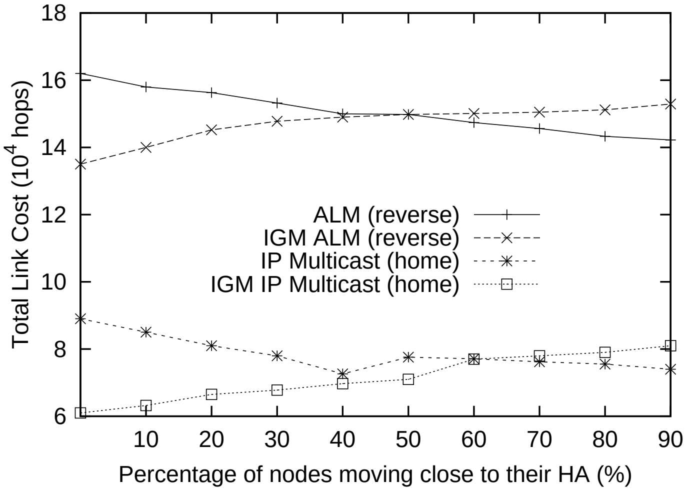

Fig. 8. Total link cost for IGM-enhanced ALM and IP multicast as the percentage of nodes that are close to their home agent varies.

The results show two main points. First, as anticipated, the higher the percentage of nodes that are located close to their home agent, the smaller the gain of the IGM-enhanced solutions. This gain disappears when we reach a balanced distribution of nodes and becomes a disadvantage when most nodes are close to their home agent. Second, even when we have scenarios in which IGM poses additional overhead, this overhead is bounded and indirectly controlled. This result is shown by the fact that even when the vast majority of nodes are close to their home agent (e.g. 90%), the performance gap between IGM and the standard schemes is approximately 10,000 hops while the gains for the opposite scenario, i.e. when 90% of nodes move far away from their home agent, are much higher, approximately 28,000 hops. The reason is that when nodes are close to their home agent, the difference in path length is not significant. But when nodes move far away, IGM switches to a direct path and the normal schemes do not. This difference results in much shorter paths for IGM and so a greater disparity in total link cost.

In addition to the link cost, we also evaluated the link stress in order to measure the distribution of data transmission load. The higher the value for this metric, the higher the number of redundant transmissions over the same set of links. Figure 9 shows our results. The results show that apart from the normal difference between ALM and IP multicast, IGM has no impact on link stress.

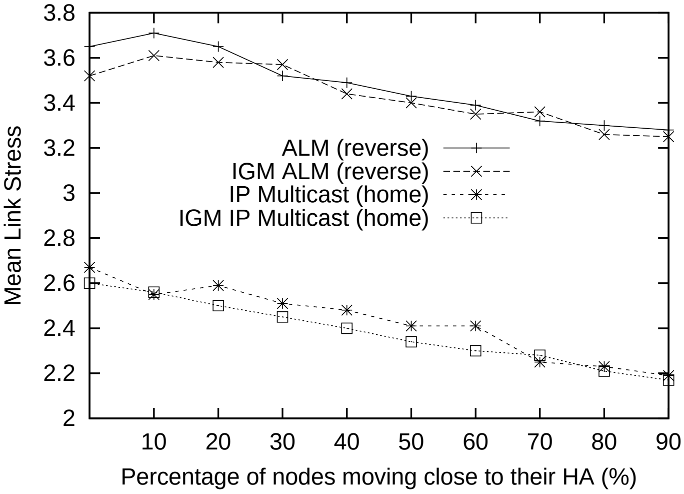

Fig. 9. Mean link stress for IGM-enhanced ALM and IP multicast as the percentage of nodes that are close to their home agent varies.

Overall, IGM is in balance with the standard schemes. It imposes a bounded amount of overhead in the worst case and offers important savings for most scenarios. We have also performed the same tests for various handover rates, but they do not show any unexpected variations.

Control Load Results. We conclude this section by describing the impact of IGM on the control load. We define *control load* as the number of binding update messages transmitted to the network by mobile nodes. Our investigation has two primary goals: to verify that IGM can considerably reduce this load by localizing node movement to inter-gateway tunnels, and to show how information on user mobility patterns can be exploited by an *intelligent* IGM scheme.

In order to test the improvement of an *intelligent* scheme over the basic scheme, we differentiated our set of simulation parameters as follows. Instead of having all nodes move with a speed causing an average rate of 5 handovers per session, we selected a random set of 10% of these nodes to move even faster, causing up to 8 handovers.

For the slower moving nodes, the experiments were straightforward. We measured the total number of router hops that were traversed by binding update messages, both for the IGM-enhanced schemes and for the standard MIPv6 schemes (both ALM and IP multicast). There is no distinction between the different ALM and IP multicast cases since the generation of binding update messages is independent of which solution is used. Our results are calculated as a ratio of the signaling load for standard MIPv6 over that of IGM.

For the faster moving nodes, we used an *intelligent* scheme so that whenever a handover occurred, each gateway recorded the number of joins and leaves. Then, as this information was passed from one gateway to the next, carried in *information reply* messages, each new gateway could classify the incoming mobile nodes according to their recent speed. If the speed was above a specific threshold, corresponding to the 10% of nodes with 8 handovers, the gateway would not send a binding update message to the home agent even if it was a more efficient path than the tunnel. The reasoning being that since the node is expected to move away very soon, there would be little value in attempting to notify the home agent and create a direct path.

Figure 10 shows results for both types of experiments. For the first type, labeled as *IGM*, we see that IGM reduces the signaling load by a ratio of 1.88 to 1.72. This difference offers a considerable amount of savings and is similar to other studies showing similar benefits for Localized Mobility Management (LMM) protocols [22], [24]. The results for the second type are also encouraging since they show an even higher improvement for the *Intelligent IGM*. Compared to the simple IGM scheme, the intelligent alternative increased the amount of binding update message savings by approximately 2%, from 1.88 to 1.92. Despite the fact that the gain disappears as the nodes move closer to their home agents, we have an indication that information at the edge of the network can be used to enhance performance.

VII. CONCLUSIONS

Compared to the one-to-one operation of unicast and the one-to-all of broadcast, *multicast* is a more efficient way of reaching a specific set of network nodes. Implementations of multicast can be broadly realized in two prominent layers: either the network or the application. Although each approach is associated with specific advantages (e.g. performance efficiency for IP multicast and deployment simplicity for ALM), the introduction of mobility introduces several new complexities.

In a previous study [4], we showed that when mobility is introduced, the performance gap between IP multicast and ALM widens considerably. Moreover, the effect of numerous other issues such as network control/management, trust, and node cooperation have two important implications. First, there

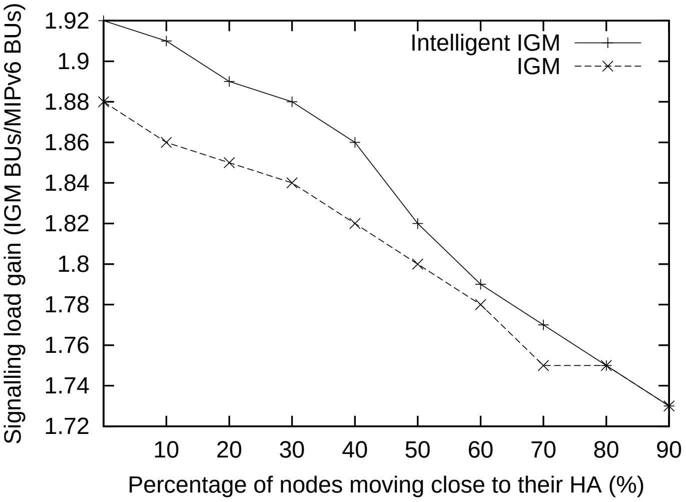

Fig. 10. Signaling load for standard and intelligent IGM as the percentage of nodes that are close to their home agent varies.

is a need to re-evaluate the tradeoffs offered by both ALM and IP multicast. Second, as numerous extensions have been proposed to the basic schemes, we are now faced with a wide and complicated spectrum of alternative deployment options. Choosing the right balance of complexity and efficiency is a challenging and multi-dimensional problem.

The contribution of this paper has been twofold. First, we summarized the most important findings of our previous studies and we have investigated the deployment-versus-performance issue for the set of possible solutions. Second, we have developed the Internet Gateway Multicast (IGM) solution as a compromise between the two extremes. Simulation results show that not only can IGM achieve satisfactory levels of performance, but important advances can be accomplished in the area of robustness and control message overhead. By recognizing the importance of simple deployment requirements, we advocate this solution as part of the general trend in pushing intelligence to the borders between the wired and wireless domains.

REFERENCES

[1] J. Saltzer, D. Reed, and D. Clark, "End-to-end arguments in system design," *ACM Transactions in Computer Systems*, vol. 2, no. 4, November 1984.

[2] K. Almeroth, "The evolution of multicast: From the MBone to inter-domain multicast to Internet2 deployment," *IEEE Network*, vol. 10, no. 1, pp. 10–20, January/February 2000.

[3] S.-W. Tan, G. Waters, and J. Crawford, "A survey and performance evaluation of scalable tree-based application layer multicast protocols," University of Kent, Tech. Rep. 9-03, July 2003.

[4] A. Garyfalos, K. Almeroth, and J. Finney, "A comparison of network and application layer multicast for mobile IPv6 networks," in *ACM Workshop on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, San Diego, CA, September 2003.

[5] A. Garyfalos, K. Almeroth, and K. Sanzgiri, "Deployment complexity versus performance efficiency in mobile multicast," in *International Workshop on Broadband Wireless Multimedia: Algorithms, Architectures and Applications (BroadWiM)*, San Jose, CA, October 2004.

[6] I. Romdhani, et. al., "IP mobile multicast: Challenges and solutions," *IEEE Communications Surveys and Tutorials*, vol. 6, no. 1, First Quarter 2004.

[7] C. Lin and K. Wan, "Mobile multicast support in IP networks," in *IEEE Infocom*, Tel Aviv, Israel, March 2000.

[8] H. Omar, T. Saadawi, and M. Lee, "Multicast support for mobile-IP with the hierarchical local registration approach," in *ACM International Workshop on Wireless Mobile Multimedia (WOWMOM)*, Boston, MA, August 2000.

[9] C. Tan and S. Pink, "Mobicast: A multicast scheme for wireless networks," *Kluwer Mobile Networks and Applications*, vol. 5, no. 4, 2000.

[10] S. Jain, R. Mahajan, and D. Wetherall, "A study of the performance potential of dht-based overlays," in *USENIX Symposium on Internet Technologies and Systems (USITS)*, Seattle, WA, March 2003.

[11] 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," in *IEEE Infocom*, San Francisco, CA, April 2003.

[12] M. Castro, P. Druschel, Y. Hu, and A. Rowston, "Exploiting network proximity in peer-to-peer overlay networks," Microsoft Research, Tech. Rep., June 2002.

[13] A. Medina, A. Lakhina, I. Matta, and J. Byers, "BRITE: An approach to universal topology generation," in *International Workshop on Modeling, Analysis and Simulation of Computer and Telecommunications Systems (MASCOTS)*, Cincinnati, OH, August 2001.

[14] S. Bhattacharyya, "An overview of source-specific multicast (SSM)," Internet Engineering Task Force (IETF), RFC 3569, July 2003.

[15] S. Deering, D. Estrin, D. Farinacci, V. Jacobson, G. Liu, and L. Wei, "PIM architecture for wide-area multicast routing," *IEEE/ACM Transactions on Networking*, pp. 153–162, April 1996.

[16] K. Ramachandran and K. Almeroth, "Mafia: A multicast management solution for access control and traffic filtering," in *IEEE/IFIP Conference on Management of Multimedia Networks and Services (MMNS)*, Belfast, Northern Ireland, September 2003.

[17] R. Chalmers and K. Almeroth, "A mobility gateway for small-device networks," in *IEEE International Conference on Pervasive Computing and Communications (PerCom)*, Orlando, FL, March 2004.

[18] R. Chalmers, G. Krishnamurthi, and K. Almeroth, "Enabling intelligent handovers in heterogeneous networks," *Journal of Wireless Personal Communications*, vol. 29, no. 3, pp. 247–261, June 2004.

[19] T. Pagtzis, C. Williams, P. Kirstein, C. Perkins, and A. Yegin, "Requirements for localised ip mobility management," in *IEEE Wireless Communications and Networking Conference (WCNC)*, New Orleans, LA, March 2003.

[20] H. Soliman, C. Castelluccia, K. Elmalki, and L. Bellier, "Hierarchical mobile ipv6 (hmipv6)," IETF Internet Draft, June 2003.

[21] T. Schmidt and M. Wahlisch, "Performance analysis of multicast mobility in a hierarchical mobile IP proxy environment," in *TERENA networking conference*, Rhodes, Greece, June 2004.

[22] S. Pack and Y. Choi, "A study on perfromance of hierarchical mobile IPv6 in IP-based cellular networks," *IEICE Transactions on Communications*, vol. E87-B, no. 3, March 2004.

[23] G. Tsirtsis, A. Yegin, C. Perkins, G. Dommerty, K. El-Malki, and M. Khalil, "Fast handovers for mobile IPv6," IETF Internet Draft, October 2004.

[24] X. Perez-Cost, M. Torrent-Moreno, and H. Hartenstein, "A performance comparison of mobile IPv6, Hierarchical Mobile IPv6, Fast Handovers for mobile IPv6 and their combination," *ACM Mobile Computing and Communications Review*, vol. 7, no. 4, October 2003.