Mobile Ad Hoc Networking Working Group Elizabeth M. Royer
INTERNET DRAFT University of California, Santa Barbara
15 July 2000 Charles E. Perkins
Nokia Research Center

Multicast Ad hoc On-Demand Distance Vector (MAODV) Routing
draft-ietf-manet-maodv-00.txt

Status of This Memo

Abstract

The multicast operation of the Ad hoc On-Demand Distance Vector (AODV) routing protocol (MAODV) is intended for use by mobile nodes in an ad hoc network. It offers quick adaptation to dynamic link conditions, low processing and memory overhead, and low network utilization. It creates bi-directional shared multicast trees

connecting multicast sources and receivers. These multicast trees are maintained as long as group members exist within the connected portion of the network. Each multicast group has a group leader whose responsibility is maintaining the group sequence number, which is used to ensure freshness of routing information.

Contents

1. Introduction

The Multicast Ad hoc On-Demand Distance Vector (MAODV) protocol enables dynamic, self-starting, multihop routing between participating mobile nodes wishing to join or participate in a multicast group within an ad hoc network. The membership of these multicast groups is free to change during the lifetime of the network. MAODV enables mobile nodes to establish a tree connecting multicast group members. Mobile nodes are able to respond quickly to link breaks in multicast trees by repairing these breaks in a timely manner. In the event of a network partition, multicast trees are established independently in each partition, and trees for the same multicast group are quickly connected if the network components merge.

One distinguishing feature of MAODV is its use of sequence numbers for multicast groups. Each multicast group has its own sequence number, which is initialized by the multicast group leader and incremented periodically. Using these sequence numbers ensures that routes found to multicast groups are always the most current ones available. Given the choice between two routes to a multicast tree, a requesting node always selects the one with the greatest sequence number.

MAODV is the multicast protocol associated with the Ad hoc On-Demand Distance Vector (AODV) routing protocol, and as such it shares many similarities and packet formats with AODV. The Route Request and Route Reply packet types are based on those used by AODV, as is the unicast Route Table. Similarly, many of the configuration parameters used by MAODV are defined by AODV. The reader is referred to the AODV Internet draft [3] for the suggested values of those parameters, as well as for details on the unicast operation of AODV.

2. Overview

Route Requests (RREQs), Route Replies (RREPs), Multicast Activations (MACTs), and Group Hellos (GRPHs) are the message types utilized by the multicast operation AODV. RREQs and RREPs are handled as specified in [3], except for certain procedures controlled by new flags (see section 5, 6). These message types are handled by UDP, and normal IP header processing applies. So, for instance, the requesting node is expected to use its IP address as the source IP address for the messages. The range of dissemination of broadcast RREQs can be indicated by the TTL in the IP header. Fragmentation is typically not required.

As long as the multicast group members remain connected (within a ''multicast tree''), MAODV does not play any role. When a node either wishes to join a multicast group or find a route to a multicast group, the node uses a broadcast RREQ to discover a route to the multicast tree associated with that group. For join requests, a route is determined when the RREQ reaches a node that is already a member of the multicast tree, and the node's record of the multicast group sequence number is at least as great as that contained in the RREQ. For non-join requests, any node with a current route to the multicast tree may respond to the RREQ. A current route is defined as an unexpired multicast route table entry whose associated sequence number for the multicast group is at least as great as that contained in the RREQ. The route to the multicast tree is made available by unicasting a RREP back to the source of the RREQ. Since each node receiving the request caches a route back to the source of the request, the RREP can be unicast back to the source from any node able to satisfy the request. Once the source node has waited the discovery period to receive RREPs, it selects the best route to the multicast tree and unicasts the next hop along that route a MACT message. This message activates the route.

Nodes monitor the link status of next hops on the multicast tree. When a link break on the multicast tree is detected, the tree branch should be immediately repaired through the use of the RREQ/RREP/MACT messages.

A multicast group leader is associated with each multicast group. The primary responsibility of this node is the initialization and maintenance of the group sequence number. A Group Hello message is periodically broadcast across the network by the multicast group leader. This message carries a multicast group and group sequence number and corresponding group leader IP address. This information is used for disseminating updated group sequence numbers throughout the multicast group and for repairing multicast trees after a previously disconnected portion of the network containing part of the multicast tree becomes reachable once again.

3. Route Tables

MAODV is a routing protocol, and it deals with route table management. Route table information must be kept even for ephemeral routes, such as are created to temporarily store reverse paths towards nodes originating RREQs. MAODV uses the following fields with each route table entry, which are based on the route table defined in the AODV Internet Draft [3]:

- Destination IP Address
- Destination Sequence Number
- Hop Count (number of hops needed to reach destination)
- Last Hop Count (described in subsection 9.2.1)
- Next Hop
- Next Hop Interface
- List of Precursors
- Lifetime (expiration or deletion time of the route)
- Routing Flags

The following information is stored in each entry of the multicast route table for multicast tree routes:

- Multicast Group IP Address
- Multicast Group Leader IP Address
- Multicast Group Sequence Number
- Next Hops

- Hop Count to next Multicast Group Member
- Hop Count to Multicast Group Leader

The Next Hops field is a linked list of structures, each of which contains the following fields:

- Next Hop IP Address
- Next Hop Interface
- Link Direction
- Activated Flag

The direction of the link is relative to the location of the group leader, i.e. UPSTREAM is a next hop towards the group leader, and DOWNSTREAM is a next hop away from the group leader. A node on the multicast tree must necessarily have only one UPSTREAM link. The IP Address of a Next Hop MUST NOT be used to forward multicast messages until after a MACT message has activated the route (see Section 9.6). The Next Hop Interface fields in the Route Table and Next Hop field of the Multicast Route Table are used for recording the outgoing interface on which the next hop can be reached (see Section 9.14).

Nodes MAY also maintain a Group Leader Table, which is an association between multicast groups and their corresponding group leaders. The fields of the group leader table are the following:

- Multicast Group IP Address
- Group Leader IP Address

This table can be used when discovering routes to multicast groups, as described in Section 9.2.

4. MAODV Terminology

This protocol specification uses conventional meanings [1] for capitalized words such as MUST, SHOULD, etc., to indicate requirement levels for various protocol features. This section defines other terminology used with MAODV that is not already defined in [2].

group leader

A node which is a member of the given multicast group and which is typically the first such group member in the connected portion of the network. This node is responsible for initializing and maintaining the multicast group destination sequence number.

group leader table

The table where ad hoc nodes keep information concerning each multicast group and its corresponding group leader. There is one entry in the table for each multicast group for which the node has received a Group Hello (see Section 9.2).

multicast tree

The tree containing all nodes which are members of the multicast group and all nodes which are needed to connect the multicast group members.

multicast route table

The table where ad hoc nodes keep routing (including next hops) information for various multicast groups.

reverse route

A route set up to forward a reply (RREP) packet back to the source from the destination or from an intermediate node having a route to the destination.

5. Route Request (RREQ) Message Format

```
    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      |J|R|G|       Reserved          |   Hop Count   |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |           Other fields as specified for AODV.......
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The format of the Route Request message remains as specified in [3], except that the following flags have been added:

J — Join flag; set when source node wants to join a multicast group.

R — Repair flag; set when a node wants to initiate a repair to connect two previously disconnected portions of the multicast tree.

When a node wishes to repair a multicast tree, it appends the Multicast Group Rebuild extension (see Section 10.2). When a node wishes to unicast the RREQ for a multicast group to the group leader, it includes the Multicast Group Leader extension (see Section 10.1).

6. Route Reply (RREP) Message Format

```
    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      |R|     Reserved        |Prefix Sz|   Hop Count   |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |           Other fields as specified for AODV.......
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The format of the Route Reply message remains as specified in [3], except that the following flag has been added:

R Repair flag; set when a node is responding to a repair request to connect two previously disconnected portions of the multicast tree.

When the RREP is sent for a multicast destination, the Multicast Group Information extension is appended (see Section 10.3).

7. Multicast Activation (MACT) Message Format

```
 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      |J|P|G|U|R|     Reserved        |   Hop Count   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  Multicast Group IP address                   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       Source IP address                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    Source Sequence Number                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The format of the Multicast Activation message is illustrated above, and contains the following fields:

Type 5

J Join flag; set when a node is joining the multicast group, as opposed to finding a route to the group for the transmission of data messages.

P Prune flag; set when a node wishes to prune itself from the tree, unset when the node is activating a tree link.

G Group Leader flag; set by a multicast tree member that fails to repair a multicast tree link breakage, and indicates to the group member receiving the message that it should become the new multicast group leader.

U — Update flag; set when a multicast tree member has repaired a broken tree link and is now a new distance from the group leader.

R — Reboot flag; set when a node has just rebooted (see Section 9.13).

Reserved — Sent as 0; ignored on reception.

Hop Count — The distance of the sending node from the multicast group leader. Used only when the 'U' flag is set; otherwise sent as 0.

Multicast Group IP Address
The IP address of the Multicast Group for which a route is supplied.

Source IP Address
The IP address of the sending node.

Source Sequence Number
The current sequence number for route information generated by the source of the route request.

To prune itself from the tree (i.e., inactivate its last link to the multicast tree), a multicast tree member sends a MACT with the 'P' flag = 1 to its next hop on the multicast tree. A multicast tree member that has more than one next hop to the multicast tree SHOULD NOT prune itself from the multicast tree.

8. Group Hello (GRPH) Message Format

```
 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      |U|M|        Reserved             |   Hop Count   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    Group Leader IP address                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  Multicast Group IP address                   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|               Multicast Group Sequence Number                 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The format of the Group Hello message is illustrated above, and contains the following fields:

Type
: 6

U
: Update flag; set when there has been a change in group leader information.

M
: Off_Mtree flag; set by a node receiving the group hello that is not on the multicast tree.

Reserved
: Sent as 0; ignored on reception.

Hop Count
: The number of hops the packet has traveled. Used by multicast tree nodes to update their distance from the group leader when the M flag is not set.

Group Leader IP Address
: The IP address of the group leader.

Multicast Group IP Address
: The IP address of the Multicast Group for which the sequence number is supplied.

Multicast Group Sequence Number
: The current sequence number of the multicast group.

9. Protocol Operation

This section describes the scenarios under which nodes generate control messages for multicast communication, and how the fields in the messages are handled.

9.1. Maintaining Multicast Tree Utilization Records

For each multicast tree to which a node belongs, either because it is a member of the group or because it is a router for the multicast tree, the node maintains a list of next hops -- i.e., those 1-hop neighbors that are likewise a part of the multicast tree. This list of next hops is used for forwarding messages received for the multicast group. A node forwards a multicast message to every such next hop, except that neighbor from which the message arrived. If there are multiple next hops, the forwarding operation MAY be performed by broadcasting the multicast packet to the node's neighbors; only the neighbors that belong to the multicast tree and have not already received the packet continue to forward the multicast packet.

9.2. Generating Route Requests

A node sends a RREQ either when it determines that it should be a part of a multicast group, and it is not already a member of that group, or when it has a message to send to the multicast group but does not have a route to that group. If the node wishes to join the multicast group, it sets the 'J' flag in the RREQ; otherwise, it leaves the flag unset. The destination address of the RREQ is always set to the multicast group address. If the node knows the group leader and has a route to it, the node MAY place the group leader's address in the Multicast Group Leader extension (Section 10.1), and unicast the RREQ to the corresponding next hop for that destination. Otherwise, if the node does not have a route to the group leader, or if it does not know who the multicast group leader is, it broadcasts the RREQ and does not include the extension field.

After transmitting the RREQ, the node waits for the reception of a RREP. The node may resend the RREQ up to RREQ_RETRIES additional times if a RREP is not received. If a RREQ was unicast to a group leader and a RREP is not received within RREP_WAIT_TIME milliseconds, the node broadcasts subsequent RREQs for that multicast group across the network. If a RREP is not received after RREQ_RETRIES additional requests, the node may assume that there are no other members of that particular group within the connected portion of the network. If it wanted to join the multicast group, it then becomes the multicast group leader for that multicast group and initializes the sequence number of the multicast group. Otherwise, if it only wanted to send packets to that group without actually joining the group, it drops the packets it had for that group and aborts the session.

When the node wishes to join or send a message to a multicast group, it first consults its Group Leader Table. Based on the existence of an entry for the multicast group in this table, the node then formulates and sends the RREQ as described at the beginning of this section.

9.2.1. Controlling Route Request broadcasts

To prevent unnecessary network-wide broadcasts of RREQs, the source node SHOULD use an expanding ring search technique as specified in [3]. When a RREP is received, the Hop Count to the group leader used in the RREP packet is remembered as Last Hop Count for that node in the routing table. When a new route to the same multicast group is required at a later time (e.g., upon a link break in the multicast tree), the TTL in the RREQ IP header is initially set to this Last Hop Count plus TTL_INCREMENT. Thereafter, following each timeout the TTL is incremented by TTL_INCREMENT until TTL = TTL_THRESHOLD is reached, as in [3].

9.3. Receiving Route Requests

When a node receives a RREQ, the node checks whether the 'J' flag of the RREQ is set. If the 'J' flag is set, the node can only respond if it is a member of the multicast tree for the indicated multicast group, and if its record of the multicast group sequence number is at least as great as that contained in the RREQ. If the 'J' flag is not set, then the node can respond if it has an unexpired route to the multicast group and the above multicast group sequence number criteria is met.

If the node does not meet either of these conditions, it rebroadcasts the RREQ from its interface(s) but using its own IP address in the IP header of the outgoing RREQ. The Destination Sequence Number in the RREQ is updated to the maximum of the existing Destination Sequence Number in the RREQ and the multicast group sequence number in the multicast route table (if an entry exists) of the current node. The TTL or hop limit field in the outgoing IP header is decreased by one. The Hop Count field in the broadcast RREQ message is incremented by one, to account for the new hop through the intermediate node.

The node always creates or updates a reverse route as specified in [3]. This reverse route would be needed in case the node receives an eventual RREP back to the node which originated the RREQ (identified by the Source IP Address).

In addition to creating or updating the route table entry for the source node, the node receiving the RREQ also creates a next hop entry for the multicast group in its multicast route table. If no entry exists for the multicast group, it creates one, and then places the node from which it received the RREQ as next hop for that group. It leaves the Activated flag associated with this next hop unset. The direction for this next hop entry is DOWNSTREAM.

9.4. Generating Route Replies

If a node receives a join RREQ for a multicast group, and it is already a member of the multicast tree for that group, the node updates its multicast route table and then generates a RREP message. The Source and Destination IP Addresses in RREQ message are copied to corresponding fields in the RREP message. The RREP contains the current sequence number for the multicast group and the IP address of the group leader. Furthermore, the node initializes the Hop Count field of the RREP to zero. Additional information about the multicast group is entered into the Multicast Group Information extension (see Section 10.3). It unicasts the RREP back to the node indicated by the Source IP Address field of the received RREQ.

A node can respond to a join RREQ only if it is a member of the multicast tree. If a node receives a multicast route request that is not a join message, it can reply if it has a current route to the multicast tree. Otherwise it continues forwarding the request. If a node receives a join RREQ for a multicast group and it is not already a member of the multicast tree for that group, it rebroadcasts the RREQ to its neighbors.

When a node receives a RREQ for a multicast group which has its own IP address in the Destination IP address of the IP header, that means that the source node expects this destination node to be the multicast group leader. In this case, if the node is in fact not the group leader, it can simply ignore the RREQ. The source node will time out after RREP_WAIT_TIME milliseconds and will broadcast a new RREQ without the group leader address specified.

Regardless of whether the multicast group leader or a multicast tree member generates the RREP, the RREP fields are set as follows:

Hop Count 0

Destination IP Address
The IP address of the multicast group.

Destination Sequence Number
The current multicast group sequence number.

Lifetime
The time for which nodes receiving the RREP consider the route to be valid (only used it the RREQ is not a join request).

The Multicast Group Information extension described in Section 10.3 is also included for join requests. If the node generating the RREP is not on the multicast tree (because the RREQ was not a join RREQ), it places its distance from the multicast tree in the Hop Count field, instead of 0.

9.5. Forwarding Route Replies

If an intermediate node receives a RREP in response to a RREQ that it has transmitted (or retransmitted on behalf of some other node), it creates a multicast group next hop entry for the node from which it received the RREP. The direction of this next hop is UPSTREAM, and the Activated flag is left unset. Additionally, the node updates the Lifetime field in the route table entry associated with the node from which it received the RREP. It then increments the Hop Count and Multicast Group Hop Count fields or the RREP and forwards this packet along the path to the source of the RREQ.

When the node receives more than one RREP for the same RREQ, it saves the route information with the greatest sequence number, and beyond that the lowest hop count; it discards all other RREPs. This node forwards the first RREP towards the source of the RREQ, and then forwards later RREPs only if they have a greater sequence number or smaller metric.

9.6. Route Activation

When a node broadcasts a RREQ message, it is likely to receive more than one reply since any node in the multicast tree can respond. The RREP message sets up route pointers as it travels back to the source node. If the request is a join request, these route pointers may eventually graft a branch onto the multicast tree. Also, because multicast data packets may be transmitted as broadcast traffic, the route to the multicast tree must be explicitly selected. Otherwise, each node with a route to the tree which receives a multicast data packet will rebroadcast the packet, resulting in an inefficient use of network bandwidth. Hence, it is necessary to activate only one of the routes created by the RREP messages. The RREP containing the largest destination sequence number is chosen to be the branch added to the multicast tree (or the path to the multicast tree, if the request was a non-join). In the event that a node receives more than one RREP with the same (largest) sequence number, it selects the first one with the smallest hop count, i.e., the shortest distance to a member of the multicast tree.

After waiting RREP_WAIT_TIME milliseconds, the node must select the route it wishes to use as its link to the multicast tree. This is accomplished by sending a Multicast Activation (MACT) message. The Destination IP Address field of the MACT packet is set to the IP address of the multicast group. If the node is joining the multicast group, it sets the join flag of this message. The node unicasts this message to the selected next hop, effectively activating the route. It then sets the Activated flag in the next hop Multicast Route Table entry associated with that node. After receiving this message, the node to which the MACT was sent activates the route entry for the link in its multicast route table, thereby finalizing the creation of the tree branch. All neighbors not receiving this message time out and delete that node as a next hop for the multicast group in their route tables, having never activated the route entry for that next hop.

Two scenarios exist for a neighboring node receiving the MACT message. If this node was previously a member of the multicast tree, it does not propagate the MACT message any further. However, if the next hop selected by the source node's MACT message was not previously a multicast tree member, it will have propagated the original RREQ further up the network in search of nodes which are tree members. Thus it is possible that this node also received more than one RREP, as noted in Section 9.5.

When the node receives a MACT selecting it as the next hop, it unicasts its own MACT to the node it has chosen as its next hop, and so on up the tree, until a node which was already a part of the multicast tree is reached.

9.7. Multicast Tree Pruning

A multicast group member can revoke its member status at any time. However, it can only actually leave the multicast tree if it is not a tree router for any other nodes in the multicast group (i.e., if it is a leaf node). If a node wishing to leave the multicast group is a leaf node, it unicasts to its next hop on the tree a MACT message with the 'P' flag set and with the Destination IP Address set to the IP address of the multicast group. It then deletes the multicast group information for that group from its multicast route table. When its next hop receives this message, it deletes the sending node's information from its list of next hops for the multicast tree. If the removal of the sending node causes this node to become a leaf node, and if this node is also not a member of the multicast group, it may in turn prune itself by sending its own MACT message up the tree.

When the multicast group leader wishes to leave the multicast group, it proceeds in a manner similar to the one just described. If it is a leaf node, it may leave the group and unicast a prune message to its next hop. The next hop acts in the manner described in

Section 9.10, since the prune message is coming from its upstream neighbor. Otherwise, if the group leader is not a leaf node, it may not prune itself from the tree. It takes the actions described in Section 9.9, where it selects one of its next hops and unicasts to it the MACT with set 'G' flag.

9.8. Repairing a Broken Link

Branches of the multicast tree become invalid if a broken link results in an infinite metric being associated with the next hop route table entry. When a broken link is detected between two nodes on the multicast tree, the two nodes should delete the link from their list of next hops for the multicast group. The node downstream of the break (i.e., the node which is further from the multicast group leader) is responsible for initiating the repair of the broken link. In order to repair the tree, the downstream node broadcasts a RREQ with destination IP address set to the IP address of the multicast group and with the 'J' flag set. The destination sequence number of the RREQ is the last known sequence number of the multicast group. The node also includes the Multicast Group Leader Extension. The Multicast Group Hop Count field of this extension is set to the distance of the source node from the multicast group leader. A node MUST have a hop count to the multicast group leader less than or equal to the indicated value in order to respond. This hop count requirement prevents nodes on the same side of the break as the node initiating the repair from replying to the RREQ.

The RREQ is broadcast using an expanding rings search, as described in Section 9.2.1. Because of the high probability that other nearby nodes can be used to rebuild the route, the original RREQ is broadcast with a TTL (time to live) field value equal to TTL_INCREMENT plus the Multicast Group Hop Count. In this way, the effects of the link breakage may be localized. If no reply is received within RREP_WAIT_TIME milliseconds, the node SHOULD increment the TTL of each successive RREQ by TTL_INCREMENT, until either a route is determined, or TTL_THRESHOLD is reached. After this point, if a route is still not discovered, each additional RREQ is broadcast with TTL = NET_DIAMETER. Up to RREQ_RETRIES additional broadcasts may be attempted after TTL = NET_DIAMETER is reached.

A node receiving this RREQ can respond if it meets the following conditions:

- It is a member of the multicast tree.

- Its record of the multicast group sequence number is at least as great as that contained in the RREQ.

- Its hop count to the multicast group leader is less than or equal to the contained in the Multicast Group Hop Count extension field.

If the originating node receives more than one RREP, route activation occurs as described in the previous section.

At the end of the discovery period, the node selects its next hop and unicasts a MACT message to that node to activate the link, as described in Section 9.6. Since the node was repairing a tree break, it is likely that it is now a different distance from the group leader than it was before the break. If this is the case, it must inform its DOWNSTREAM next hops of their new distance from the group leader. It does this by broadcasting a MACT message with the 'U' flag set, and the Hop Count field set to the node's new distance from the group leader. This 'U' flag indicates that multicast tree nodes should update their distance from the group leader. When a node on the multicast tree receives the MACT message with the 'U' flag set, it determines whether this packet arrived from its UPSTREAM neighbor. If it did not, the node discards the packet. Otherwise, it increments the Hop Count value contained in the MACT packet and updates it distance to the group leader to be this node value. If this node has one or more downstream next hops, it in turn must send a MACT message with a set 'U' flag to its next hops, and so on.

When a link break occurs, it is possible that the tree will be repaired through different intermediate nodes. If the node UPSTREAM of the break is not a group member, and if the loss of that link causes it to become a leaf node, it sets a prune timer to wait for the link to be repaired. This PRUNE_TIMEOUT should be larger than RREP_WAIT_TIMEOUT to give the link time to be repaired. If, when this timer expires, the node has not received a MACT message selecting it to be a part of the repaired tree branch, it prunes itself from the tree by sending a MACT with set 'P' flag to its next hop, as previously described.

9.9. Tree Partitions

It is possible that after a link breaks, the tree cannot be repaired due to a network partition. If the node attempting to repair a tree link break does not receive a response after RREQ_RETRIES attempts, it can be assumed that the network has become partitioned and the multicast tree cannot be repaired at this time. In this situation, if the node which initiated the route rebuilding is a multicast group member, it becomes the new multicast group leader for its part of the multicast tree partition. It increments the group sequence number and then broadcasts a Group Hello (GRPH) for this multicast group. The 'U' flag in the GRPH is set, indicating that there has been a change in the group leader information. All nodes receiving this message update their group leader table to indicate the new group leader information. Nodes which are a part of the multicast tree also update the group leader and sequence number information for that group in their multicast route table.

On the other hand, if the node which had initiated the repair is not a multicast group member, there are two possibilities. If it only has one next hop for the multicast tree, it prunes itself from the tree by unicasting a MACT message, with the 'P' flag set, to its next hop. The node receiving this message notes that the message came from its upstream link, i.e., from a node that is closer to the group leader than it is. If the node receiving this message is a multicast group member, it becomes the new group leader and broadcasts a GRPH message as indicated above. Otherwise, if it is not a multicast group member and it only has one other next hop link, it similarly prunes itself from the tree. This process continues until a multicast group member is reached.

The other possibility is that the node which initiated the rebuilding is not a group member and has more than one next hop for the tree. In this case, it cannot prune itself, since doing so would partition the tree. It instead selects one of its next hops and unicasts a MACT with the 'G' flag set to that node. This flag indicates that the next group member to receive this message should become the new group leader. It then changes the direction of that link to be

UPSTREAM. If the node receiving the MACT is a group member, then this node becomes the new group leader. Otherwise, the node unicasts its own MACT message with the 'G' flag set to one of its next hops, and changes the direction of that link. Once a group member is reached, the new group leader is determined.

9.10. Reconnecting Two Trees

In the event that a link break cannot be repaired, the multicast tree remains partitioned until the two parts of the network become connected once again. A node from one partition of the network knows that it has come into contact with a node from the other partition of the network by noting the difference in the GRPH message multicast group leader information. The multicast group leader with the lower IP address initiates the tree repair. For the purposes of this explanation, call this node GL1. GL1 unicasts a RREQ with both the 'J' and 'R' flags set to the group leader of the other network partition (GL2), using the node from which it received the GRPH as the next hop. This RREQ contains the current value of GL1's multicast group sequence number. If any node that receives the RREQ is a member of GL2's multicast tree, it MUST forward the RREQ along its upstream link, i.e. towards GL2. This prevents any loops from being formed after the repair. Upon receiving the RREQ, GL2 takes the larger of its and the received multicast group sequence number, increments this value by one, and responds with a RREP. This is the group leader which becomes the leader of the reconnected multicast tree. The 'R' flag of the RREP is set, indicating that this RREP is in response to a repair request.

As the RREP is propagated back to GL1, nodes add the incoming and outgoing links to the multicast route table next hop entries if these entries do not already exist. The nodes also activate these entries, thereby adding the branch on to the multicast tree. If a node that was previously a member of GL1's tree receives the RREP, it MUST forward the packet along its link to its previous group leader (GL1). It then updates its group leader information to reflect GL2 as the new group leader, changes the direction of the next hop link associated with GL1 to DOWNSTREAM, and sets the direction of the link on which it received the RREP to UPSTREAM. When GL1 receives the RREP, it updates its group leader information and sets the link from which it received the RREP as its upstream link. The tree is now reconnected. The next time GL2 broadcasts a GRPH, it sets the 'U' flag to indicate that there is a change in the group leader information and group members should update the corresponding information. All network nodes update their group leader table to reflect the new group leader information.

9.11. Maintaining Local Connectivity

Each node on the multicast tree MUST keep track of its neighbors which are next hops on the multicast tree. This is done as specified in [3]. If a link to the next hop cannot be detected by any of the methods specified therein, the forwarding node MUST assume that the link is broken, and take corrective action by following the methods specified in Section 9.8.

Whenever the multicast tree is used for the transmission of data packets, a node on the tree must hear from each of its next hop neighbors every RETRANSMIT_TIME. The node can hear from its neighbors through one of the methods previously described, or through the reception of a rebroadcast of a data packet from these next hops, if the data packets are being transmitted by broadcast. If a node does not hear from one of its multicast tree next hops within RETRANSMIT_TIME, and if the multicast tree is actively being used to transmit data packets, then the node MUST assume the link to its next hop is broken, and proceed as described in Section 9.8. Decreasing the amount of time between when a node must hear from its next hops from ALLOWED_HELLO_LOSS * HELLO_INTERVAL milliseconds to RETRANSMIT_TIME allows for the quickest detection of broken links and fewer data packet losses.

9.12. Group Hello Messages

If a node sends a RREQ to join a multicast group ('J' flag set) and after RREQ_RETRIES attempts does not receives a response, it then becomes the multicast group leader. The node initializes the multicast group sequence number and then broadcasts a GRPH message to inform network nodes that it is now the group leader for the multicast group. To ensure nodes maintain consistent and up-to-date information about who the multicast group leaders are, any node which is a group leader for a multicast group broadcasts such a GRPH across the network every GROUP_HELLO_INTERVAL milliseconds. The contents of the GRPH fields are set as follows:

U Flag 0

M Flag 0

Hop Count 0

Group Leader IP Address
The IP Address of the group leader.

Multicast Group IP Address
The IP Address of the Multicast Group for which the node is the group leader.

Multicast Group Sequence Number
One plus the last known sequence number of the multicast group.

Nodes receiving the GRPH increment the Hop Count field by one before forwarding the message. When a node not on the multicast tree receives the GRPH message, it sets the 'M' flag. This indicates that this incarnation of the message has traveled off the multicast tree, and hence cannot be used by group members to verify their distance from the group leader. The 'U' flag is set by the group leader whenever there has been a change in group leader information. It informs nodes that they should update the group leader information associated with the indicated multicast group.

A node receiving the GRPH message first verifies that is has not already received the GRPH message with the same group IP address/group sequence number pair in the last BCAST_ID_SAVE milliseconds. If it has, it silently discards the message. Otherwise, the node updates its group leader table to reflect the current multicast group IP address/group leader IP address combination, and increments the Hop Count value in the message. If the node receiving the GRPH message is a member of the multicast tree, it also updates this information in its multicast route table. If the 'M' flag is unset, then this node can also use the Hop Count value to verify its current distance from the group leader.

After processing the GRPH message, the node buffers the group IP address/ group sequence number combination to avoid multiple processings of the same GRPH message. It then rebroadcasts the message to its neighbors.

9.13. Actions After Reboot

A node participating in the multicast tree that reboots (or restarts the routing daemon) loses all of its multicast tree information. Upon reboot, the rebooted node does not know whether it was previously a member of the multicast tree. Hence, it should unconditionally broadcast a MACT message with set Reboot ('R') flag to inform neighboring nodes that it has lost its multicast group information. When a node on the multicast tree receives the reboot MACT message, it checks whether this message came from one of its next hops on the multicast tree. If so, one of two situations exists.

If the reboot MACT came from a downstream link, the node deletes that link from its list of next hops and sets a prune timer according to the guidelines in Section 9.8. Otherwise, if the reboot MACT came from a node's upstream link, it must rebuild the tree branch as is also indicated in Section 9.8.

9.14. Interfaces

Because MAODV should operate smoothly over wired, as well as wireless, networks, and because it is likely that MAODV will also be used with multi-homed radios, the interface over which packets arrive must be known to MAODV whenever a packet is received. This includes the reception of RREQ, RREP, MACT, and GRPH messages. Whenever a packet is received from a new neighbor, the interface on which that packet was received is recorded into the route table entry for that neighbor, along with all the other appropriate routing information. Similarly, whenever a route to a new destination is learned, the interface through which the destination can be reached is also recorded into the destination's route table entry.

When multiple interfaces are available, a node receiving and rebroadcasting a RREQ or GRPH message rebroadcasts that message on all interfaces except the one on which it was received. If a node is participating on the multicast tree, and if multicast data packets are being sent as broadcast traffic, a multicast tree node should only rebroadcast a packet on those interfaces which have multicast tree next hops.

10. Extensions

This section specifies for multicast extensions for RREQ and RREP messages, conforming to the format defined in [3].

10.1. Multicast Group Leader Extension Format

This extension is appended to a RREQ by a node wishing to repair a multicast tree.

```
 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      |    Length     | Multicast Group Leader IP ... |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   ... Address (continued)     | Previous Hop IP Address ...   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   ... (continued)             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Type 3

Length 8

Multicast Group Leader IP Address
The IP Address of the Multicast Group Leader.

Previous Hop IP Address
The IP Address of the node which previously received the RREQ. This field is used when the RREQ is unicast to the group leader when a node wishes to join a multicast group.

This extension is used when unicasting the RREQ to the group leader. Each node receiving the RREQ updates the Previous Hop IP Address field to reflect its address.

10.2. Multicast Group Rebuild Extension Format

This extension is appended to a RREQ by a node wishing to repair a multicast tree.

```
 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      |    Length     |  Multicast Group Hop Count    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Type 4

Length 2

Multicast Group Hop Count
: The distance in hops of the node sending the RREQ from the Multicast Group Leader.

This extension is used for rebuilding a multicast tree branch. It is used to ensure that only nodes as least as close to the group leader as indicated by the Multicast Group Hop Count field respond to the request.

10.3. Multicast Group Information Extension Format

The following extension is used to carry additional information for the RREP message (see Section 6) when sent to establish a route to a multicast destination.

```
 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      |    Length     |   Multicast Group Hop Count   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                Multicast Group Leader IP Address              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Type 5

Length 6

Multicast Group Hop Count
The distance of the node from the Multicast Group Leader.

Multicast Group Leader IP Address
The IP Address of the current Multicast Group Leader.

This extension is included when responding to a RREQ to join a multicast group. The node responding to the RREQ places its distance from the group leader in the Multicast Group Hop Count field.

11. Configuration Parameters

This section gives default values for some important values associated with MAODV protocol operations. A particular mobile node may wish to change certain of the parameters, in particular the NET_DIAMETER, MY_ROUTE_TIMEOUT, ALLOWED_HELLO_LOSS, RREQ_RETRIES, and possibly the HELLO_INTERVAL. In the latter case, the node should advertise the HELLO_INTERVAL in its Hello messages, by appending a Hello Interval Extension to the RREP message. Choice of these parameters may affect the performance of the protocol.

Multicast Parameters	Value
GROUP_HELLO_INTERVAL	5,000 Milliseconds
MTREE_BUILD	2 * REV_ROUTE_LIFE
PRUNE_TIMEOUT	ACTIVE_ROUTE_TIMEOUT
RETRANSMIT_TIME	750 Milliseconds

Additional AODV Parameters
ACTIVE_ROUTE_TIMEOUT
ALLOWED_HELLO_LOSS
BCAST_ID_SAVE
DELETE_PERIOD
HELLO_INTERVAL
NET_DIAMETER
NEXT_HOP_WAIT
NODE_TRAVERSAL_TIME
REV_ROUTE_LIFE
RREP_WAIT_TIME
RREQ_RETRIES
TTL_START
TTL_INCREMENT
TTL_THRESHOLD

Please see [3] for a discussion of and the appropriate values for the Additional AODV Parameters.

12. Security Considerations

Currently, MAODV does not specify any special security measures. Route protocols, however, are prime targets for impersonation attacks, and must be protected by use of authentication techniques involving generation of unforgeable and cryptographically strong message digests or digital signatures. It is expected that, in environments where security is an issue, that IPSec authentication headers will be deployed along with the necessary key management to distribute keys to the members of the ad hoc network using MAODV.

References

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

[2] C. Perkins. Mobile Ad Hoc Networking Terminology. IETF Internet Draft, draft-ietf-manet-term-00.txt, October 1997. (Work in Progress).

[3] C. E. Perkins, E. M. Royer, and S. R. Das. Ad Hoc On Demand Distance Vector (AODV) Routing. IETF Internet Draft, draft-ietf-manet-aodv-06.txt, July 2000. (Work in Progress).

Author's Addresses

Questions about this memo can be directed to:

Elizabeth M. Royer
Dept. of Electrical and Computer Engineering
University of California, Santa Barbara
Santa Barbara, CA 93106
+1 805 893 7788
+1 805 893 3262 (fax)
eroyer@alpha.ece.ucsb.edu

Charles E. Perkins
Communications Systems Laboratory
Nokia Research Center
313 Fairchild Drive
Mountain View, CA 94303
USA
+1 650 625 2986
+1 650 691 2170 (fax)
charliep@iprg.nokia.com