State Analysis and Aggregation Study for Multicast-based Micro Mobility

Ahmed Helmy
Electrical Engineering – Systems Department
University of Southern California, Los Angeles, CA 90089

***Abstract*- We propose *intra-domain multicast-based mobility* as a solution for IP mobility. Our architecture addresses problems with existing IP mobility proposals; mainly scalability and handoff performance. In our scheme *mobility proxies* are used to allocate per-domain multicast addresses to mobiles for use in micro mobility. State aggregation is studied as an essential element to improve scalability of our approach. We introduce a simple, yet very efficient aggregation algorithm, based on *bit-wise lossy aggregation.* An important result obtained indicates that state tends to be concentrated in less than 20% of the nodes and that our scheme is *extremely efficient* in reducing the state in those nodes. We show that our scheme achieves much higher aggregation gain than conventional *prefix-based* aggregation.**

I. INTRODUCTION

IP mobility addresses the problem of changing network point-of-attachment transparently during movement. Mobile IP (MIP)[4][5] is the current IP mobility standard. However, several studies [1][3][7] have shown that Mobile IP has poor performance during handoff due to communication overhead with the home agent. *Micro-mobility* techniques attempt to improve handoff performance by either using per-domain foreign agents[7][26][27] (or hierarchical approaches) or by using complex caching and forwarding techniques between the previous location and the new location[5][24][25]. In this paper, we introduce a *new multicast-based mobility* scheme for micro-mobility and show that it provides efficient handoff while, at the same time, providing a simpler solution than other micro-mobility approaches.

In multicast-based mobility each mobile node is assigned a multicast address to which it joins through base stations it visits throughout its movement. Handoff is performed using standard join/prune mechanisms. Multicast-based architecture for inter-domain mobility[1] suffers scalability problems concerning multicast state growth with the growth in number of mobile nodes. The architecture also requires ubiquitous multicast deployment and complex security measures. To alleviate these problems, we propose an *intra-domain* multicast-based mobility solution, in which a mobile node is assigned a domain-wide multicast address that it uses for *micro mobility*. The allocated multicast address is locally scoped to a domain. This allows for a domain-wide address allocation scheme, in which a group of *mobility proxies* allocate multicast addresses for visiting mobiles. These addresses are *locally-scoped* and are used temporarily by the mobiles for micro mobility while moving within the domain. Mobile proxies perform inter-domain mobility on behalf of visiting mobiles, then *multicast*-tunnel the packets to the mobile. The multicast address of a mobile does not change while moving within the domain.

The multicast address allocation scheme is performed per-domain. This provides potential for multicast state aggregation opportunities. We thoroughly study and evaluate various multicast aggregation techniques. Our analysis shows that *bit-wise lossy* aggregation achieves aggregation gain much higher than the traditional *prefix-based* aggregation schemes. We observe that multicast state distribution in our case is *non-uniform* among network nodes, and that our scheme achieves *substantial state reduction* for nodes with high state concentration.

The rest of this paper is organized as follows. Section II gives an overview of multicast-based mobility, its promise and its problems. Section III presents our inter-domain multicast-based architecture for micro-mobility. Section IV discusses state aggregation, while Section V presents simulation results and analysis. Related work is discussed in Section VI and Section VII concludes.

II. MULTICAST-BASED MOBILITY

In multicast-based mobility, each mobile node (MN) is assigned a multicast address. The MN, throughout its movement, would join this multicast address through the locations it visits. Nodes wishing to send to the MN send their packets to a *multicast* address, instead of sending their packets to a unicast address. Because the movement will be to a geographical vicinity, it is highly likely that the join from the new location (to which the mobile has recently moved) will traverse a small number of hops to reach the already-established multicast distribution tree. Hence, performance during handoff will be improved drastically. An overview of this architecture is given in Figure 1. As the MN moves, it joins to the assigned multicast address through the new base station. Once the MN starts receiving packets through the new location, it sends a prune message to the old base station to stop the flow of the packets down that path. Thus completing the smooth handoff process.

In[1] we show that multicast-based mobility incurs less than half the handoff delays than does MIP[4], with almost half the network overhead and end-to-end delays.

In spite of such promise, many issues need to be addressed to realize multicast-based mobility. These issues include scalability of multicast state, multicast address allocation, requiring ubiquitous deployment of multicast, and security overhead during handoff.

Scalability of Multicast State. A mobile node is assigned a multicast address to which it joins during its movement. The state created in the routers en-route from MN to the sender is source-group (S,G) state. With growth in number of mobile nodes and number of groups (G), the number of states kept in the routers increases. If there are ‘x’ senders, each sending to ‘y’ MNs, on average, with average path length of ‘l’ then the network routers create ‘$x.y.l$’ (S,G) states. This does not scale.

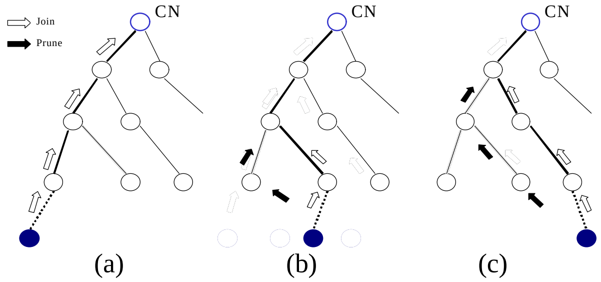

(a) (b) (c)

Figure 1. Multicast-based mobility. As the MN moves, as in (b) and (c), the MN joins the distribution tree through the new location and prunes through the old location.

Multicast Address Allocation. The problem of multicast address allocation[10] is exasperated by requiring each MN to have a globally-unique multicast address. Aside from the fact that the multicast address space is restricted for IPv4, using a global multicast address for each MN may be wasteful, and requiring uniqueness may not be practical.

Ubiquitous Multicast Deployment. To implement inter-domain multicast-based mobility, inter-domain multicast routing needs to be deployed. This requirement restricts the applicability of our inter-domain mobility architecture.

Security Overhead. Mobility setting is prone to remote redirection attacks, where a malicious node redirects to itself packets destined to MN. Authentication should be used with messages revealing information about mobile nodes. Security measures are complex and may incur a lot of overhead. If such measures are invoked with every handoff, they may overshadow benefits of efficient handoff mechanisms.

To alleviate these problems, we propose an intra-domain multicast-based mobility solution, described next.

III. INTRA-DOMAIN ARCHITECTURAL OVERVIEW

In our intra-domain architecture, a mobile node is assigned a multicast address to which it joins while moving. The multicast address, however, is assigned only within a domain and is used for intra-domain micro mobility. While moving between domains, an inter-domain mobility protocol is invoked (e.g., Mobile IP).

When a mobile node moves into a new domain, it contacts the entry point base station (BS). The BS performs per-domain security, then assigns a unicast care-of-address (CoA) for the mobile node to use in that subnet. As in Figure 2, the BS then sends a *request* to a mobility proxy (MP) to obtain a multicast address for the visiting MN.

Upon receiving the request the MP performs inter-domain handoff on behalf of the MN. Also, the MP assigns a multicast address *(G)* for the visiting MN, sends a *reply* message to the BS and keeps record of this mapping. The mapping is used for packet encapsulation later on.

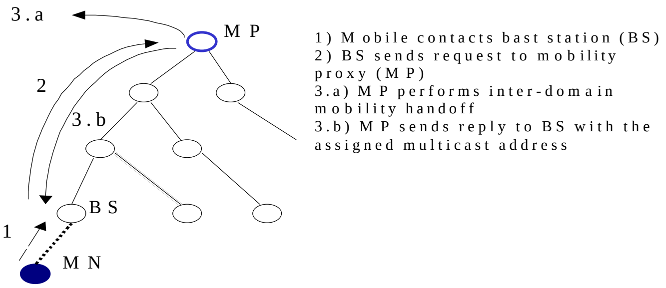

Figure 2. Sequence of actions as the mobile node moves into a domain.

The visiting MN then joins group *G*. The joins are sent to (*MP*,*G*) and are processed as per the underlying multicast routing. The MN moves within the domain using the same locally-scoped multicast address. Handoff is performed using standard join/prune mechanisms and lightweight intra-domain security. When packets are sent to the MN, they are forwarded to the MP using inter-domain mobility. The packets are then encapsulated by the MP, based on the mapping, and multicast to the MN.

Architectural Discussion

In terms of *scalability*, our scheme attempts to address the limitations of the inter-domain multicast-based mobility. In terms of multicast state scalability we note that the multicast state growth is O(*G*) for the architecture presented in this study, as opposed to O(*S*x*G*) in [1][2]. However, there may still be state concentration on certain paths in the network. To further improve scalability of multicast state we investigate several aggregation techniques in the next section. This is quite essential to achieve a scalable solution. *Address allocation* is performed by the mobility proxies on a per-domain basis, the multicast address assignment is now a local mechanism, and the multicast addresses are locally scoped within the domain. This facilitates address allocation and provides per-domain privacy as the multicast packets are not forwarded out of the domain. With regards to *incremental multicast deployment*, our architecture allows for incremental deployment of multicast, based on per-domain approach. This way, the best handoff performance can be attained using our architecture without requiring inter-domain multicast. *Security overhead* during handoff is reduced by using lightweight intra-domain security mechanisms while moving within a domain.

Robustness is crucial to ensure proper operation in the face of crashes and failures. To avoid single-point-of-failure scenarios we provide mechanisms to enhance our protocol robustness. Instead of having only one mobility proxy (MP) per-domain, we propose to have multiple MPs (typically, five to ten per-domain). These MPs are typically placed at the border of the domain. Each MP sends periodic *liveness* messages to a well-known domain-specific group called *MP-group*. All base station routers join this group and receive the liveness messages. Each such router maintains a *live-MP* list and maintains a timer for each MP that is reset by the liveness message from that MP. When a base station router is first contacted by a visiting MN, it performs a hash procedure to select one of the MPs from the MP-list. We use a hash procedure to avoid distributing explicit mapping (which does not scale). The hash procedure assigns a weight to each MP*i* using hash(MN, MP*i*), then selects the highest weight MP to which it sends the request message. This scheme has two advantages. First, it distributes the visiting MNs equally over the MP-list. Second, if MP fails only those MNs that hashed to it are re-hashed, other MNs are not affected. See [31][21] for more detail.

IV . STATE AGGREGATION

The main problem with multicast-based mobility is *scalability of multicast state* with the increase in number of visiting mobile nodes. This is especially a problem where state concentration is expected to occur, as in the mobility proxies. Hence, it is quite crucial to use an effective multicast state aggregation technique to alleviate such a problem.

Most previous work on state aggregation uses *prefix aggregation (PxA)*. That is, two states can be aggregated only if they have the same address prefix. For example, the two addresses 128.125.50.2 and 128.125.50.3 can be aggregated as one entry as 128.125.50.2/31, where 31 is the mask length. This has proven to be efficient for aggregating unicast routing tables in the Internet, since a domain/subnet has a specific unicast prefix. It is not clear, however, if this benefit applies for multicast addresses that are not geographically significant. We propose another kind of aggregation called the *bit-wise aggregation (BA)*. As the name suggests BA works with bits instead of prefixes. For example, 128.125.0.2 and 128.125.1.2 may be aggregated as 128.12.0.2\9, where 9 is the position of the aggregated bit. We perform analysis to understand behavior of these schemes. We define *aggregation ratio (AR)* as the number of states before aggregation (*x*) to the number of states after aggregation (*y*); i.e., *AR=x/y*. Both schemes have identical AR for *in-order* numbers. Figure **3** shows the AR when the numbers are random. The following table presents the results:

	Av. prefix	Av. bit-wise	Av. bit-wise/prefix
80% population	1.40	1.84	1.32
100% population	2.48	1.98	1.19

Note the interesting *cross-over-point at 80% population*. The overall average AR for PxA is '2.48' and for BA is '1.98'. Up to 80% of the population, however, BA outperforms PxA by a factor of 1.32. Hence, we choose bit-wise over prefix aggregation for our scheme.

We further classify multicast aggregation as *perfect (PA)* or *lossy aggregation (LA)*. A multicast state consists of {Src, Grp, iif, oifList}, where iif is the incoming interface and oif is the outgoing interface. Src is the source of the multicast (the MP) and iif points towards the MP. In PA, groups can only be aggregated if the oifList if the same. For LA, however, states are aggregated even though the interfaces may be different. LA achieves better aggregation at the expense of extra network overhead, as the data packets may be sent down an extra link that does not reach a receiver. We study lossy bit-wise (LBA) and perfect bit-wise aggregation (PBA).

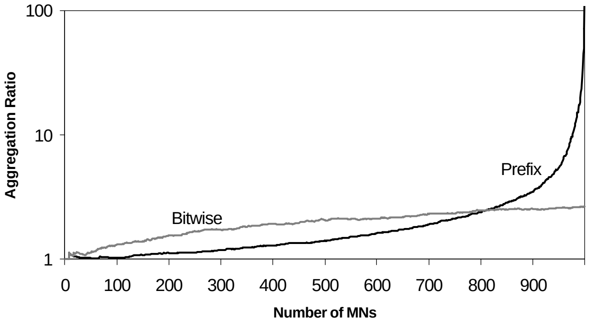

Figure 3. Aggregation ratio for random numbers. Bit-wise aggregation outperforms prefix aggregation up to 80% of the number population..

V. SIMULATION AND ANALYSIS

The first step to solve the scalability problem of multicast state is to understand the state distribution in the routers. We then apply aggregation and analyze the state reduction obtained under the different aggregation techniques. Aggregation gain, in general, depends on several factors, including topology, MP placement, number of MNs, among others. We study and evaluate this problem across different dimensions of various network sizes and number of mobile nodes and mobility proxies.

A. Simulation Setup

We use the network simulator (NS-2)[15] for simulation. Two sets of simulation scenarios were investigated. In the first set, called *dynamic scenarios*, 1000 MNs randomly enter the domain, and move to random nodes within the domain, each time joining through the new location and pruning through the old location, thus capturing the dynamics of the multicast tree. Up to 250k moves were simulated. In the second set of scenarios, called *snapshot scenarios,* MNs enter the domain at random entry nodes and at random times, but they do not move. Thus simulating a snapshot of the domain where nodes may exist at random locations. This approach allows us to scale our simulations to up to 250k MNs. In both simulation scenarios, we use up to 4 mobility proxies placed at backbone nodes. We have simulated several topologies likely to represent intra-domain networks (see Table 1).

Table 1 Simulation Topologies. TS: transit stub, ARPA: arpanet based on real data.

name	nodes	links	av deg	name	nodes	links	av deg
ARPA	47	68	2.89	TS-200	200	372	3.72
TS-100	100	185	3.7	TS-250	250	463	3.72
TS-150	150	276	3.71	TS-300	300	559	3.73

B. Analysis and Results

We first discuss analysis of a topology with 100 nodes and 1 MP. This illustrates our analysis method to understand state distribution and aggregation gains. Then we present results for various topologies and multiple MPs.

i) 100 Nodes with 1 MP: The first topology used for the simulation is that given in Figure 4, with 100 nodes, transit-stub structure, and one mobility proxy (MP) placed at node 0.

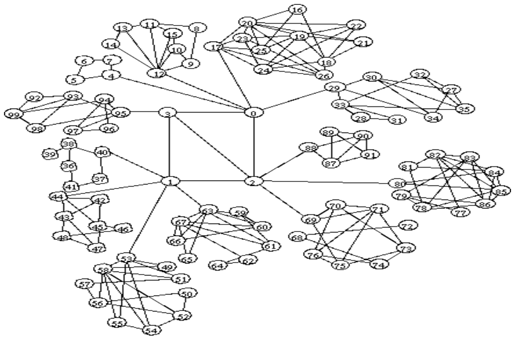

Figure 4. 100 node transit-stub topology (TS-100)

For *dynamic scenarios,* Figure 5 (a) shows the multicast state distribution across the nodes. We notice that much of the multicast state in the network is concentrated at the backbone nodes 0-3. Only 17-20% of the nodes hold more than the average number of states. Also, 40-60% hold less than 1% of the total number of MNs and 66-71% hold less than 2%. That is, we observed a *very high concentration of states in only a small fraction of the nodes*.

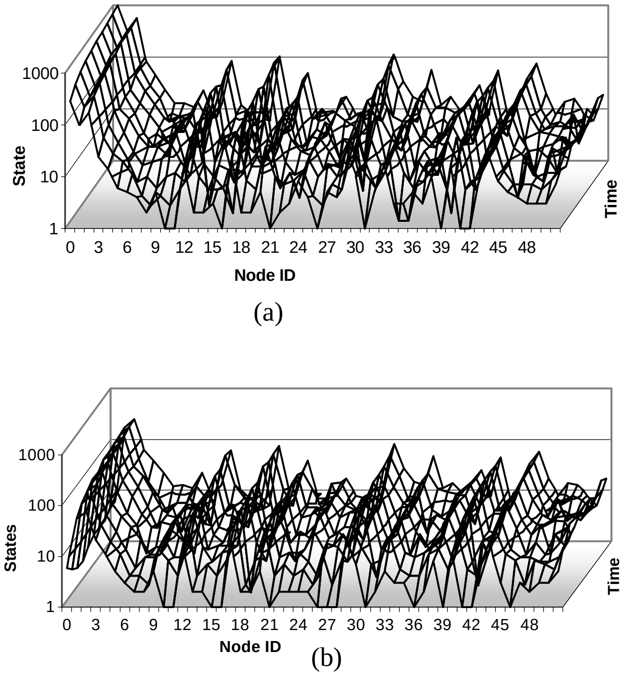

(a)

(b)

Figure 5. State Distribution: (a) without aggregation, (b) with lossy aggregation (data shown for 50 nodes starting from 250 MNs for clarity.)

When applying *lossy bit-wise* aggregation to the above simulations, we obtain the results shown in Figure 5 (b). It is clear that nodes where aggregation is most effective are those nodes 0-3 with maximum state. *The average AR for the 20% of nodes with maximum state was 10.07 (i.e, 90% reduction).*

The overall number of states over the 100 nodes is given in Figure 6. As shown, lossy aggregation obtains good state reduction (factor of 2, or 50% reduction, for average number of states and around 1.5 for 90th percentile). Also, we noticed a significant decrease in *variance* of states across the nodes.

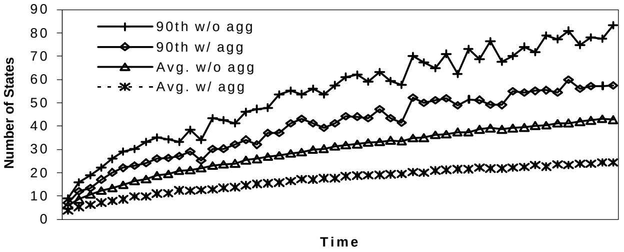

Figure 6. Overall average and 90th percentile.

For *snapshot scenarios*, with 250k MNs, the state distribution across time is given in Figure 7. Again, we see concentration of the state at nodes 0 through 3. We also observe surges in other nodes (the darker areas of the graph).

We examine the state distribution at the last snapshot. The average state per node is 10,830 states. However, only 20% of the nodes had 10k or more states, and around 60% of the nodes have around 2500 states (i.e., 1% of the total number of MNs). This is consistent with our earlier findings and is a strong indication that the state distribution is skewed, with potential for efficient aggregation in nodes with large number of states, where state reduction is mostly needed.

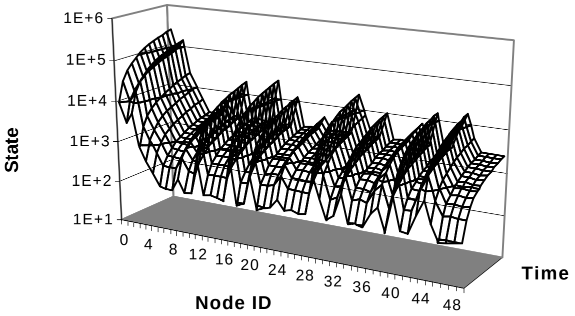

Figure 7. Distribution of state across nodes and time, for 250k MNs (data is shown for 50 nodes and starts from 10k MNs, for clarity).

To further understand the *aggregation* performance, we apply both lossy and perfect aggregation techniques to the snapshot scenarios (up to 40k MNs). For both techniques, we measure the average AR, 90th percentile and maximum state ratios[1]. As shown in Figure 8, these ratios increase with the increase of number of MNs. Also, it is clear that the lossy aggregation achieves better ratios than perfect aggregation. For lossy aggregation the average AR approaches 2 for large number of MNs, whereas for perfect aggregation AR approaches 1.4.

ii) Various Topologies with Multiple MPs: We now investigate lossy and perfect aggregation techniques over several topologies. We also analyze aggregation trends with multiple mobility proxies. We simulated snapshot scenarios with 10k MNs. Figure 9 (a) shows AR results for lossy

---

[1] Max state ratio=Max State Before Aggregation/Max State After Aggregation, and similarly for the 90th percentile ratio.

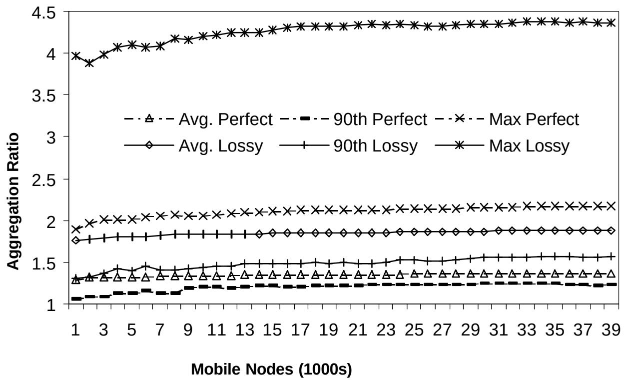

Figure 8. Aggregation ratios for lossy and perfect aggregation techniques.

aggregation. The average AR per node ranges from 1.25 (for 300 nodes 4 MPs) to 1.99 (for 50 nodes with 1 MP). We note several important trends; for the same number of MNs, as the number of nodes in the topology increases, the state concentration in the nodes decreases and the AR decreases. Also, as the number of MPs increases, the concentration of states in the nodes decreases and the AR decreases.

Simulation results for the perfect aggregation are given in Figure 9 (b). The average aggregation ratio ranges from 1.14 (for 300 nodes with 4 MPs) to 1.43 (for 50 nodes with 1 MP). Evidently, lossy aggregation achieves better AR. The trends for both aggregation techniques are quite similar.

VI. RELATED WORK

Several architectures have been proposed to provide IP mobility support. In Mobile IP (MIP) [4] a mobile node (MN) is assigned a *home* address and *home agent* (HA). When the MN moves, it acquires a care-of-address (COA) from a foreign agent (FA). MN *registers* its COA with the HA. Packets destined to MN are sent to HA, then are *tunneled* to MN (*triangle routing). MIPv6* [6] avoids triangle routing by sending *binding updates* to the sender, containing COA of MN. Overhead during handoff, however, renders this scheme unsuitable for *micro* mobility. In[8] a scheme based on dynamic DNS updates is proposed. When MN moves it updates DNS mapping for its host name. This incurs handoff latency due to DNS update delays. Multicast-based mobility (M&M) is proposed in [1][2]. This approach avoids triangle routing and reduces handoff latency. The study in[1] shows superiority of handoff for M&M over Mobile IP protocols. These schemes, however, suffer from issues of scalability of multicast state, address allocation and dependency on inter-domain multicast. We address these issues in our work.

Micro mobility approaches[23] include cellular IP[17] and Hawaii[18]. A domain-gateway registers its address with the HA and forwards the packets to MN. These approaches need special signaling to update mobile-specific routes and require changing unicast routing in all routers. In cellular IP[17], signaling is data-triggered to create paths by having routers snoop on packets. Hawaii[18] proposes a separate routing protocol and requires explicit signaling from the mobiles. These approaches create distribution tree using extra routing entries for the mobile. Our approach builds upon existing multicast mechanisms as opposed to re-creating them. Approaches based on seamless handoff[5][24][25] between old and new access routers, involve fairly complex signaling and buffering procedures. Approaches using hierarchy[7][26][27] employ a gateway per-domain and need to keep a location database to map identifiers into locations. This mapping suffers scalability and robustness problems. Multicast-based mobility incurs less handoff delays and is simpler as it re-uses existing standard multicast mechanisms.

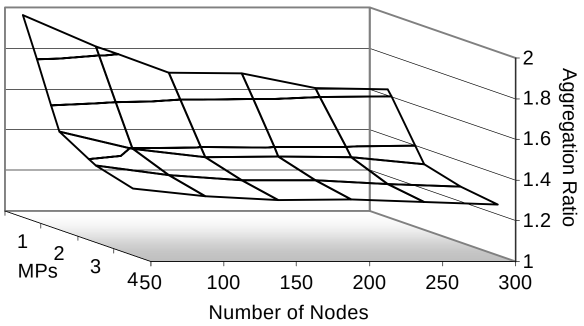

(a) lossy bit-wise aggregation

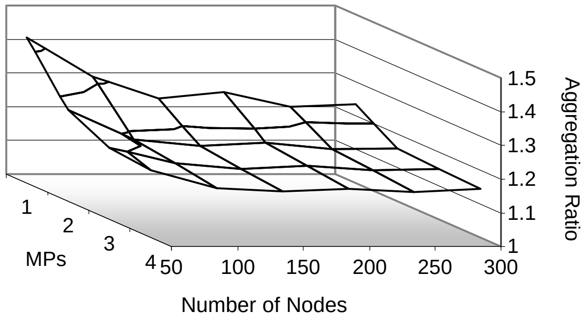

(b) perfect bit-wise aggregation

Figure 9. Aggregation ratio with various topologies and multiple MPs

In the area of multicast state aggregation, [19] proposes an interface-centric model for aggregation. This approach, however, benefits from having a large number of group members, which does not apply in our case. [29] studies strict, pseudo-strict and lossy *prefix* aggregations for wide-area multicast routing. We show that *bit-wise* aggregation usually achieves better gains than prefix aggregation. Our study is the first to address state aggregation for IP mobility and one of very few to study multicast aggregation.

VII. CONCLUSIONS

In this paper, we presented a new *intra-domain multicast-based* protocol for supporting micro mobility. Our scheme uses *mobility proxies* to assign domain-scoped multicast addresses to visiting mobiles. A mobile uses its assigned address during its movement throughout the domain. In our architecture we address serious drawbacks of inter-domain multicast-based mobility approaches. Particularly, we address issues of multicast state *scalability*, multicast address allocation, incremental multicast deployment and overhead of security during handoff.

The main contribution of our paper is the work on *multicast state aggregation*. Unlike previous work, our extensive simulations and analysis show that, for multicast aggregation, *bit-wise* aggregation is a better choice than *prefix* aggregation. Furthermore, we observe that multicast state tends to be distributed *unevenly* across the nodes in the topology. For example, 20% or less of the nodes had more than the average state per node, and up to 60% of the nodes had states/entries less than 1% of the number of MNs. Such state concentration facilitates efficient aggregation.

We have shown through extensive simulation over various topologies and multiple mobility proxies that *bit-wise lossy* aggregation obtains the best aggregation gains. Average aggregation ratios between 1.25 and 2 were obtained in our simulations. This translates into 20% to 50% reduction in multicast state. The average ratio goes up to 10 (i.e., 90% reduction) for the top 20% nodes in state concentration.

Our findings indicate that the aggregation ratio increases with the increase in number of visiting mobile nodes, the decrease in number of mobility proxies, and the decrease in number of nodes in the topology.

We hope that the understanding developed in this paper will help design scalable efficient solutions for IP mobility.

REFERENCES

[1] A. Helmy, "A Multicast-based Protocol for IP Mobility Support", ACM SIGCOMM Second International Workshop on Networked Group Communication (NGC), Palo Alto, November 2000.

[2] J. Mysore, V. Bharghavan, "A New Multicasting-based Architecture for Internet Host Mobility", *Proceedings of ACM MobiCom*, 1997.

[3] S. Seshan, H. Balakrishnan, R. Katz, "Handoffs in Cellular Wireless Networks: The Daedalus Implementation and Experience", *Kluwer Journal on Wireless Networks*, 1995.

[4] C. Perkins, "IP Mobility Support", IETF *RFC 2002*, Oct 96.

[5] C. Perkins and D. Johnson, "Mobility Support in IPv6", *Proceedings of MobiCom'96*, November 1996.

[6] C. Perkins, D. Johnson, "Route Optimization in Mobile IP", *Internet Draft, IETF*, February 2000.

[7] R. Caceres, V. Padmanabhan, "Fast and Scalable Handoffs for Wireless Internetworks" *MobiCom* '96.

[8] A. Snoeren, H. Balakrishnan, "An End-to-End Approach to Host Mobility", *ACM MobiCom*, Aug 00.

[9] D. Estrin, D. Farinacci, A. Helmy, D. Thaler, S. Deering, V. Jacobson, M. Handley, C. Liu, P. Sharma, "Protocol Independent Multicast – Sparse Mode: Protocol Specification", *IETF RFC 2362*, March '98.

[10] S. Kumar, P. Radoslavov, D. Thaler, C. Alaettinoglu, D. Estrin, M. Handley, "The MASC/BGMP Architecture for Inter-domain Multicast Routing", *Proceedings of ACM SIGCOMM*, August 1998.

[11] B. Cain, S. Deering, I. Kouvelas, A. Thyagarajan, "Internet Group Management Protocol, version3", *Internet draft of the IETF*, March 2000.

[12] K. Calvert, M. Doar, E. Zegura, "Modeling Internet Topology", *IEEE Comm*, p. 160-163, June 1997.

[13] G. Phillips, S. Shenker, H. Tangmunarunkit, "Scaling of Multicast Trees: Comments on the Chuang-Sirbu scaling law", *SIGCOMM*'99.

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

[15] L. Breslau, D. Estrin, K. Fall, S. Floyd, J. Heidemann, A. Helmy, P. Huang, S. McCanne, K. Varadhan, Y. Xu, H. Yu, "Advances in Network Simulation", *IEEE Computer*, vol. 33(5), p. 59-67, May 2000.

[16] E. Fleury, Yih Huang, P.K. McKinley. "On the performance and feasibility of multicast core selection heuristics", Int'l Conference on Computer Communications and Networks, Page(s): 296 –303, 1998.

[17] A. Campbell, J. Gomez, S. Kim, A. Valko, C. Wan, Z. Turanyi, "Design, implementation, and evaluation of cellular IP" IEEE Personal Communications , Volume: 7 Issue: 4 , Page(s): 42 –49, Aug. 2000.

[18] R. Ramjee, T. La Porta, L. Salgarelli, S. Thuel, K. Varadhan, L. Li, "IP-based access network infrastructure for next-generation wireless data networks", IEEE Personal Communications , Volume: 7 Issue: 4 , Page(s): 34 –41, Aug. 2000.

[19] D. Thaler, M. Handley, "On the aggregatability of multicast forwarding state", IEEE INFOCOM 2000, Page(s): 1654 -1663 vol.3, 2000.

[20] P. Radoslavov, D. Estrin, R. Govindan, "Exploiting the Bandwidth-Memory Tradeoff in Multicast State Aggregation", USC-CS-TR99-697, Computer Science Department, USC, July 1999.

[21] D. Estrin, M. Handley, A. Helmy, P. Huang, D. Thaler, "A Dynamic Bootstrap Mechanism for Rendezvous-based Multicast Routing", Proceedings of IEEE INFOCOM '99, New York, March 1999.

[22] D. Zappala, "Alternate path routing for multicast", INFOCOM 2000

[23] A. T. Campbell and J. Gomez IP Micro-Mobility Protocols, ACM SIGMOBILE Mobile Computer and Communication Review, 2001

[24] A. O'Neill, G. Tsirtsis, S. Corson, Edge Mobility Architecture, Internet-draft, draft-oneill-ema-02.txt, July 2000.

[25] G. Dommety, et al, "Fast Handovers for Mobile IPv6", Internet Draft, draft-ietf-mobileip-fast-mipv6-02.txt, July 2001.

[26] E. Gustafsson, A. Jonsson, C. Perkins, Mobile IP Regional Registration, draft-ietf-mobileip-reg-tunnel-02, March 2000.

[27] H. Soliman, C. Castelluccia, K. Elmalki, L. Bellier, "Hierarchical Mobile IPv6", draft-ietf-mobileip-hmipv6-04.txt, July 2001.

[28] P. Radoslavov, D. Estrin, R. Govindan, "Exploiting the Bandwidth-Memory Tradeoff in Multicast State Aggregation",USC-CSTR, July99.

[29] J. Kempf, P. Calhoun, C. Pairla, Foreign Agent Assisted Handover, draft-calhoun-mobileip-proactive-fa-01.txt, June 00.

[30] G. Dommety, "Local and Indirect Registration for Anchoring Handoffs",draft-dommety-mobileip-anchor-handoff-01, Dec 00.

[31] A. Helmy, "State Analysis and Aggregation Study for Multicast-based Micro Mobility", USC-CS-TR-02-753, Jan 02.