Hierarchical Level-based IP Multicasting for Tactical Networks

John Lee and Matthew Cheng
Telcordia Technologies (Formerly Bellcore)
331 Newman Springs Road, Red Bank, NJ 07701
{jolee, mcheng} @ research.telcordia.com

***Abstract* - In this paper*, a new IP multicast routing protocol, called "Hierarchical Level-based IP Multicasting" (HLIM), is proposed for tactical networks. Unlike the existing IP multicasting schemes, HLIM can provide a shared and shortest-path multicast tree without using a concentration point like a core or rendezvous point. It supports not only host mobility (movements of IP hosts) but also network mobility (movement of IP routers with/without hosts). The shortest paths between mobile sources and receivers can still be maintained after a host or network movement.**

I. INTRODUCTION

The existing multicasting schemes are basically differentiated by how a tree (called "spanning tree" or "multicast tree") is set up to forward multicast packets from a sender to the receivers within a multicast group. Such schemes can be classified into two main categories: source-based tree (e.g., distance vector multicast routing protocol (DVMRP), multicast open shortest path first (MOSPF), and protocol independent multicast – dense mode (PIM-DM) [1]-[3]) and shared tree (e.g., core based tree (CBT) and protocol independent multicast – sparse mode (PIM-DM) [4]-[6]).

In the source-based tree, a multicast tree is created from each sender to all the receivers of a multicast group. This approach maintains the shortest paths. However, it is less scalable, especially when there are many sources involved in the same multicast group, since it starts with creating a broadcast tree and the tree is not shared among different sources in the same group.

On the other hand, the shared tree approach mitigates the scalability problem by establishing a single multicast tree shared by all sources and receivers of a group. The drawback, however, becomes apparent when optimal routing is emphasized because the shared tree scheme provides the shortest path only between the root of the tree and the receivers, not between the source and the receivers. Also, this approach requires complex mechanisms to advertise the addresses of the roots to every multicast router.

In addition to these shortcomings, existing IP multicasting schemes have no provision for host and network mobility since the schemes have been mainly derived for fixed IP networks. Recently, some ideas for supporting host mobility have been proposed. The basic idea is to extend the current IP multicasting schemes to adapt to mobile environments with the help of mobile IP [7][15]. However, such extension cannot completely resolve the mobility problems incurred by the movements of multicasting hosts. For example, in the source-based approach, the shortest path between a sender and receivers would not be maintained as the sender moves. To date, the network mobility issue has not been addressed at all. Therefore, a new IP multicasting scheme that can support both host and network mobility efficiently is highly desirable.

Our multicasting protocol design, Hierarchical Level-based IP Multicasting (HLIM), can overcome the shortcomings of the conventional IP multicasting schemes and also support both host and network mobility. HLIM assumes that the IP multicast routers are arranged hierarchically as in tactical networks. Each IP multicast router is assigned a level number corresponding to its level in the hierarchy. In addition to the conventional group identifier, a HLIM multicast address carries a scope information, which defines the highest and lowest levels that a multicast packet can reach.

A HLIM multicast tree is established along the shortest paths from the sources to the receivers within the group. It is also shared by all the sources and receivers within the group. As a source or receiver moves, the ongoing multicast session is maintained by extending the tree to the new locations. When a sub-network of routers and hosts move, all the ongoing multicast sessions in the sub-network are maintained. Moreover, only the parent router of the sub-network is aware of the movement and invokes the appropriate operations. The movement is transparent to all other routers and hosts inside the sub-network. After a host or network movement, the extended multicast tree remains shortest-path and shared.

In Section II, the network architecture assumed in our protocol design is described. The basic concept of HLIM and its operations are explained in Section III, IV and V, followed by the conclusions in Section VI.

II. TACTICAL NETWORK ARCHITECTURE

A tactical IP network (TI) is a four-level hierarchical mobile network where IP routers are hierarchically connected as shown in Fig. 1 [9]. Both IP hosts and routers are mobile in the network. The highest level is the brigade level (level 4) and the lowest level is the platoon level (level 1). An IP router at one level (e.g., brigade) interconnects directly with multiple (e.g., 3 to 5) IP routers at the next lower level (e.g., battalion). We refer to such a direct interconnection between one higher-level and multiple lower-level IP routers as a "core" tactical IP subnet. The TI can be pictured as an integration of the "core" subnets in a hierarchical manner.

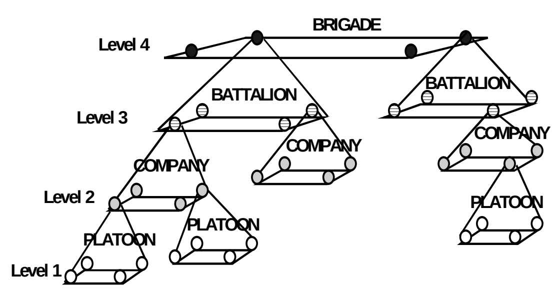

Fig. 1. Tactical IP Network

Fig. 2 illustrates a core tactical IP subnet with local subnets attached at each router. The router at the higher level is labeled H_1, while the four routers at the lower level are labeled L_1 – L_4. Each router (except those located at the lowest or highest level in the hierarchy) has three interfaces: local interface (L_IF), down interface (D_IF), and up interface (U_IF). A router uses its L_IF to handle traffic from and to its local subnets, D_IF for its immediate

* This paper is the result of work done under the Contract DAAB07-98-C-D007 from U.S. Army CECOM.

descendant routers, and U_IF for its immediate ascendant router and its adjacent routers at the same level. In each subnet (core or local), some wireless multiple access technology with broadcast/multicast capabilities is used to interconnect the IP hosts and/or routers.

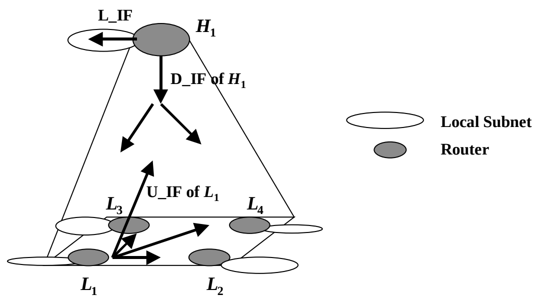

Fig. 2. Three possible interfaces in tactical IP subnets.

Though only one router is shown in a local subnet (or at the higher level in a core subnet) in the previous figures, it is possible to have more than one IP router to handle IP traffic between the local subnet and the ascendant/descendant routers as shown in Fig. 3. In unicast, this would not produce any duplicate messages because one preferred IP router among them is selected by the unicast routing protocol. However in multicasting, if a designated IP router is not assigned in advance, unnecessary duplicate multicast traffic may result. To avoid this situation, the HLIM protocol elects only one of the routers to handle the multicast traffic. The elected router is denoted as a hierarchical designated router (HDR). The HDR which handles multicast traffic for its local subnet is referred as a local HDR. The HDR which handles multicast traffic flowing between core subnets of adjacent levels is referred as a global HDR. The other non-HDR routers within a subnet ignore any multicast control or data packets received. Note that the other non-HDR routers may serve as backups for the elected HDRs, in case that a HDR moves away, has a malfunction, or is destroyed.

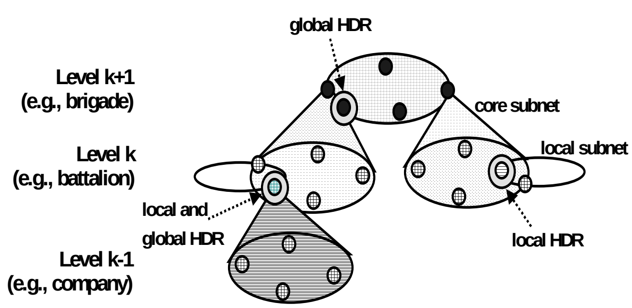

Fig. 3. Concept of hierarchical designated routers (HDR).

III. THE BASIC CONCEPT OF HLIM

One basic element in HLIM is a newly defined scope for multicasting. The main idea is to define a multicast scope region flexibly based on the hierarchical structure of the TI, and utilize such scope to help establish the shortest path and shared multicast tree. A HLIM scope region is uniquely identified by a scope boundary, SB(L,H), and the root identity of the region, "root ID" (RID). The SB(L,H) is comprised of two hierarchical levels, where L sets the lower boundary and H sets the upper boundary that a multicast packet generated within the region can reach. The RID for a given SB(L,H) is the identity of the HDR at level H+1 in that region (denoted as HDR(H+1)). It can be some identifier assigned to the root HDR of the region (which should be unique among all HDRs at the same level of the TI), or the IP address of the HDR.

The HLIM scope region information is included as part of a HLIM group address. A HLIM group address can therefore be denoted as G(L,H,RID,APL_ID), where APL_ID identifies a specific multicast application/service. It implies not only a specific multicast application/service but also its scope boundaries and unique scope region. Consequently, any multicast packet carries its own scope region information and the routers forward the packets only within the defined scope region. These HLIM group address fields may be carried completely within an IPv4 or IPv6 class D address [11], or in a combination of a class D address and some other fields in an IP header (e.g., IP option [16]). It should be noted that the RID is needed only because HLIM supports mobile hosts and allows "handover" of ongoing multicast session during movement. As long as a host remains in its own scope region, the packets to/from the host will not interfere with the packets bearing the same SB(L,H) and APL_ID information at another place. However, when the host moves outside its original scope region, the packets to/from the host will interfere with the packets carrying the same SB(L,H) and APL_ID information at the new place if the RID is not included.

Fig. 4 shows some possible HLIM scope regions. The RID for a given SB(L,H) is denoted as HDR_ID(H+1). For example, the RID for SB(1,3) is HDR_ID(4). Note that there is no RID for SB(L=1, 2, 3, or 4, H=4). In this case, *null* is used as the RID in the group address. Multicast packets associated with, for example, SB(2,3) (the shaded scope region) are not allowed to move beyond the HDRs at level 3 and below the HDRs at level 2 in normal cases. The packets will be forwarded beyond the region only when a mobile user who has already established a multicast session at its original location moves outside the scope region and is willing to continue its on-going multicast session. As shown in the figure, a scope region is defined in a way that the multicast packets from any host are delivered to any other host within the same scope region without traversing other scope regions. Each HDR should keep a list of its ascendant HDR_IDs (e.g., a HDR(2) keeps HDR_ID(3) and HDR_ID(4)). Such list is used by hosts and HDRs to obtain RIDs and determine the appropriate HLIM operations when a mobile host/router moves to this HDR.

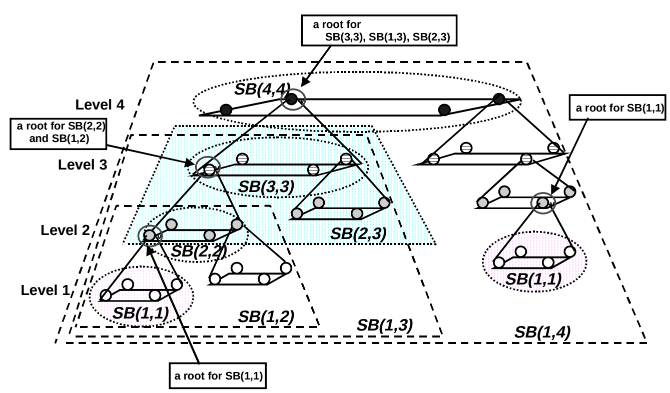

Fig. 4 HLIM scope regions in TI.

In the following two sections, we will describe the HLIM operations for a mobile host/router to establish a multicast session and to handover an on-going session during movement. The HLIM protocol also contains other functions for: 1) autoconfiguration of

D_IF and U_IF for a HLIM router during startup, 2) election of local/global HDRs when necessary (e.g., during startup, or when the old one moves away, malfunctions or is destroyed), 3) dynamic maintenance of an established multicast tree, 4) increased reliability support for the HLIM control messages, and 5) updating the list of ascendant HDR_IDs among the descendant HDRs when there is a change in an ascendant HDR (e.g., when a new one is elected). Due to the length limitation of this paper, the details of these auxiliary functions are omitted.

IV. HLIM REGULAR OPERATIONS

This section describes the HLIM operations for a host/router to start a new multicast session. To facilitate the explanation of the HLIM operations, we assume a HDR is elected to be both a local and global HDR, and its IP address is used as a HDR_ID.

A. Local Host Joining a Group

The current IGMP [11] can be applied to the regular operations for local host joining a multicast group, with enhancements introduced by the HLIM design.

A host at level k obtains a SB(L,H) and APL_ID for initiating a multicast session through some advertising mechanisms such as SAP (session announcement protocol) [13] and SDP (session description protocol) [14], and the corresponding RID from its HDR_ID list which has been advertised by its affiliating HDR(k). Once the host has the fully qualified multicast address G(L,H,RID,APL_ID), it issues an IGMP report for the group address. Upon receiving the report, the HDR(k) invokes the tree joining process if $L \le k \le H$ (i.e., the HDR is located within the scope region). Otherwise, it ignores the report since the group requested by the host doesn’t belong to the scope region.

B. Establishing the HLIM Multicast Tree

Once the HDR(k) receives a IGMP report, it sets up a transient forwarding cache which consists of a group address and the corresponding outgoing interfaces. Once the cache is set, the HDR(k) sends a HLIM_Join message to its immediate parent HDR(k+1) and then waits for a HLIM_ACK message from it. The HLIM messages are multicast with IP TTL of 1.

When the HDR(k+1) receives a HLIM_Join from its D_IF, the HDR(k+1) checks if $L \le k+1 \le H$. If that is not true or the message is received from its U_IF, it ignores the message. Otherwise it relays the HLIM_Join to its parent HDR(k+2) if it is not an on-tree HDR (i.e., no existing forwarding cache for the group). If it is an on-tree HDR, it issues a HLIM_ACK back to its child HDR which issued the HLIM_Join. If k+1=H (i.e., the HDR(k+1) is located at the highest level of the scope region), it sets up a *confirmed* forwarding cache for the group and sends a HLIM_ACK to its child HDR. Once the HDRs with the *transient* forwarding cache receives a HLIM_ACK, they switch the *transient* forwarding cache to *confirmed* forwarding cache and become on-tree HDRs for the group.

Fig. 5 shows examples of HLIM join operations. The IGMP_Report from the invalid receiver is discarded since the receiver is located outside of the scope region. The join message from the *receiver 1* is forwarded up to the highest HDR on the scope region, HDR(4), and this HDR issues a HLIM_ACK down to the HDR to which the receiver is attaching. After that, the join message from the *receiver 2* traverses only up to the HDR(3) which is an on-tree HDR and a HLIM_ACK is returned by this on-tree HDR.

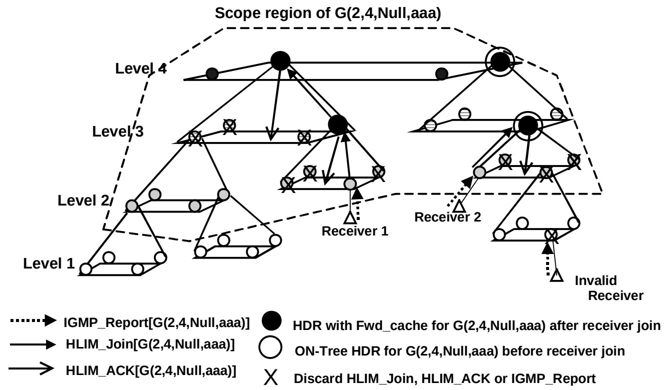

Fig. 5. HLIM regular operations for a host joining a group.

C. Forwarding Multicast Packets

When a HDR(k) receives a multicast packet (denoted as MP[G(L,H,RID,APL_ID), S]), it forwards the packet to the outgoing interfaces indicated in the forwarding cache of the addressed group (if any), except the incoming interface of the packet. After that, if the HDR(k) is inside the scope region of the addressed group, it forwards the packet to its U_IF if it has not done so and records the U_IF in the forwarding cache. This forwarding operation is designed to work not only in the normal cases but also the host and network mobility cases.

Fig. 6 shows how multicast packets are forwarded in normal cases. Note that the HLIM provides the shortest path from a source to each of the receivers and the multicast tree is shared by all sources and receivers of the same group.

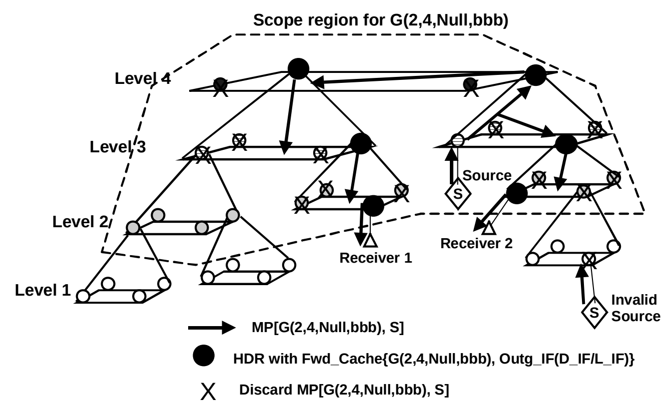

Fig. 6. HLIM Forwarding Operation.

V. HLIM OPERATIONS FOR MOBILITY

The basic concept to support mobility in HLIM is to figure out whether a mobile node (hosts or network) has moved within or outside the original scope region of an on-going multicast session. If the mobile node has moved within the scope region, the regular operations as described in previous subsection are applied. If the mobile node has moved outside the scope region, the mobility operations are invoked. The main idea for the mobility operations is to find a binding point (BPT) and extend the established multicast tree for a mobile host or network through the BPT.

The BPT is the HDR that provides linkages between the original scope region and the new location of a mobile node through the shortest path. It is located at the bottom of the original scope region for the mobile node which has moved beneath the scope region. It is the root of the original scope region if the mobile node has moved outside of (but not below) the scope region. Through the BPT, a mobile node outside its original scope region can continue to receive or send multicast packets from the sources or to the receivers in the same multicast session. It should be noted that all the multicast packets are still delivered to/from the new location of the mobile node through the shortest path. Two examples of BPTs are shown in Fig. 7.

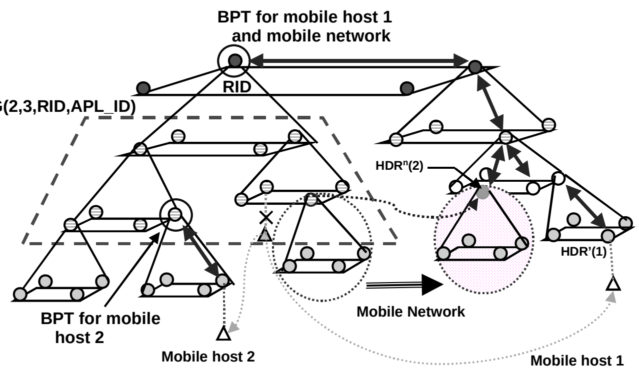

Fig. 7. Example BPTs for mobile nodes.

As described before, each HDR keeps its HDR_ID list and updates its HDR_ID list whenever any of its ascendant HDRs is changed. If a HDR (or HDR with its descendant networks) moves to an new location, it will get a new set of HDR_ID list. When the movement of a mobile host (or network) while engaging a multicast session is detected, HLIM determines the movement type (i.e., inside or outside the scope region) of the mobile host (or network) with reference to the scope region of the on-going multicast session. Such determination takes place at the new HDR to which the mobile host has moved (or at the topmost HDR on the mobile network). For example, the determination of movement type for the *mobile host 1* and *mobile network* shown in Fig. 7 takes place at HDR'(1) and $HDR^n(2)$, respectively.

A. Host Mobility Operations

1) IGMP Operations for a Mobile Host: The following new messages are added to the current IGMP to support host mobility: Explicit_M_Report for a mobile receiver, Fwd_M_Report and Fwd_M_Reply for a mobile source.

When a mobile host detects its movement to a new location while engaging a multicast session, it first issues an Explicit_M_Report[G(L,H,RID,APL_ID)] message (if it is a receiver) or a Fwd_M_Request[G(L,H,RID,APL_ID), S] message (if it is a source) to its new HDR at level k. Upon receiving the Explicit_M_Report, the HDR finds out whether the mobile host moved within or out of the original scope region by the following comparisons:

- $H = N$ (where N is the number of levels in the hierarchy) and $L \le k$
- $H < N$ and $L \le k \le H$ and RID = HDR_ID(H+1)

If one of these conditions satisfies, the HDR(k) determines that the mobile host has moved within the original scope region and invokes the regular join operations. Otherwise, the HDR(k) starts the mobile join operations by sending a mobile join message, HLIM_M_Join[G(L,H,RID,APL_ID), RID, Backward], toward the RID. Note that the direction indicator flag, Backward, implies that the mobile join is invoked by a mobile receiver and therefore the *incoming interface* of the mobile join will be recorded in the forwarding cache for the group.

Upon receiving the Fwd_M_Request, the HDR finds out whether the mobile source has moved within or out of the original scope region as described before. If the mobile source has moved within the original scope region, the HDR replies by sending a Fwd_M_Reply. Otherwise, the HDR sends a HLIM_M_Join[G(L,H,RID,APL_ID), RID, Forward] toward the RID. The direction indicator flag, Forward, implies that the message is initiated by a mobile source and therefore the *outgoing interface* of the mobile join message will be recorded in the forwarding cache for the group. After receiving a Fwd_M_Reply (which will be returned after the HLIM mobility join operations in the network are completed), a mobile source can resume sending multicast packets for the on-going multicast session to its new HDR.

2) Extending the HLIM Tree for a Mobile Host: The mobile joining operations occur outside the scope region. A HLIM_M_Join message traverses toward a RID. When the message reaches an on-tree HDR outside the scope region or the BPT, a HLIM_M_ACK message is returned and the extended multicast tree can be established along the path.

Some example operations for host mobility are shown in Fig. 8. Since *mh1* has moved within the scope region, only the regular join operations are invoked. Since *mh2* and *mh3* have moved out of the scope region, the mobile join operations are performed between the new HDR serving *mh2* and the BPT2, and between the new HDR serving *mh3* and the BPT3.

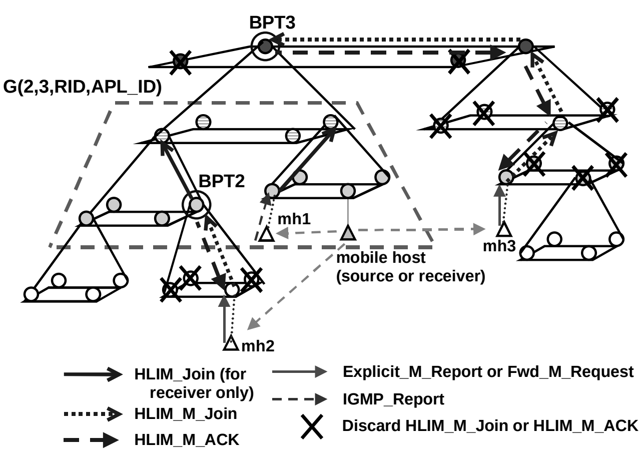

Fig. 8. Join operations for a mobile host after movement.

Fig. 9 depicts how multicast packets are forwarded from mobile sources to receivers after their movements.

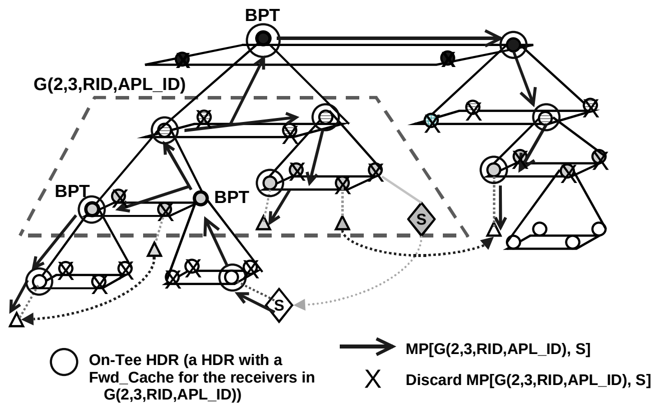

Fig. 9. Example multicast forwarding for mobile sources and receivers.

B. Network Mobility Operations

We assume that a network movement in the TI is allowed only within the same level. A router at level k can therefore only move to a core subnet comprised of parent routers with level k+1 and their child routers with level k. It can move with its whole descendant sub-networks (including its local subnet), its local subnet only, or alone. Our design goal is to hide any network mobility from the hosts or routers within the moving network. In other words, no operations are required from the hosts or routers inside the moving network. Only the “topmost” router interacts with the outside of the moving network to hand over any on-going multicast sessions to the new location.

Only the movement of an “on-tree” HDR (i.e., the HDR with one or more forwarding caches established by HLIM_Join, HLIM_M_Join or HLIM forwarding) will invoke mobility operations to hand over all the on-going multicast sessions through the router.

When a network moves to a new location, each HDR within the moving network updates its HDR_ID list at the new location. The topmost HDR on the moving network determines the movement type of each on-going multicast session by the comparison applied in the host mobility operation. Once the movement types are determined, appropriate join operations are invoked as for host mobility and therefore the extended multicast tree can be established for all the on-going multicast sessions in the moving network.

VI. CONCLUSIONS

This paper describes a new design of mobile IP multicasting protocol for tactical networks named “Hierarchical Level-based IP Multicasting” (HLIM). Both host and network mobility is taken into account in the HLIM design. HLIM takes advantage of the hierarchical structure and broadcast nature in a tactical IP network. A new flexible scope is added to the multicast addresses to better manage the multicast traffic. HLIM combines the benefits of both source-based tree and shared tree multicasting protocols. It always establishes the shortest path tree between the sources and receivers of the same group. All the sources and receivers of the same group share the same tree without a concentration point required in PIM or CBT. Such a tree is maintained dynamically to adapt to any movement or loss of the hosts or routers involved in the multicast session.

For future research, HLIM can be extended as an inter-domain IP multicasting protocol for the general Internet with the help of a protocol, such as the ATM Private Network-Network Interface (PNNI) protocol [8], to form a logical hierarchical network out of an arbitrary network. Also, with the help of MMWN (multimedia support for mobile wireless networks) [12], which can establish a logical hierarchical network in an ad-hoc network, HLIM can be extended for an ad-hoc network.

ACKNOWLEDGMENT

The authors express great appreciation to Charles Graff and Michael Bereschinsky (of U.S. Army CECOM) for their continued support and active involvement in this work. Their knowledge, especially on tactical networks proved invaluable in this project. The authors had many interesting discussions with them and John Strange (also of U.S. Army CECOM) throughout this project.

REFERENCES

[1] T. Pusateri, “Distance Vector Multicast Routing Protocol”, Internet Draft. ftp://ftp.isi.edu/internet-drafts/draft-ietf-idmr-dvmrp-v3-08.txt, Feb. 1999.

[2] J. Moy, “Multicast Extensions to OSPF”, RFC 1584, March 1994.

[3] S. Dearing, et. al., “Protocol Independent Multicast Version 2, Dense Mode Specification,” Internet Draft, ftp://ftp.isi.edu/internt-drafts/draft-ietf-pimv2-dm-0.txt, Nov. 1998.

[4] D. Estrin, et. al., “Protocol Independent Multicast-Sparse Mode (PIM-SM) : Protocol Specification”, RFC 2362, June 1998.

[5] A. Ballardie, “Core Based Trees (CBT) Multicast Routing Architecture”, RFC 2201, September 1997.

[6] A. Ballardie, “Core Based Trees (CBT version2) Multicast Routing – Protocol Specification,” RFC 2189, Sep. 1997.

[7] C. E. Perkins, “Mobile IP Design Principles and Practices”, Addison Wesley, 1998.

[8] ATM Forum, Private Network-Network Interface (PNNI) Specification Version 1.0, Mar. 1996.

[9] C. Graff, “Future Network Architecture for Tactical Army Networks”, IEEE talk, Viewgraphs available from graff@doim6.monmouth.army.mil.

[10] R. Hinden and S. Deering, “IP Version 6 Addressing Architecture”, RFC2373, July 1998.

[11] W. Fenner, “Internet Group Management Protocol, version 2 (IGMPv2),” RFC 2236, Nov. 1997.

[12] R. Ramanathan and M. Steenstrup. “Hierarchically-organized, multi-hop mobile wireless networks for quality-of-service support”, ACM Mobile Networks & Applications (MONET), June 1998.

[13] Handley, M., “SAP: Session announcement protocol,” Internet Draft, Internet Engineering Task Force, Nov. 1996. Work in progress.

[14] Handley, M. and V. Jacobson, “SDP: session description protocol”, RFC 2327, April 1998.

[15] G. Xylomenos and G. C. Pdyzos, “IP Multicast for Mobile Hosts”, IEEE Comm. Mag., Jan. 1997.

[16] Gary R. Wright and W. Richard Stevens, “TCP/IP Illustrated, Volume 2”, Addison Wesley, 1995.

0-7803-5538-5/99/$10.00 (c) 1999 IEEE