IETF MANET Working Group
INTERNET-DRAFT
Expiration: May 17, 2003

Mario Gerla
Xiaoyan Hong
Li Ma
University of California, Los Angeles
Guangyu Pei
Rockwell Scientific Company

November 17, 2002

Landmark Routing Protocol (LANMAR) for Large Scale Ad Hoc Networks

<draft-ietf-manet-lanmar-05.txt>

Status of This Memo

This Internet-Draft is a submission to the IETF Mobile Ad Hoc Networks (MANET) Working Group. Comments on this draft may be sent to the Working Group at manet@itd.nrl.navy.mil, or may be sent directly to the authors.

Abstract

The Landmark Routing Protocol (LANMAR) utilizes the concept of landmark for scalable routing in large, mobile ad hoc networks. It relies on the notion of group mobility: i.e., a logical group (for example a team of coworkers at a convention) moves in a coordinated fashion. The existence of such logical group can be efficiently reflected in the addressing scheme. It assumes that an IP like address is used consisting of a group ID (or subnet ID) and a host ID, i.e. <Group ID, Host ID>. A landmark is dynamically elected in each group. The route to a landmark is propagated throughout the network using a Distance Vector mechanism. Separately, each node in the network uses a scoped routing algorithm (e.g., FSR) to learn about routes within a given (max number of hops) scope. To route a packet to a destination outside its scope, a node will direct the packet to the landmark corresponding to the group ID of such destination. Once the packet approaches the landmark, it will typically be routed directly to the destination. A solution to nodes outside of the scope of their landmark (i.e., drifters) is also addressed in the draft. Thus, by summarizing in the corresponding landmarks the routing information of remote groups of nodes and by using the truncated local routing table, LANMAR dramatically reduces routing table size and routing update overhead in large networks. The dynamic election of landmarks enables LANMAR to cope with mobile environments. LANMAR is well suited to provide an efficient and scalable routing solution in large, mobile, ad hoc environments in which group behavior applies and high mobility renders traditional routing schemes inefficient.

Contents

Status of This Memo 1

Abstract 1

1. Introduction 3

2. Changes 4

3. Terminology 6
    3.1. General Terms 6
    3.2. Specification Language 7

4. Protocol Applicability 7
    4.1. Networking Context 7
    4.2. Protocol Characteristics and Mechanisms 7

5. Protocol Overview 9
    5.1. Protocol Descriptions 9
    5.2. Landmark Election 10

5.3. Drifters 10

6. Protocol Specifications 11
6.1. Data Structures 11
6.1.1 Landmark Status tuple 11
6.1.2 Landmark Distance Vector 12

6.1.3 Drifter List 12
6.2. LANMAR Update Message Format 12
6.2.1 Description of the fields 13
6.2.2 Propagation of LANMAR Update Messages 14
6.3. Processing Landmark Updates 14
6.3.1 Originating a Landmark in a Subnet 14
6.3.2 Receiving Landmark Updates 15
6.3.3 Winner Competition 15
6.3.4 Dealing with Stale LMDV Entries 16
6.4. Processing Drifter Updates 16
6.4.1 Originating a Drifter Entry 16
6.4.2 Receiving Drifter Updates 16
6.4.3 Removing a Drifter Entry 17
6.5. Operations Regarding to Lost Neighbors 17

7. Data Packet Forwarding 17

8. Combining LANMAR with a Host Protocol 17
8.1. Share Neighbor Information 18
8.1.1 Inform the Host Protocol about a Neighbor 18
8.1.2 Being Informed about a Lost Neighbor 18
8.2. Scoped Routing Operations 18
8.2.1 Destination-Sequenced Distance Vector 18
8.2.2 Fisheye State Routing Protocol 18
8.2.3 Optimized Link State Routing Protocol 18
8.2.4 Topology Broadcast Based on Reverse-Path Forwarding 19

9. Implementation Status 19

Acknowledgments 19

References 19

Chair's Address 20

Authors' Addresses 20

1. Introduction

This document describes the Landmark Routing Protocol (LANMAR) [1,2] developed by the Wireless Adaptive Mobility (WAM) Laboratory [4] at Computer Science Department, University of California, Los Angeles.

The concept of landmark routing was first introduced in wired area networks [6]. The original scheme required predefined multi-level hierarchical addressing. The hierarchical address of each node reflects its position within the hierarchy and helps to find a route to it. Each node knows the routes to all the nodes within its hierarchical partition. Moreover, each node knows the routes to various landmarks at different hierarchical levels. Packet forwarding is consistent with the landmark hierarchy and the path is gradually refined from the top level hierarchy to lower levels as a packet approaches its destination.

LANMAR borrows the concept of landmark and extends it to the wireless ad hoc environment. LANMAR scheme does not require predefined hierarchical address, but it uses the notion of landmarks to keep track of logical subnets in which the members have a commonality of interests and are likely to move as a group (e.g., brigade in the battlefield, a group of students from same class and a team of co-workers at a convention). Each such logical group has an elected landmark. For each group, the underlying scoped routing algorithm will provide accurate routing information for nodes within scope. The routing update packets are restricted only within the scope. The routing information to remote nodes (nodes outside the node's scope) is summarized by the corresponding landmarks. Thus, the LANMAR scheme largely reduces the routing table size and the routing update traffic overhead. It greatly improves scalability.

In addition, in order to recover from landmark failures, a landmark node is elected in each subnet. Landmark election provides a flexible way for the LANMAR protocol to cope with a dynamic and mobile network. The protocol also provides a solution for nodes that are outside the scopes of the landmarks of their logical groups (drifters).

The LANMAR runs on top of a proactive routing protocol. It requires that the underlying routing protocol support the scoped subnetworking. Many MANET proactive routing protocols (e.g., DSDV, FSR, OLSR and TBRPF) can be the host protocol with only minor modifications (Section 8). The main advantage of LANMAR is that the routing table includes only the nodes within the scope and the landmark nodes. This feature greatly improves scalability by reducing routing table size and update traffic O/H.

Thus the Landmark Routing Protocol provides an efficient and scalable routing solution for a mobile, ad hoc environment while keeping line and storage overhead (O/H) low. Moreover, the election provides a much needed recovery from landmark failures.

2. Changes

Major changes from version 04 to version 05:

- Clarified the relation between LANMAR and a host protocol, which includes: Removed "Neighbor List" from LANMAR protocol and description about related operations; Added operation descriptions about sharing neighbor information with the host protocol; Retained earlier descriptions about modifications on

making various MANET proactive routing protocols to route within scope, and TBRPF is added. All these combination efforts are organized in Section 8.

- Removed "Landmark Flag" field from LANMAR Update message format. Instead, "Packet type" field is added.

- Added "last heard time" field in LMDV. Also added operations regarding to updating this field and timing out stale entries.

- Removed the description implying the dependency of LANMAR on the local topology information, (as that is not the case any more,) from descriptions of propagating and receiving a LMU.

- Editorial changes.

Major changes from version 03 to version 04:

- Removed "neighbor landmark flag" field from neighbor list. Clarified the operations when a neighbor is lost.

- Clarified the processing of landmark update messages, especially, the operations when an infinite distance metric occurs. Operation regarding to an infinite distance metric is also added in data forwarding.

- A separate section describing the operation before sending a landmark update message is added.

- Reported current implementation status.

- Editorial changes.

Major changes from version 02 to version 03:

- A drifter sequence number is used in drifter list to indicate each new occurrence of a drifter.

- Processing of lost neighbors is added.

- A separate section describing the modifications made to various proactive protocols. Operations of these protocols will then only perform within a certain hop distances.

- Editorial changes.

Major changes from version 01 to version 02:

- Update of Status of This Memo.

Major changes from version 00 to version 01:

- A destination sequence number for each landmark is used to ensure loop-free updates for a particular landmark.

- Landmark updates are propagated in separate messages, instead of being piggybacked on local routing updates. This modification decouples landmark routing from the underlying proactive routing protocol.

3. Terminology

3.1. General Terms

This section defines terminology used in LANMAR.

node

A MANET router that implements Landmark Routing Protocol.

neighbor

Nodes that are within the radio transmission range.

scope

A network area that is centered at each node and bounded by a certain maximum hop distances.

host protocol

Also known as local routing protocol, i.e., a proactive protocol that works together with the Landmark Routing Protocol, but only operates within the scope of each node.

underlying protocol

This term is used interchangeably with host protocol.

scoped routing protocol

A routing protocol that only exchanges routing information up to a certain hop distance (scope).

subnet

Logical groups of nodes that present similar motion behavior.

group

This term is used interchangeably with subnet.

3.2. Specification Language

The keywords MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD, SHOULD NOT, RECOMMENDED, MAY, and OPTIONAL in this document are to be interpreted as described in RFC 2119 [5].

4. Protocol Applicability

4.1. Networking Context

LANMAR is best suited for large scale mobile ad hoc wireless networks. The landmark scheme on top of a scoped routing algorithm has large advantages in reducing routing update packet size and keeping reasonably accurate routes to remote nodes. Moreover, the fact that the route error is blurred by distance obviously reduces the sensitivity to network size.

LANMAR is also suited for high mobility ad hoc wireless networks. This is because in a mobile environment, a change on a link far away from the source does not necessarily cause a change in the routing table at the source since all the information about remote nodes is summarized by landmarks.

4.2. Protocol Characteristics and Mechanisms

* Does the protocol provide support for unidirectional links?(if so, how?)

    No.

* Does the protocol require the use of tunneling? (if so, how?)
    No.

* Does the protocol require using some form of source routing? (if so, how?)

    No.

* Does the protocol require the use of periodic messaging? (if so, how?)

    Yes. The LANMAR periodically broadcasts landmark information to its neighbors.

* Does the protocol require the use of reliable or sequenced packet delivery? (if so, how?)

  No. As the packets are sent periodically, they need not be

  sent reliably.

* Does the protocol provide support for routing through a multi-technology routing fabric? (if so, how?)

  Yes. It is assumed that each node's network interface is assigned a unique IP address.

* Does the protocol provide support for multiple hosts per router? (if so, how?)

  Yes. The router that has multiple hosts can use network ID of these hosts as the address to participate LANMAR.

* Does the protocol support the IP addressing architecture? (if so, how?)

  Yes. Each node is assumed to have a unique IP address (or set of unique IP addresses in the case of multiple interfaces). The LANMAR references all nodes/interfaces by their IP address. This version of the LANMAR also supports IP network addressing (network prefixes) for routers that provide access to a network of non-router hosts.

* Does the protocol require link or neighbor status sensing (if so, how?)

  No.

* Does the protocol have dependence on a central entity? (if so, how?)

  No.

* Does the protocol function reactively? (if so, how?)

  No.

- Does the protocol function proactively? (if so, how?)

  Yes. The LANMAR proactively maintains landmark information at each node.

- Does the protocol provide loop-free routing? (if so, how?)

  Yes. For in-scope destinations, the protocol uses routing paths learned from the host protocol. If the host protocol provides loop-free routing, e.g., FSR and DSDV, so does LANMAR. For out-scope destinations, only routes to landmarks are used. Because these routes are DSDV, it is loop free. When a packet approaches the destination, in-scope routes are used again.

- Does the protocol provide for sleep period operation?(if so, how?)

  Yes. However, this requires TDMA MAC layer support. The router can be scheduled to sleep during idle periods.

- Does the protocol provide some form of security? (if so, how?)

  Yes. When a node broadcasts routing update message, only entries of in-scope nodes and landmarks are included. This will prevent other remote nodes from being heard.

- Does the protocol provide support for utilizing multi-channel, link-layer technologies? (if so, how?)

  Yes. In fact, the multi-channel can be used to separate routing messages from user data packets.

5. Protocol Overview

5.1. Protocol Descriptions

As mentioned in Section 1, the landmark concept we adopt here uses the notion of logical subnets in which the members have a commonality of interests and are likely to move as a group. Each logical subnet has one node serving as a landmark of that subnet. The protocol requires that the landmark of each subnet have the knowledge of all the members in its group. The protocol deploys a routing scope at each node. The size of the scope is a parameter measured in hop distance. It is chosen in such a way that if a node is at the center of a subnet, the scope will cover the majority of the subnet members. If the shape of a subnet is likely to be a cycle, the center node's scope will cover all the members of the subnet. If this center node is elected as a landmark, it fulfills the requirement of the protocol.

LANMAR is supported by two complementary, cooperating routing schemes: (a) a local, "myopic" routing scheme (operating within the limited scope) that maintains routes to nearby destinations and; (b) a "long haul" distance vector routing scheme that propagates landmark information. An elected landmark uses a destination sequence number to ensure its routing entry update loop-free. Thus, Each node maintains a distance vector for landmarks of all the subnets. The size of the landmark distance vector equals to the number of logical subnets and thus landmark nodes.

When there are members drifting off their landmark's scope, the landmark will create separate distance vector entries for them. The distance vectors for landmarks and drifters are exchanged among neighbors in periodical routing update messages.

As a result, in-scope destinations can be routed through accurate

routes obtained from local routing table. A data packet directed to a remote destination initially aims at the landmark of that remote group; as it gets closer to the landmark, it may eventually switch to the accurate route to the destination provided by local host protocol at some nodes within the scope of the destination.

5.2. Landmark Election

Dynamic election/re-election of landmark node is essential for LANMAR to work in a wireless mobile environment. Basically, each node tracks other nodes of its group in its scope and computes weight, e.g. the number of the nodes it has found. At the beginning of the LANMAR, no landmark exists. Protocol LANMAR only uses the host protocol functionality. As host routing computation progresses, one of the nodes will learn (from the system routing table) that more than a certain number of group members (say, T) are in its scope. It then proclaims itself as a landmark for this group and adds itself to the landmark distance vector. Landmarks broadcast the election weights to the neighbors jointly with the landmark distance vector update packets.

When more than one node declares itself as a landmark in the same group, as the landmark information floods out, each node will perform a winner competition procedure. Only one landmark for each group will survive and it will be elected. To avoid flapping between landmarks (very possible in a mobile situation), we may use hysteresis in the replacement of an existing landmark. I.e., the old Landmark is replaced by the new one only if its weight is, say less than 1/2 of the weight of the current election winner. Once ousted, the old leader needs the full weight superiority to be reinstated.

This procedure is carried out periodically in the background. At steady state, a landmark propagates its presence to all other nodes like a sink in DSDV. It is extremely simple and it converges (by definition). In a mobile environment, an elected landmark may eventually lose its role. The role shifting is a frequent event. In a transient period, there exist several landmarks in a single group. The transient period may be actually the norm at high mobility. This transient behavior can be drastically reduced by using hysteresis.

When a landmark dies, its neighbors will detect the silence after a given timeout period. A new round of landmark election will then start from new qualified members and flood over the group.

5.3. Drifters

Typically, all members in a logical subnet are within the scope of the landmark, thus the landmark has a route to all members. It may

happen, however, that some of the members drift off the scope, for example, a tank in a battalion may become stranded or lost. To keep track of such drifters, i.e., to make the route to them known to the landmark, the following modification to the routing table exchange is necessary. Each node, say i, on the shortest path between a landmark L and a drifter l associated with that landmark keeps a distance vector entry to l. Note that if l is within the scope of i, this entry is already included in the routing table of node i. When i transmits its distance vector to neighbor, say j, then j will retain the entry to member l only if d(j,l) is smaller than the scope or d(j,L) is smaller than d(i,L). The latter condition occurs if j is on the shortest path from i (and therefore from l) to L. This way, a path is maintained from the landmark to each of its members, including drifters. The procedure starts from l, at the time when a node finds it becomes a drifter. It informs the landmark hop by hop about its presence.

The occurrences of drifters are dynamic in a mobile network. In order to timely remove the staled drifter information, the time when a node hears a drifter is recorded. A node monitors whether it becomes a drifter periodically and refreshes its occurrence along the path towards the landmark.

6. Protocol Specifications

This section discusses the operation of LANMAR routing protocol. The sending and receiving of landmark updates are in a proactive nature. The routing packets are processed separately from ordinary data packets.

6.1. Data Structures

Each node has a unique logical identifier defined by a subnet field and a host field. The host field is unique in the subnet and might in fact coincide with the physical address. The logical identifier can also be an IP address when the subnet address logically reflects the grouping of nodes.

As LANMAR runs on top of a host routing protocol, it shares the routing table with the host protocol, i.e., both host protocol and LANMAR contribute their portion to the system routing table. LANMAR's routing table portion includes a landmark distance vector (LMDV) and a drifter distance vector (DFDV).

In addition to the routing tables, each node has a landmark status tuple. LANMAR does not maintain a separate neighbor list. Instead, it interacts with the host protocol for possible table updates caused by neighbor changes. In this draft, we only describe data structures that pertain to LANMAR.

6.1.1. Landmark Status Tuple

Each node has not only a logical identifier, which basically is its address, but also a landmark status tuple. The tuple includes a flag which indicates whether the node is a landmark or not, a election weight (the number of group members the node detects within its scope) and a sequence number. When a node is elected, the status tuple will be copied to its landmark distance vector. The sequence number is advanced. These are the three fields for a tuple:
- Landmark flag
- Weight: Number of group members in its scope
- Sequence number

6.1.2. Landmark Distance Vector

Landmark distance vector (LMDV) gives the next hop information to all landmarks in the network. Every subnet has an entry in LMDV. The latest route information to the landmark of each subnet is learned when a landmark update message is received. LMDV functions as a part of the routing table. It has the following fields:
- Landmark status tuple
- Next hop address
- Distance
- Last heard time

6.1.3. Drifter List

The drifter list (DFDV) of LANMAR provides the next hop information to the drifters known to the current node. The entries are updated with landmark update message. The latest time a drifter is heard is recorded in DFDV. The DFDV works as a part of routing table. It has the following fields:
- Destination drifter address
- Next hop address
- Distance
- Drifter sequence number
- Last heard time

6.2. LANMAR Update Message Format

There is only one message type of LANMAR protocol: LANMAR Update (LMU). The messages are periodically exchanged with neighbors. They update the landmark distance vector LMDV and the drifter list DFDV. It is possible that LMDV or DFDV is empty, so no entries of the empty table will be included. The processing of the LMDV and DFDV will be describe separately. The following format does not include the node's identifier because it can be obtained from IP Header.

```
    0                   1                   2                   3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

```
   | Type  |  N_landmarks  |  N_drifters   |       Reserved        |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                       Landmark Address 1                      |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                      Next Hop  Address 1                      |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |  Distance 1   |  N_members 1  |       Sequence Number 1       :
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   :                             . . .                             :
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                        Drifter Address 1                      |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                      Next Hop  Address 1                      |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |  Distance 1   |   Drifter Sequence Number 1   |     ...       :
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   :                             . . .                             |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

6.2.1. Description of the fields

Type

This field indicates the packet type. Currently LANMAR only has one packet type, i.e., LMU. The field is also needed to distinguish LANMAR routing control packets from the host protocol routing packets.

N_landmarks

The number of entries of the landmark distance vector.

N_drifters

The number of entries of the drifter list.

Reserved

The bits are set to '0' and are ignored on reception.

Landmark Address 1, Next Hop Address 1, Distance 1, N_members 1 and Sequence Number 1

The first entry in the landmark distance vector.
Landmark Address 1, N_members 1 and Sequence Number 1 are the status tuple of the destination landmark.
Next Hop Address 1 and Distance 1 is the next hop and distance to the landmark.

These fields are repeated N_landmarks times for each entry in landmark distance vector.

Drifter Address 1, Next Hop Address 1, Distance 1 and Drifter Sequence Number 1

The first entry in the drifter list.
Next Hop Address 1 and Distance 1 are the next hop and distance to the Drifter Address 1.

These fields are repeated N_drifters times for each entry in the drifter list.

The length of the message is limited to the maximum IP packet size. In that case, multiple packets may be required to broadcast all the entries.

6.2.2. Propagation of LANMAR Update Messages

LANMAR update messages (LMUs) are propagated periodically. The frequency could be set according to mobility. Propagation jitters must be used for each message transmission to reduce collisions. Before sending each LMU, a node first performs drifter operations (described in Section 6.4.1) to check whether it is a drifter. An existing drifter node increases its drifter sequence number by 2. Then a node recalculates the current number of group members within its scope from the system routing table. The new number is recorded at the election weight field of its landmark status tuple. And if it is a landmark, the corresponding entry in the LMDV must be updated with this new weight. For a landmark node, its sequence number must increase by 2 both in its status tuple and in its LMDV entry. Then, LMDV will be searched for stale entries. Operations are given in 6.3.4. DFDV will be searched for stale entries too. Operations are given in 6.4.3. Finally, the node assembles in the LMU its LMDV and DFDV.

6.3. Processing Landmark Updates

Landmark update information is a part of the LANMAR periodic routing update message. The update information includes sender's LMDV. Landmark update message is used for landmark election and building paths to landmarks.

6.3.1. Originating a Landmark in a Subnet

Before propagating of each LMU, when a node obtains a new weight, It will check if this number is greater than a threshold T. The node qualifies as a landmark when one of the following conditions is met.
  (1) it is the only landmark for the group so far;
  (2) the existing landmark has an infinite distance metric;
  (3) it wins the election (see 6.3.3) when competing with the existing landmark.
When the node becomes a landmark, it increases its sequence

number by 2. Its current landmark status tuple will be inserted into the LMDV or the existing landmark is replaced with the new winner. This landmark entry will be broadcast to neighbors with the next update packet.

6.3.2. Receiving Landmark Updates

When a node receives a landmark update message, it recalculates the current number of group members within its scope from the system routing table first. The new number is recorded at the election weight field of its landmark status tuple. And if it is a landmark, the corresponding entry in the LMDV must be updated with this new weight. The node compares its LMDV entries with the entries in the incoming LMDV message for each subnet. There are three situations to consider:
(1) An incoming landmark entry corresponding to a new subnet and its distance metric is less than infinity, the entry

will be copied.

(2) An incoming entry having the same landmark as an existing one (in node's LMDV), it will be accepted only if one of the following conditions is met.
   (a) it contains a larger sequence number and the distance metric is less than infinity;
   (b) it contains a larger sequence number and the distance metric equals to infinity and it happens to be the next hop in the already existing entry;
   (c) it has the same sequence number with the existing entry, but a smaller distance metric.

(3) If an incoming entry contains a different landmark for the same subnet as recorded in LMDV, only the landmark that does not have an infinite distance and is elected through a winner competition algorithm (see 6.3.3) is accepted. The LMDV entry will be kept/updated according to the outcomes of the competition.

During the processing of the current LMU, each inserted or updated LMDV entry is stamped with receiving time in its "last heard time" field. After a LMU is processed, the LMDV will be search for stale entries. Operations are given in 6.3.4.

6.3.3. Winner Competition

When more than one node declares itself as a landmark in the same group, a simple solution is to let the node with the largest number of group members win the election and in case of tie, lowest ID breaks the tie. The other competing nodes defer.

To use hysteresis in replacing an existing landmark, let us assume the competing node's number of members is M, the existing landmark's number of members is N and a factor value S. When M is greater than N*S, then the competing node replaces the existing landmark. Or, when N reduces to a value smaller than a threshold T,

then it gives up the landmark role. A tie occurs when M falls within an interval [N*1/S, N*S], then the node with a larger member number wins the election. If a tie occurs again with equal member number, i.e., M equals to N, it is broken using lowest ID. A tie can always be broken using lowest ID as the address is used as ID and it is unique.

6.3.4. Dealing with Stale LMDV Entries

Each entry in LMDV is time stamped of its last receiving time. Every time before a LMU message is propagated or after a LMU is processed, the LMDV table is checked for staled entries. If such an entry is found, it must be marked an infinite distance metric and the sequence number be increased by 1. Thus, the lost landmark entry will overwrite the entries at downstream nodes. A new elected landmark will replace the lost one.

6.4. Processing Drifter Updates

Drifter update information is a part of the LANMAR periodical routing update message. The update information is the drifter list (DFDV) of the sender. The computation of the DFDV at each node includes checking the node itself to see whether it is a drifter and recording paths to other drifters.

6.4.1. Originating a Drifter Entry

By checking the distance to the landmark of its group, each node easily knows whether it has become a drifter. If the distance is larger than the scope, the node will put itself into its drifter list. This drifter information will be sent back to the landmark hop by hop along the shortest path to it which can be learned from the LMDV. For each drifter, only the node on its shortest path to the landmark needs to receive its information, so before the entry is broadcast, the drifter or the intermediate nodes look for the next hop to drifter's landmark in their LMDVs first. Then the next hop is included in LMU within the drifter entry. Each drifter also maintains a drifter sequence number. Each time a node finds itself a drifter, the sequence number will be increased by 2. The DFDV will be propagated with the next update packet.

6.4.2. Receiving Drifter Updates

Upon receiving an update packet, the DFDV part is retrieved and processed. If an entry of incoming DFDV indicates that the current node is its next hop to the landmark, i.e., the current node is on the drifter's shortest path to the landmark, the current node will insert or update its drifter list. The receiving time is stamped in the DFDV. The node sending the update packet is recorded in DFDV as the next hop to the drifter. The reverse path to the drifter is thus built up. The procedure ends when the landmark receives the drifter entry. The updated DFDV will be propagated with the next update packet.

6.4.3. Removing a Drifter Entry

Each entry in DFDV is time stamped of its last receiving time. Every time before the DFDV is sent or routing by DFDV is needed, the table is checked for staled entries. If such an entry is found, it is removed.

6.5. Operations Regarding to Lost Neighbors

When a lost neighbor is reported by the host protocol (the host protocol may discover by checking staled entries in a neighbor list or by receiving a feedback from the MAC layer protocol), LMDV and DFDV will be searched. If the lost neighbor happens to be the next hop to a landmark, the corresponding table entry in LMDV must be marked an infinite distance metric and the sequence number be increased by 1. Thus, the new link break information will overwrite the entries at downstream nodes. Till the landmark propagates the next new update message with a sequence number increased by 2, new routes will build up. If the lost neighbor happens to be the next hop to a drifter, the corresponding table entry in DFDV is removed.

7. Data Packet Forwarding

Data packets are relayed hop by hop. Each node's system routing table contains routing entries from the host protocol's routing table (called local routing table), drifter distance vector and landmark distance vector. The tables are looked up sequentially for the destination entry. If the destination is within a node's scope, the entry can be found directly in the local routing table and the packet is forwarded to the next hop node. Otherwise, the drifter list DFDV is searched for the destination. If the entry is found, the packet is forwarded using the next hop address from DFDV. If not, the logical subnet field of the destination is retrieved and the LMDV entry of the landmark corresponding to the destination's logical subnet is searched. If the distance metric is not an infinity, the data packet is then routed towards the landmark using the next hop address from LMDV. If all these attempts are failed, the data packet is dropped. When the data packet is routed towards a landmark, it is not necessary for the packet to pass through the landmark. Rather, once the packet gets within the scope of the destination on its way towards the landmark, it is routed to the destination directly.

8. Combining LANMAR with a Host Protocol

Current LANMAR and a host protocol have separate data structures, separate timers and separate routing messages, resulting in LANMAR interfering little with the local routing information. However, they may still have correlation when they work together. Operations needed for combination are listed and explained in this section.

8.1. Share Neighbor Information

As LANMAR does not maintain a neighbor table, it may share information about neighbors with the host protocol.

8.1.1. Inform the Host Protocol about a Neighbor

When a node receives a landmark update, LANMAR may inform the host protocol about the sender in case the sender is a new comer. The host protocol, if a neighbor table is maintained, may update its neighbor table accordingly.

8.1.2. Being Informed about a Lost Neighbor

LANMAR may accept the notification from a host protocol regarding to the loss of a neighbor, which leads to a search and possible updating in LMDV and DFDV (see section 6.5).

8.2. Scoped Routing Operations

8.2.1. Destination-sequenced Distance Vector routing protocol

Distance Vector type routing protocols use smaller routing tables (comparing to Link State type) and generate lower routing overhead. Destination-sequenced Distance Vector Routing (DSDV) [3] uses destination sequenced sequence numbers to prevent the forming of loops. The protocol can work together with LANMAR. The modifications include containing only the destinations within the local scope in the periodic routing update messages and turning off the triggered updates.

8.2.2. Fisheye State Routing protocol

Fisheye State Routing (FSR) [7][8] is easy to adapt to a host protocol. A two level Fisheye scope is used when FSR is used as a host protocol. For nodes within the scope, the updating is in a certain frequency. But for nodes beyond the scope, the update frequency is reduced to zero. As a result, each node maintains accurate routing information for in-scope nodes.

8.2.3. Optimized Link State Routing protocol

Optimized Link State Routing (OLSR) [9] provides the facility for scope-limited flooding of messages. The generic message

format contains a Time To Live field, which gives the maximum number of hops that a message will travel. Each time a message is retransmitted, the Time To Live field is decreased by 1. When the value of this field is reduced to zero, the massage will not be forwarded any more.

OLSR can be one of the underlying protocol of LANMAR. The Time To Live field is set to the scope defined in LANMAR when a message is originated.

8.2.4. Topology Broadcast Based on Reverse-Path Forwarding

Topology Broadcast Based on Reverse-Path Forwarding (TBRPF) [10] can be adapted to scoped routing operations as easy as that when constructing the "reportable node set" RN, only the nodes that are within scope are included. The source tree so built up at a node will reach only up to a limited hop distance. To achieve this, the hop distance should be one of the link metrics.

9. Implementation Status

LANMAR version 1 (according to version 3 of the draft, but excluding the drifter operations) has been implemented in Linux and is in use for laboratory experiments.

Acknowledgments

This work was supported in part by NSF under contract ANI-9814675, by DARPA under contract DAAB07-97-C-D321, and by ONR under contract N00014-01-C-0016.

References

[1] G. Pei, M. Gerla and X. Hong, LANMAR: Landmark Routing for Large Scale Wireless Ad Hoc Networks with Group Mobility, Proceedings of IEEE/ACM MobiHOC 2000, Boston, MA, Aug. 2000.

[2] M. Gerla, X. Hong, G. Pei, Landmark Routing for Large Ad Hoc Wireless Networks, Proceedings of IEEE GLOBECOM 2000, San Francisco, CA, Nov. 2000.

[3] C.E. Perkins and P. Bhagwat, Highly Dynamic Destination-Sequenced Distance-Vector Routing (DSDV) for Mobile Computers, In Proceedings of ACM SIGCOMM'94, London, UK, Sep. 1994, pp. 234-244.

[4] UCLA Wireless Adaptive Mobility (WAM) Laboratory. http://www.cs.ucla.edu/NRL/wireless

---

[5] S. Bradner. Key words for use in RFCs to Indicate Requirement Levels. RFC 2119, March 1997.

[6] P. F. Tsuchiya, The Landmark Hierarchy: a new hierarchy for routing in very large networks, Computer Communication Review, vol.18, no.4, Aug. 1988, pp. 35-42.

[7] G. Pei, M. Gerla, and T.-W. Chen, Fisheye State Routing: A Routing Scheme for Ad Hoc Wireless Networks, Proceedings of ICC 2000, New Orleans, LA, Jun. 2000.

[8] G. Pei, M. Gerla, and T.-W. Chen, Fisheye State Routing in Mobile Ad Hoc Networks, Proceedings of Workshop on Wireless Networks and Mobile Computing, Taipei, Taiwan, Apr. 2000.

[9] P. Jacquet, P. Muhlethaler, A. Qayyum, A. Laouiti, L. Viennot and T. Clausen, Optimized Link State Routing Protocol, Internet

Draft, IETF MANET Working Group, draft-ietf-manet-olsr-04.txt, Mar. 2002.

[10] R. G. Ogier, F. L. Templin, B. Bellur, M. G. Lewis, Topology Broadcast Based on Reverse-Path Forwarding (TBRPF), Internet Draft, IETF MANET Working Group, draft-ietf-manet-tbrpf-05.txt, Mar. 2002.

Chair's Address

The MANET Working Group can be contacted via its current chairs:

M. Scott Corson
Flarion Technologies, Inc.
Bedminster One
135 Route 202/206 South
Bedminster, NJ 07921
USA

Phone: +1 908 947-7033
Email: corson@flarion.com

Joseph Macker
Information Technology Division
Naval Research Laboratory
Washington, DC 20375
USA

Phone: +1 202 767-2001
Email: macker@itd.nrl.navy.mil

Authors' Addresses

Questions about this document can also be directed to the authors:

Mario Gerla
3732F Boelter Hall
Computer Science Department
University of California
Los Angeles, CA 90095-1596

USA

Phone: +1 310 825-4367
Fax: +1 310 825-7578
Email: gerla@cs.ucla.edu

Xiaoyan Hong
3803F Boelter Hall
Computer Science Department
University of California
Los Angeles, CA 90095-1596
USA

Phone: +1 310 825-4623
Fax: +1 310 825-7578
Email: hxy@cs.ucla.edu

Li Ma
3803D Boelter Hall
Computer Science Department
University of California
Los Angeles, CA 90095-1596
USA

Phone: +1 310 825-1888
Fax: +1 310 825-7578
Email: mary@cs.ucla.edu

Guangyu Pei
Rockwell Scientific Company
1049 Camino Dos Rios
P.O. Box 1085
Thousand Oaks, CA 91358-0085
USA

Phone: +1 805 373-4639
Fax: +1 805 373-4383
Email: gpei@rwsc.com