

Internet-Draft

HOLSR

July 2009

Abstract

This document describes the Hierarchical Optimized Link State Routing (HOLSR) mechanism for heterogeneous mobile ad hoc networks. In this specification a heterogeneous mobile ad hoc network is defined as a network of mobile routers that are characterized by different communication capabilities, such as communication channels, processing powers or energy levels.

The HOLSR mechanism is an extension to the OLSRv2 protocol. HOLSR takes advantage of the router's distinct communications capabilities to reduce the routing control overhead in large heterogeneous ad hoc networks, thus improving the performance of the routing mechanism. More precisely, HOLSR defines a hierarchy in the network and presents a routing scheme for this hierarchical structure with a better scalability.

Internet-Draft

HOLSR

July 2009

Table of Contents

1. Introduction	5
2. Terminology	5
3. Applicability Statement	7
4. Protocol Overview and Functioning	8
4.1. Hierarchy Building	8
4.2. Hierarchy Routing	11
5. Protocol Parameters	11
5.1. Message Intervals	11
5.1.1. CIA Messages	12
5.1.2. HTC Messages	12
5.2. Information Validity Times	13
5.2.1. CIA Validity Times	13
5.2.2. HTC Validity Times	13
5.3. Cluster Intersection Depth	14
5.4. Cluster Critical Path	14
6. Protocol Details	14
6.1. Description of HOLSR Logical Topology Levels	15
6.1.1. Example	15
6.2. Cluster Formation	16
6.2.1. Particular Case: Cluster Intersections	17
6.2.2. Cluster ID Announcement Validity Timer	17
6.3. Cluster Interaction	18
6.3.1. Types of Hierarchical Topology Control Message	18
6.3.2. HTC and TC Message Propagation	19
6.4. Clusters Decommission	19
6.5. HOLSR Repositories	20
6.6. HOLSR Messages	21
6.6.1. CIA Messages	21
6.6.1.1. CIA Message TLVs	22
6.6.2. HTC Messages	22
6.6.2.1. HTC Message TLVs	23
6.7. HOLSR Algorithm	23
6.7.1. Processing and Broadcasting CIA Message	23
6.7.2. Processing and Forwarding HTC Message	24
6.7.3. Cluster Intersection Handling	25
7. Proposed Values for Parameters	26
7.1. Message Intervals Parameter Defaults	26
7.1.1. CIA Messages	26
7.1.2. HTC Messages	27
7.2. Information Validity Times Parameters	27
7.2.1. CIA Messages	27
7.2.2. HTC Messages	27
7.3. Cluster Intersection Depth	27
7.4. Cluster Critical Path	27
8. Contributors	27
9. Acknowledgements	28

Internet-Draft

HOLSR

July 2009

10. IANA Considerations	28
10.1. Message Types	28
10.2. TLV Types	29
10.3. HTC_MSG_TYPE Values	29
11. Security Considerations	30
12. References	30
12.1. Normative References	30
12.2. Informative References	30
Appendix A. Example of Data Transfer using Clusters	30
Appendix B. Example of Cluster Decommission Handling	33
Appendix C. Packet and Message Layout	35
C.1. Packet and Message Options	35
C.2. Example CIA Message	35
C.3. Example HTC Message	36
Authors' Addresses	37

Internet-Draft

HOLSR

July 2009

1. Introduction

The Hierarchical Optimized Link State Routing Protocol (HOLSR) is designed for heterogeneous ad hoc networks, and is an extension to the OLSRv2 [I-D.ietf-manet-olsrv2] routing protocol. HOLSR has two main functionalities:

- o Hierarchy building: HOLSR organizes routers into hierarchical levels according to their capacities (e.g. radio, energy, processing) and dynamically groups routers into clusters at each level.
- o Hierarchical routing: HOLSR provides a hierarchical routing that supports random movement of the routers among and between clusters, and can dynamically adapt to topology changes at different levels. Within each cluster, the OLSRv2 routing protocol, as specified in [I-D.ietf-manet-olsrv2], is used.

The main improvements realized by the HOLSR protocol are a reduction in the amount of topology control information that needs to be exchanged at different levels of the hierarchical network topology, and the efficient use of high capacity routers. Another significant benefit is a reduction in routing computational cost: if a link in one part of the network is broken, only those routers within the cluster in which the link exists need to re-calculate the routing table, while routers in other clusters are not affected. HOLSR is versatile in that it does not require a logical hierarchical addressing scheme but can accommodate one if required.

In an HOLSR network structure, the routers equipped with high capacity links form the higher-level network, while the routers with low capacity links form the lower-level network. At a given level, cluster heads are selected from the routers with higher capabilities. The choice of the cluster head or the establishment of its capability criteria (e.g. multiple interfaces, more processing power, increased memory, etc.), is out of scope of this document. It is left to the implementer or the administrator, who can select the best criteria adapted to the specific network deployment. Routers can move from one cluster to another, and can associate with the nearest cluster head, while the cluster heads propagate updated membership information among each other using Hierarchical Topology Control (HTC) messages.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and

Internet-Draft

HOLSR

July 2009

"OPTIONAL" in this document are to be interpreted as described in [RFC2119].

This document uses OLSR terminology as defined in OLSRv1 [RFC3626] and OLSRv2 [I-D.ietf-manet-olsrv2]. MANET specific terminology is to be interpreted as described in [RFC5444] and NHDP [I-D.ietf-manet-nhdp].

Additionally, this document uses the following terminology:

Cluster

A small group or gathering of MANET routers that have the same cluster head.

Cluster Head

A MANET router can become a cluster head at level *i* if:

- * it possesses a "higher capability" besides the one it shares with its cluster neighbor routers, and
- * it participates in communication level *i*+1.

The definition of "higher capability" is a deployment issue, out of scope of this document.

Cluster Member

A MANET router that is associated with a cluster head.

CIA

Cluster Id Announcement. A control message initiated periodically by a cluster head to declare its leadership and invite other routers to join its cluster. Cluster members generate CIA messages including a distance field set to their distance (in terms of number of hops) to the cluster head). The message is generated by cluster members until it reaches the edge of the cluster.

Cluster Edge

Cluster Edge. Distinguishes routers belonging to the cluster from others.

Cluster Intersection Depth

When routers from one cluster can hear Cluster Id Announcements from another cluster, it is said that these two clusters intersect, i.e. a router on the edge of one cluster is in radio communication range from the edge of a neighbor cluster. The depth of the intersection depends on the permissions given to edge routers from different clusters to inter-communicate without going

Internet-Draft

HOLSR

July 2009

through (and bypassing) their respective cluster heads. For a depth value of '1', edge routers from two different clusters within radio range can communicate, but they cannot forward messages from each other within their own cluster. For a depth value of '2', the same two edge routers are permitted to forward neighbor cluster's messages once (Time to live (TTL) = '1'). For a depth value of '3', TTL = '2', etc.

HTC

Hierarchical Topology Control (TC). A control message sent periodically that is used to propagate the membership information of a cluster to the higher hierarchical level routers.

OLSRv2

The Optimized Link State Routing protocol version 2 is an update to OLSRv1 as published in [RFC3626].

Router

A MANET router which implements the Hierarchical OLSRv2 protocol as specified in this document.

Topology level

A logical level that is composed of links and routers that share similar characteristics (e.g. bandwidth, transmission range, energy capacity, processing capacity). The higher the level gets, the "better" the characteristics are.

3. Applicability Statement

Many contemporary ad hoc wireless networks are heterogeneous in nature, being comprised of mobile devices with heterogeneous capacities. For instance, these devices can have distinct communications capabilities with respect to data rate, radio range, frequency band, etc. They can also have different energy capacities and different processing capacities. We can cite two examples of networks where HOLSR provides high benefits: military networks and sensor networks.

In military networks for instance, soldiers, tanks and headquarters might each be given wireless communication equipment that is appropriate to their communication needs. Soldiers are usually equipped with wireless communication devices characterized by limited resources in terms of radio bandwidth, energy capacity and processing capacity. Those devices can only handle limited transmission range and have restricted communications bandwidth. Vehicles, on the other hand, are outfitted with more powerful equipments providing extended communication coverage with higher communication bandwidth

Internet-Draft

HOLSR

July 2009

capability.

In data gathering applications supported by wireless sensor networks, the aim is to maximize network lifetime. Energy efficient strategies are used such as clustering. Cluster heads are selected among the routers with a sufficient energy level, and rotate in order to balance energy consumption. Cluster heads can also be chosen with an additional applicative criterion: they represent routers having the same data value to transmit to the data sink.

HOLSR is also beneficial to mesh networks in general, where there can be limited or no mobility such as sensor networks. HOLSR can manage a combination of mobile and non-mobile components.

Scalability is one of the most important factors governing the efficiencies of heterogeneous wireless networks. Scalability may be defined as the ability of a network to adjust or maintain its performance when its size increases. That also includes the increase in traffic load that is handled. Yet under a "flat" routing protocol, the performance of an ad hoc network tends to degrade as the number of mobile routers increases. When a flat routing protocol is used, the resulting control overhead increases, depending on the number of interfaces possessed by each router since more links need to be propagated for each router in order to share the topology information. Concerning data gathering applications, the use of hierarchical routing enables reducing the amount of data transferred to the sink, and hence to maximize network lifetime.

4. Protocol Overview and Functioning

HOLSR consists of 2 main functionalities: enabling the creation of a hierarchy structure on top of a standard OLSR network; and providing the control messaging tools to transmit the hierarchy topology and enable routing through the use of gateway routers, defined as cluster heads.

The criteria used to define and select cluster heads unto which to construct the HOLSR hierarchical structure is outside the scope of this protocol.

4.1. Hierarchy Building

In HOLSR, a hierarchical structure is constructed over the network based on its member routers' capabilities. In the following description, the criteria selected for higher capabilities is linked to the amount and properties of each router's wireless interfaces.

Internet-Draft

HOLSR

July 2009

Let us demonstrate the idea using figures 1 and 2, composed of 23 MANET routers. In figure 1, routers are denoted by letters, and their preceding character determines their capabilities and available interfaces. Thus a "." identifies regular MANET routers, a "#" identifies routers with 2 interfaces ("low" and "medium" range radio); a "&" identifies routers with 3 interfaces (low, medium, and "long" range radio); and a "@" identifies routers with 2 interfaces (medium and long range radios).

For this example, it is decided that the routers with multiple interfaces be used as cluster heads. From then on, MANET routers start organizing themselves into clusters around their cluster heads. There are 4 cluster heads (a, f, p, and t) that have dual low and medium transmission range interfaces. These form 4 intersecting clusters at level 1 (see figure 2). Routers are identified with their cluster id (capital character followed by a number, or 'H' for a cluster head), cluster A members are now AH, A1, A2, and A3 (respectively a, b, c and d), etc. Clusters at level 2 are formed by routers that possess at least one interface with medium transmission range. This requirement is matched by routers a, f, p, r, and u. From these routers, the ones that contain an additional long transmission range radio interface (level 3) become cluster heads at level 2 (this applies to routers f and r). Hence 2 clusters are created at level 2 (E and F with cluster heads EH and FH). The level 3 "overall" cluster is formed by routers f and r, since they share the long transmission range capabilities. No cluster head is required at level 3 since 1) there is no higher level; 2) the 2 members are in transmission range of each other.

Internet-Draft

HOLSR

July 2009




Figure 1 - Typical MANET Example

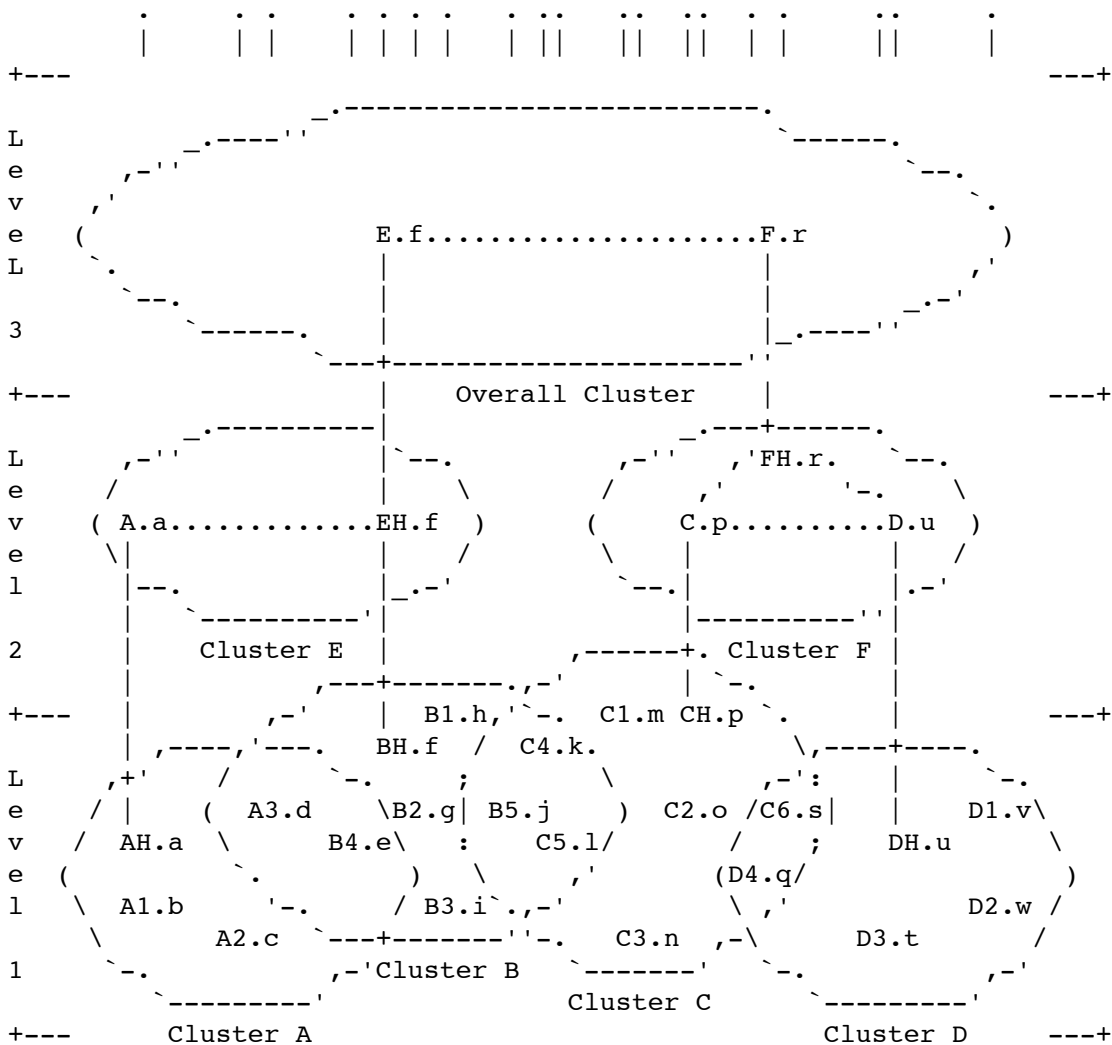


Figure 2 - Hierarchical MANET Example

Internet-Draft

HOLSR

July 2009

Cluster formation of the hierarchy structure is made possible by the addition of Cluster Id Announcement (CIA) messages. Every router selected as a cluster head sends CIA messages periodically to declare its leadership and invite other routers to join its cluster. Section 6 provides complete details.

4.2. Hierarchy Routing

In HOLSR, at a given level, routers can move from one cluster to another and associate with the nearest cluster head. For routing purposes, the information about cluster membership is transmitted to higher levels by the cluster heads. The following denotes the routing structure within HOLSR:

- o Within each cluster, OLSRV2 [I-D.ietf-manet-olsrv2] runs unaltered, providing routing topology information about router destinations within the same cluster. TC message propagation is limited within the local cluster.
- o Between clusters, the information about cluster membership and topology information is transmitted to higher levels by the cluster heads using Hierarchical Topology Control (HTC) messages.
- o For routers located in overlapping regions of multiple clusters, OLSRV2 TC messages can (with restrictions) be propagated not only within the cluster but to the neighbor clusters as well. This feature is controlled by the C_INTER_DEPTH parameter defined in Section 5.3. Section 6.7.3 provides complete details.
- o Traffic from one cluster with destination to another cluster will travel up the levels through cluster heads (acting as gateways) and then down to reach their destination. Routers set their cluster head as default gateway in their routing table. Hence, every traffic request for destinations outside of their cluster will go through their cluster head.

5. Protocol Parameters

The parameters used in this specification are those defined in OLSRV2 [I-D.ietf-manet-olsrv2] plus those described in this section.

5.1. Message Intervals

The following parameters regulate CIA and HTC message transmissions by a router.

Internet-Draft

HOLSR

July 2009

5.1.1. CIA Messages

CIA messages are transmitted periodically to advertise a router's cluster head id and its distance from that cluster head. The frequency of these advertisements is regulated by the parameters CIA_INTERVAL and CIA_MIN_INTERVAL.

Specifically, these parameters are as follows:

CIA_INTERVAL

the maximum time between the transmission of two successive CIA messages by a given cluster head. If using periodic transmission of CIA messages for a given cluster head, these SHOULD be at a separation of CIA_INTERVAL, and MAY be modified by jitter as specified in [RFC5148].

CIA_MIN_INTERVAL

the minimum interval between transmission of two successive CIA messages by a given cluster head. (This minimum interval MAY be modified by jitter, as defined in [RFC5148].)

The following constraints apply to these parameters:

- o CIA_INTERVAL > 0
- o CIA_MIN_INTERVAL >= 0
- o CIA_INTERVAL >= CIA_MIN_INTERVAL
- o If INTERVAL_TIME message TLVs as defined in [RFC5497] are included in CIA messages, then CIA_INTERVAL MUST be representable as described in [RFC5497].

5.1.2. HTC Messages

HTC messages are transmitted periodically by a cluster head to advertise its membership information to routers at the higher hierarchical level. The frequency of these advertisements is regulated by the parameters HTC_INTERVAL and HTC_MIN_INTERVAL. HTC messages are an extension to the TC messages from OLSRv2 [I-D.ietf-manet-olsrv2], and as such they MAY use the same message interval parameters, and the same message validity times:

HTC_INTERVAL

the maximum time between the transmission of two successive HTC messages by a cluster head. When no HTC messages are sent in response to local network changes under the cluster head, then HTC messages SHOULD be sent at a regular interval HTC_INTERVAL,

Internet-Draft

HOLSR

July 2009

possibly modified by jitter as specified in [RFC5148].

HTC_MIN_INTERVAL

the minimum interval between transmission of two successive HTC messages by a cluster head. (This minimum interval MAY be modified by jitter, as defined in [RFC5148].)

The following constraints apply to these parameters:

- o HTC_INTERVAL > 0
- o HTC_MIN_INTERVAL >= 0
- o HTC_INTERVAL >= HTC_MIN_INTERVAL
- o If INTERVAL_TIME message TLVs as defined in [RFC5497] are included in HTC messages, then HTC_INTERVAL MUST be representable as described in [RFC5497].

5.2. Information Validity Times

5.2.1. CIA Validity Times

The following parameters manage the validity time of CIA information:

CIA_HOLD_TIME

the value in the VALIDITY_TIME message TLV included in the CIA message.

The following constraints apply to this parameter:

- o CIA_HOLD_TIME >= CIA_INTERVAL
- o CIA_HOLD_TIME MUST be representable as described in [RFC5497].

5.2.2. HTC Validity Times

HOLSR HTC messages SHOULD use the same validity time parameters as for TC messages in OLSRv2 [I-D.ietf-manet-olsrv2].

OLSRv2 TC parameters for validity times are (refer to OLSRv2 [I-D.ietf-manet-olsrv2] for definitions):

- o T_HOLD_TIME
- o A_HOLD_TIME

Internet-Draft

HOLSR

July 2009

- o RX_HOLD_TIME
- o P_HOLD_TIME
- o F_HOLD_TIME

5.3. Cluster Intersection Depth

The following parameter regulates the communication between clusters that are intersecting. This parameter has a direct impact on the decision to route traffic through cluster heads are directly between clusters at the same level.

C_INTER_DEPTH

the maximum depth that routers from different clusters in an intersected area can include routers from another cluster than their own in their routing tables (routers can only select one cluster head ID).

The following constraints apply to this parameter:

C_INTER_DEPTH >= 0

5.4. Cluster Critical Path

The following parameter is used to identify cluster heads located on critical paths on the HOLSR architecture. A router is considered to be on a critical path if its failure causes network partition.

C_CRITICAL_PATH

- identifies a cluster head as being on the critical path of the hierarchical structure.

The following constraints apply to this parameter:

C_CRITICAL_PATH = [0,1], where 0 is 'false', and 1 is 'true'

6. Protocol Details

This section describes the HOLSR routing protocol by specifying:

- o HOLSR logical topology levels
- o Cluster formation
- o Cluster interaction

Internet-Draft

HOLSR

July 2009

- o Cluster decommission
- o HOLSR repositories
- o HOLSR Algorithm

6.1. Description of HOLSR Logical Topology Levels

Based on different communication capacities, routers are organized into logical topology levels. A logical topology level is constituted by routers having similar capacities (e.g. radio bandwidth, radio interface, energy capacity, processing capacity). The grouping and ordering of routers with similar capacities (to determine which to assign to higher or lower layers) is an administrator and deployment issue. Since one router can have multiple communication capacities (by possessing different wireless interfaces, for example), it can be visible as a router in different logical topology levels. Usually, low-power routers are equipped with only one interface offering limited data rate and transmission range. Such routers participate at the lowest topology Level 1.

Routers at topology level $i+1$, with i greater than or equal to 1, have higher capacities than routers at level i . Any router at level i must be able to communicate, via a multi-hop path if necessary, with a router at level i that is able to communicate with a router at level $i+1$. In the same manner, if the level $i+2$ exists, then any router at level $i+1$ must be able to communicate with a router at level $i+1$ that is able to communicate with a router at level $i+2$.

6.1.1. Example

An example of a three levels hierarchy can be described as follows. Routers at topology level 2 are equipped with a wireless interface that can provide a higher data rate and longer transmission range than the interfaces present on the level 1 routers. Some of the routers in topology level 2 possess up to two wireless interfaces, one of which is a wireless interface capable of communicating with level 1 routers. These routers can relay messages at topology level 2 using a frequency band or a medium access control (MAC) protocol that can differ from the one used for communication at topology level 1.

Topology level 3 routers represent routers with the highest capacity. These routers can be equipped with up to three wireless interfaces capable of communicating in turn with level 1 and 2 routers and with other level 3 routers via high-speed point-to-point direct wireless links. It should be noted that routers in level 3 or 2 can have fewer than three or two wireless interfaces, respectively, and the

Internet-Draft

HOLSR

July 2009

only requirement is that level 3 or 2 routers should be equipped with at least one wireless interface for communications at level 3 or 2, respectively. For example, a level 3 router could possess only one high-speed point-to-point interface for ensuring communication with level 3 routers.

The rules governing these first three levels can be scaled to additional topology levels, as the HOLSR network increases in size.

6.2. Cluster Formation

Under the HOLSR mechanism, in each logical topology level, (the last one possibly excepted), the mobile routers form multiple clusters upon cluster heads' invitation, which in turn can be organized into a hierarchical architecture. Unlike (flat) OLSRv2, HOLSR routers can restrict exchange of Topology Control (TC) messages within clusters, thus reducing the amount of control information that is broadcast throughout the entire network. The topology control information between clusters is exchanged via specialized HOLSR routers designated as cluster heads.

The HOLSR cluster heads are selected among those higher-capacity routers. However, the choice of the cluster head is not ruled by HOLSR, it is left to the administrator of the network. Similarly, the level at which a router exists or can exist is determined and assigned administratively, and according to specific network deployment requirements.

Mobile routers can be organized into clusters at different topology levels. A cluster comprises a group of mobile routers (at the same topology level) that are associated with a common cluster head.

At each level, the multiple mobile routers that have been selected as cluster heads declare their status and thereby invite other routers to join their cluster by periodically sending out Cluster ID Announcement (CIA) messages. CIA messages contain two fields: cluster head, which identifies the interface address of the cluster head generating the CIA message, and distance (in hops) to that cluster head. CIA messages are sent with Time to live (TTL) = '1', i.e. they only travel one hop. When a cluster head generates a CIA message, it identifies itself within the cluster head field, with distance being 0.

After reception of a CIA message, the routers in proximity to the cluster head will join that cluster, given that these routers are not part of any cluster, or that the distance to this advertised cluster head is less than the distance to another cluster head they might already be associated with.

Internet-Draft

HOLSR

July 2009

Once a router joins a cluster, it begins to broadcast CIA messages, where the cluster head field contains the ID of the cluster head it is associated with, and the distance field containing the distance, in hops, from that router to the cluster head.

In a nutshell, every router within a cluster, once associated with a cluster head, generates CIA messages with ID field set to cluster head ID, the distance field set to the distance from the router and its cluster head (in hops), and a TTL value of '1'.

6.2.1. Particular Case: Cluster Intersections

A given router may receive CIA messages from different cluster heads, indicating that it is located in the overlapping regions of multiple clusters. In such cases, the router joins whichever cluster is closer in terms of hop count and generates CIA messages only with the ID of the cluster head chosen.

There may also be scenarios where multiple routers with the same capabilities are located in the vicinity of each other. In this situation, a lower-level router may receive CIA messages with the same hop count from various cluster heads. The router will on this case attach itself to the cluster head from which it receives the first CIA and remain with that cluster head for as long as the topology does not change. It should be noted that given the random movement of mobile routers, a router might find a cluster head closer than the one to which it is currently attached. In this case, the mobile router will proceed to change its cluster and attach itself to the closest cluster head. Following this process, each router in the lowest level joins a selected cluster, and the mechanism is in turn applied at the different topology levels.

6.2.2. Cluster ID Announcement Validity Timer

A built-in diagnostic feature helps ensure the robustness of the HOLSR clustering mechanism: each router registers a timeout value for the CIA messages received. Should a cluster head become inactive or move away, routers in its cluster will stop receiving the CIA messages. Eventually the CIA message timeout value will expire and cached CIA information will become invalid. At this point the mobile routers will start to process CIA messages from another cluster and join that cluster, should an opportunity present itself. If no more CIA messages are received, then the routers will behave following the OLSRv2 specifications.

Internet-Draft

HOLSR

July 2009

6.3. Cluster Interaction

One of the main challenges in the HOLSR hierarchical architecture is support for exchange of topology information between clusters at the same or different topology levels without introducing too much overhead. In HOLSR, a cluster head acts as a gateway through which messages from cluster members are relayed to other parts of the network. Under the HOLSR network architecture each cluster head at level i needs to advertise the membership information of routers under its hierarchy to level $i+1$. This is done with the use of the Hierarchical TC message.

6.3.1. Types of Hierarchical Topology Control Message

A Hierarchical Topology Control (HTC) message is generated by a cluster head and used to transmit the membership information of a cluster (under this cluster head) to the higher hierarchical level routers. Three basic types of HTC messages are used:

Full membership

Full membership messages are periodically transmitted by a cluster head to provide information about its cluster members, including members of any lower-level clusters beneath it.

Update

Update HTC messages provide information with respect to cluster membership changes (i.e., the update HTC messages are used when mobile routers join or leave a cluster).

Request

As HTC messages carry a sequence number field, it is possible to determine whether any HTC packet loss has occurred, in which case a request for retransmission of a full membership HTC message is sent by the receiving router.

In topological terms, the higher a given router is located, the more information it obtains about the network. Routers at the highest topology level possess full knowledge of all the routers in the network; consequently, the sizes of their routing tables are as large as they would be under OLSRv2. Routers at lowest topology level are limited in scope; and consequently the size of their routing tables are reduced. I.e. They set their cluster head as default gateway in their routing tables, and all traffic route requests for routers outside their cluster is routed through the cluster head; which is exactly the way it is intended to be.

Internet-Draft

HOLSR

July 2009

6.3.2. HTC and TC Message Propagation

Routers at each hierarchical level independently select MPRs in their respective cluster level. HTC is generated by a cluster head at level i for the level $i+1$, if this level exists. It is forwarded by MPRs at level $i+1$ within this single cluster. At each hierarchical level, TC messages are generated independently. The propagation of the TC is restricted within a cluster, unless a router is located in the overlapping regions of several clusters. Therefore, an HOLSR router's location directly determines the required scope of its knowledge of network topology: for routers located toward the center of the cluster, TC propagation is limited to the local cluster; for routers located in overlapping regions of multiple clusters, the TC message is propagated not only within the local cluster but one hop further to the neighbor routers in other clusters as well. This approach offers two main advantages:

- o The control message reflecting local movement is restricted within the local area, which largely reduces protocol overhead as well as routing table computation overhead
- o Nearby routers in different clusters at the same topology level can communicate directly without having to follow the strict clustering hierarchy, which decreases delay and reduces the load on the cluster head.

6.4. Clusters Decommission

To increase robustness of the HOLSR architecture and to avoid possible network partitions, cluster heads are permitted, in very specific cases, to abandon their "head" status.

If a router that was initialized as a cluster head satisfies the requirements for decommissioning, it starts to emit CIA message announcements with the distance parameter (CLUSTER_HEAD_DIST) set to '255'. Cluster members, with matching cluster head ID, receiving such a message will de-associate from that cluster head. A router that is de-associated will then be able to process CIA messages from neighbor clusters.

The requirements to decommission are (i) to have the parameter C_CRITICAL_PATH set to "true" on one level/interface, and (ii) to lose all connections/links on that interface. It is up to the network administrator to determine what routers on what interfaces could cause network partition if failure occurs.

A decommissioned cluster head will revive its "head" status when the interface with the C_CRITICAL_PATH set to "true" has at least one

Internet-Draft

HOLSR

July 2009

link successfully reinstated. The cluster head will then publish CIA announcements with the distance parameter set to '0'. As a result, routers closer to this cluster head will re-associate with it.

6.5. HOLSR Repositories

In addition to the repositories maintained by OLSRv2, HOLSR maintains the following repositories:

- o A Cluster Information Base.
It contains one entry describing the head of the cluster to which that routers belongs, its relative distance (in hop counts) to the cluster head, and an expiry time. This repository is populated and updated upon receipt of a Cluster Identifier Announcement (CIA) message. The repository entry is timestamped with the time of the last update +<expiry_time>. When the repository entry expires (as consequence of not receiving any CIA message for an <expiry_time> period), the entry is deleted, and the router loses its association to its cluster head. Particular cases:
 - * At the cluster head: Cluster ID is one-self, distance is '0', and expiry_time is 0 or ignored.
 - * At the cluster members: Cluster Head ID is cluster head, distance is]0, 255[, and expiry_time is > 0.
- o Hierarchical Topology Information Base (cluster membership).
This repository is populated and updated upon receipt of Hierarchical Topology Control messages. It contains one entry for each HTC message originator (cluster head) received. Each entry is tupled with the list of member routers, the sequence number it was assigned by the cluster head, and the interface it was retrieved from. Each repository entry is timestamped with the time of the last update +<expiry_time>. Particular cases:
 - * At the cluster head: one of the HTC originator entry contains information about this cluster head router membership. This entry is generated by the router without reception of an HTC message.
 - * At the cluster member routers that are not cluster heads themselves: none of the originator entries matches the router's own address.

Internet-Draft

HOLSR

July 2009

6.6. HOLSR Messages

HOLSR uses the Generalized MANET Packet/Message Format as defined in [RFC5444], and adds 2 new message types and 4 new TLVs, as described below.

6.6.1. CIA Messages

At each level, the multiple mobile routers that have been defined to become cluster heads declare their status and invite other routers to join their cluster by periodically sending out Cluster ID Announcement (CIA) messages. CIA messages MUST NOT be forwarded, as their expected recipients are the immediate neighbors of the emitting router(s).

CIA messages MUST contain the following 2 TLV messages:

- o CLUSTER_HEAD_ID message TLV identifying the interface address of the cluster head generating the CIA message
- o CLUSTER_HEAD_DIST message TLV specifying the distance in hops to that cluster head. There are 3 particular cases for CLUSTER_HEAD_DIST values;
 - * CLUSTER_HEAD_DIST = '0': Identifies and is generated by the cluster head.
 - * CLUSTER_HEAD_DIST is '255': Cluster decommission message generated by the cluster head first, and consequently by its member routers.
 - * CLUSTER_HEAD_DIST is > '0' and < '255': Generated by routers associating with the cluster head, and which are not themselves cluster heads.

CIA messages MAY contain the following 2 TLV messages:

- o VALIDITY_TIME message TLV as specified in [RFC5497], representing CIA_HOLD_TIME for the transmitting MANET interface. The options included in time in [RFC5497], for representing zero and infinite time MUST NOT be used.
- o INTERVAL_TIME message TLV as specified in [RFC5497], representing CIA_INTERVAL for the transmitting MANET interface. The options included in time in [RFC5497], for representing zero and infinite time MUST NOT be used.

Internet-Draft

HOLSR

July 2009

6.6.1.1. CIA Message TLVs

In a CIA message, a router must include a single CLUSTER_HEAD_ID message TLV and a single CLUSTER_HEAD_DIST, as specified in Table 1.

Type	Value Length	Value
CLUSTER_HEAD_ID	1 octet	Specifies the cluster head ID selected by this router (generator of the CIA message)
CLUSTER_HEAD_DIST	1 octet	Specifies the distance, in hops, between the cluster head and this router (generator of the CIA message)

Table 1

6.6.2. HTC Messages

A Hierarchical TC (HTC) message is used to transmit membership information of a cluster to higher hierarchical level routers.

There are 3 basic types of HTC messages:

- o FULL_MEMBERSHIP
Full membership HTC messages are periodically transmitted by a cluster head to provide information about its cluster members, including members of any lower-level clusters beneath it.
- o UPDATE
Update HTC messages provide information with respect to cluster membership changes (i.e., the update HTC messages are used when mobile routers join or leave a cluster). As HTC messages carry a sequence number field, it is possible to determine whether any HTC packet loss has occurred, in which case a request for retransmission of a full membership HTC message is sent by the receiving router. HTC forwarding is enabled by MPRs, and is restricted within a cluster.
- o REQUEST
Request HTC messages are used to request retransmission of full membership HTC messages.

An HTC message MUST contain an HTC message type TLV:

Internet-Draft

HOLSR

July 2009

- o HTC_MSG_TYPE message TLV that is one of FULL_MEMBERSHIP, UPDATE or REQUEST.

An HTC message of type FULL_MEMBERSHIP or UPDATE MUST contain a sequence number message TLV. REQUEST type message do not need to include this TLV:

- o HTC_SEQ_NUM message TLV, specified by an incrementing counter on the HTC message generation at the cluster head router.

6.6.2.1. HTC Message TLVs

TLV messages used within the HTC messages:

Type	Value Length	Value
HTC_MSG_TYPE	1 octet	Specifies the type of HTC message
HTC_SEQ_NUM	2 octets	Specifies the sequence number of the HTC generator counter at the cluster head router

Table 2

6.7. HOLSR Algorithm

6.7.1. Processing and Broadcasting CIA Message

A CIA message is periodically generated and sent by a router that has been configured as cluster head. A CIA message generated from a cluster head has its distance value set to 0.

Routers that join the cluster identified by the cluster head also generate CIA messages periodically. CIA message generated from these routers are sent with their distance value equal to the hop count between themselves and the cluster head.

When a router receives a CIA message, three cases are possible:

- o the router does not yet belong to a cluster: it joins the cluster by (i) populating its cluster head repository and (ii) starting to broadcast CIA messages, inviting other routers to join this cluster. The cluster head distance value stored in its repository is the value received from the CIA message processed plus one.

Internet-Draft

HOLSR

July 2009

- o the router already belongs to a cluster: it checks if its distance to the cluster head indicated in the received CIA message is smaller than its distance to its current cluster head. If yes, it joins the new cluster, updates its cluster head repository and starts broadcasting CIA messages with the newly selected cluster ID (and distance set accordingly). Otherwise, it continues to broadcast CIA messages with the distance and cluster ID values from its repository.
- o the router receives a decommission message (CLUSTER_HEAD_DIST = 255). If the router's cluster head matches the cluster's ID from the CIA message, it will de-associate from that cluster head. The router will forward that decommission message until it associates itself with another cluster head.

6.7.2. Processing and Forwarding HTC Message

Cluster heads generate different types of HTC messages. For all message types, only one sequence number is stored. Every time an HTC message is generated, the sequence number is incremented and included in that message (regardless of its type).

Cluster heads generate HTC messages of type FULL_MEMBERSHIP and UPDATE.

Cluster members generate HTC messages of type REQUEST.

Concerning HTC messages, we distinguish:

- o a HTC message with full membership that is periodically generated and sent by a cluster head of level i to the level $i+1$. This message contains a sequence number that is incremented, as well as the list of routers in the cluster membership (hierarchical topology information base) repository.
- o an update HTC message is generated by a cluster head of level i to notify cluster membership changes to level $i+1$. This message contains a sequence number that is incremented, as well as the new routers having joined the cluster and the old routers having left. It is computed from the cluster membership repository, after addition or removal of a member router entry.
- o a request HTC message is generated by a router that has received a HTC message with a sequence number higher than its last one plus one. This router asks for transmission of a full membership message.

Upon receipt of a full membership HTC message, a router checks the

Internet-Draft

HOLSR

July 2009

received sequence number:

- o if it is smaller or equal, the message is silently discarded.
- o otherwise, the received membership information replaces the old one in the cluster membership repository. The message is forwarded according to the OLSRv2 forwarding rule.

Upon receipt of an update HTC message, a router checks the received sequence number:

- o if it is smaller or equal, the message is silently discarded.
- o if it is greater by 1 than the sequence number maintained in the repository, the received membership information is used to update the information in the cluster membership repository. The message is forwarded according to OLSRv2 forwarding rule.
- o otherwise, if it is greater by more than 1, a request HTC message is generated and sent to get the missing HTC messages.

Upon receipt of a request HTC message, a router checks if it is the originator of the HTC message for which a request is made (usually the originator router is a cluster head):

- o if yes, it sends a full membership message back to the sender.
- o otherwise, this message is forwarded according to OLSR forwarding rule.

6.7.3. Cluster Intersection Handling

Routers located within the intersecting domain of 2 or more clusters are allowed to communicate directly (by-passing their respective cluster heads) by setting the C_INTER_DEPTH parameter to a value greater than 0.

In the intersection domain, routers SHOULD process all HELLO and TC messages, even those not pertaining to their own cluster Id. The use of a neighbor cluster's HELLO and TC message information to build up the routing tables is constrained by the C_INTER_DEPTH parameter.

If C_INTER_DEPTH is greater than 1, an MPR in the intersecting domain is allowed to forward a TC message originating from a router not pertaining to its own cluster.

Internet-Draft

HOLSR

July 2009

Example where Clusters A and B have routers in an intersection vicinity.

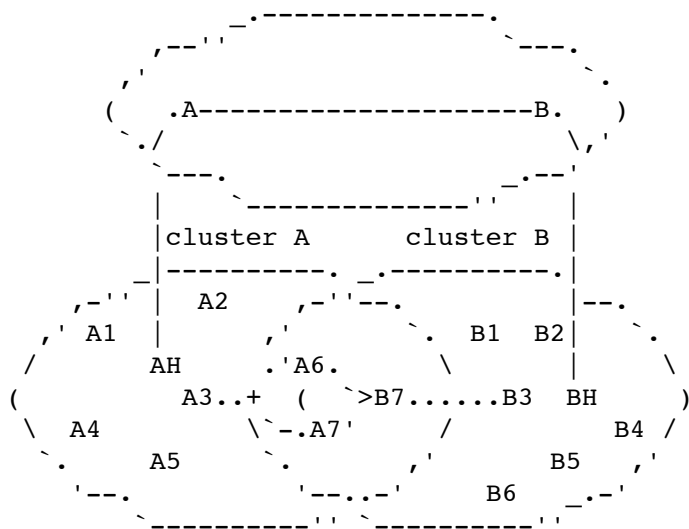


Figure 3 - Intersecting Clusters Example

Clusters A and B intersect because routers in the intersection area (A6, A7, B7) are able to hear CIA announcements from 2 different cluster heads (AH and BH).

If C_INTER_DEPTH is set to 1, then the direct communication between A6 <--> B7, and A7 <--> B7 is enabled. If router A7 wants to communicate with router B3, it will have to do so going through its cluster head AH (A7 -> A3 -> AH -> BH -> B3).

If C_INTER_DEPTH is set to 2, then router A7 is allowed to communicate directly with router B3 (bypassing its cluster head) by using the path provided by router B7; Hence A7 -> B7 -> B3.

7. Proposed Values for Parameters

This section lists the parameters and their default proposed values used in the specification of the protocol.

7.1. Message Intervals Parameter Defaults

7.1.1. CIA Messages

By default these parameters are set to be identical to those used within NHDP [I-D.ietf-manet-nhdp] for HELLO messages and intervals.

Internet-Draft

HOLSR

July 2009

- o CIA_INTERVAL = 2 seconds
- o CIA_MIN_INTERVAL = CIA_INTERVAL/4 = 0.5 seconds

7.1.2. HTC Messages

By default these parameters are set to be identical to those used within OLSRv2 [I-D.ietf-manet-olsrv2] for TC messages and intervals.

- o HTC_INTERVAL = 5 seconds
- o HTC_MIN_INTERVAL = HTC_INTERVAL/4 = 1.25 seconds

7.2. Information Validity Times Parameters

7.2.1. CIA Messages

- o CIA_HOLD_TIME = 3 x CIA_INTERVAL

7.2.2. HTC Messages

Since HTC validity times parameter refer to TC parameters, their default parameters SHOULD use the same values as the ones used in OLSRv2 [I-D.ietf-manet-olsrv2].

- o T_HOLD_TIME = 3 x TC_INTERVAL
- o A_HOLD_TIME = T_HOLD_TIME
- o RX_HOLD_TIME = 30 seconds
- o P_HOLD_TIME = 30 seconds
- o F_HOLD_TIME = 30 seconds

7.3. Cluster Intersection Depth

- o C_INTER_DEPTH = 1

7.4. Cluster Critical Path

- o C_CRITICAL_PATH = 0 (false)

8. Contributors

This specification is the result of the joint efforts of the following contributors, listed alphabetically.

Internet-Draft

HOLSR

July 2009

- o Thomas Heide Clausen, LIX, France, <T.Clausen@computer.org>
- o Ying Ge, CRC, Canada, <ying.ge@crc.ca>
- o Yannick Lacharite, CRC, Canada <yannick.lacharite@crc.ca>
- o Pascale Minet, INRIA, France, <pascale.minet@inria.fr>
- o Maoyu Wang, CRC, Canada, <maoyu.wang@crc.ca>

9. Acknowledgements

The authors would like to acknowledge the authors of OLSRv2 for their guidance in the production of this specification.

We want to acknowledge the work done by Ying Ge (CRC), who implemented a version HOLSR for OLSRv1 for her numerous inputs, reviews and comments. Also the authors would like to gratefully acknowledge the following people for intense technical discussions, early reviews and comments on the specification and its components (listed alphabetically): Ulrich Herberg (LIX), Louise Lamont (CRC), Henning Rogge (FGAN), and the entire IETF MANET working group.

This document was written with the xml2rfc tool described in [RFC2629].

10. IANA Considerations

10.1. Message Types

This specification defines two message types, which must be allocated from the "Assigned Message Types" repository of [RFC5444] with assignment as specified in Table 3.

Name	Type	Description
CIA	TBD	Cluster ID Announcement message
HTC	TBD	Hierarchical Topology Control message

Table 3

Internet-Draft

HOLSR

July 2009

10.2. TLV Types

This specification defines four message TLV types, which must be allocated from the "Message TLV Types" namespace defined in [RFC5444]. IANA are requested to make allocations in the 8-127 range for these types. This will create four new type extension registries with assignments as specified in Table 4. Specifications of these TLVs are in Section 6.6.1 and Section 6.6.2.

Name	Type	Type Extension	Description
CLUSTER_HEAD_ID	TBD	0	Specifies the cluster head ID selected by this router (generator of the CIA message)
		1-255	RESERVED
CLUSTER_HEAD_DIST	TBD	0	Specifies the distance, in hops, between the cluster head and this router (generator of the CIA message)
		1-255	RESERVED
HTC_MSG_TYPE	TBD	0	Specifies the type of HTC message
		1-255	RESERVED
HTC_MSG_NUM	TBD	0	Specifies the sequence number of the HTC generator counter at the cluster head router
		1-255	RESERVED

Table 4

10.3. HTC_MSG_TYPE Values

The values which HTC_MSG_TYPE can use are the following:

- o FULL_MEMBERSHIP = 0
- o UPDATE = 1
- o REQUEST = 2

Internet-Draft

HOLSR

July 2009

11. Security Considerations

All drafts are required to have a security considerations section.

12. References

12.1. Normative References

- [I-D.ietf-manet-olsrv2]
Clausen, T., Dearlove, C., and P. Jacquet, "The Optimized Link State Routing Protocol version 2", draft-ietf-manet-olsrv2-09 (work in progress), July 2009.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5148] Clausen, T., Dearlove, C., and B. Adamson, "Jitter Considerations in Mobile Ad Hoc Networks (MANETs)", RFC 5148, February 2008.
- [RFC5444] Clausen, T., Dearlove, C., Dean, J., and C. Adjih, "Generalized Mobile Ad Hoc Network (MANET) Packet/Message Format", RFC 5444, February 2009.
- [RFC5497] Clausen, T. and C. Dearlove, "Representing Multi-Value Time in Mobile Ad Hoc Networks (MANETs)", RFC 5497, March 2009.

12.2. Informative References

- [I-D.ietf-manet-nhdp]
Clausen, T., Dearlove, C., and J. Dean, "MANET Neighborhood Discovery Protocol (NHDP)", draft-ietf-manet-nhdp-10 (work in progress), July 2009.
- [RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, June 1999.
- [RFC3626] Clausen, T. and P. Jacquet, "Optimized Link State Routing Protocol (OLSR)", RFC 3626, October 2003.

Appendix A. Example of Data Transfer using Clusters

This appendix gives an example of data transfer in a hierarchical network with three levels. It reuses the figure presented in the Protocol Overview (Section 4).

Internet-Draft

HOLSR

July 2009

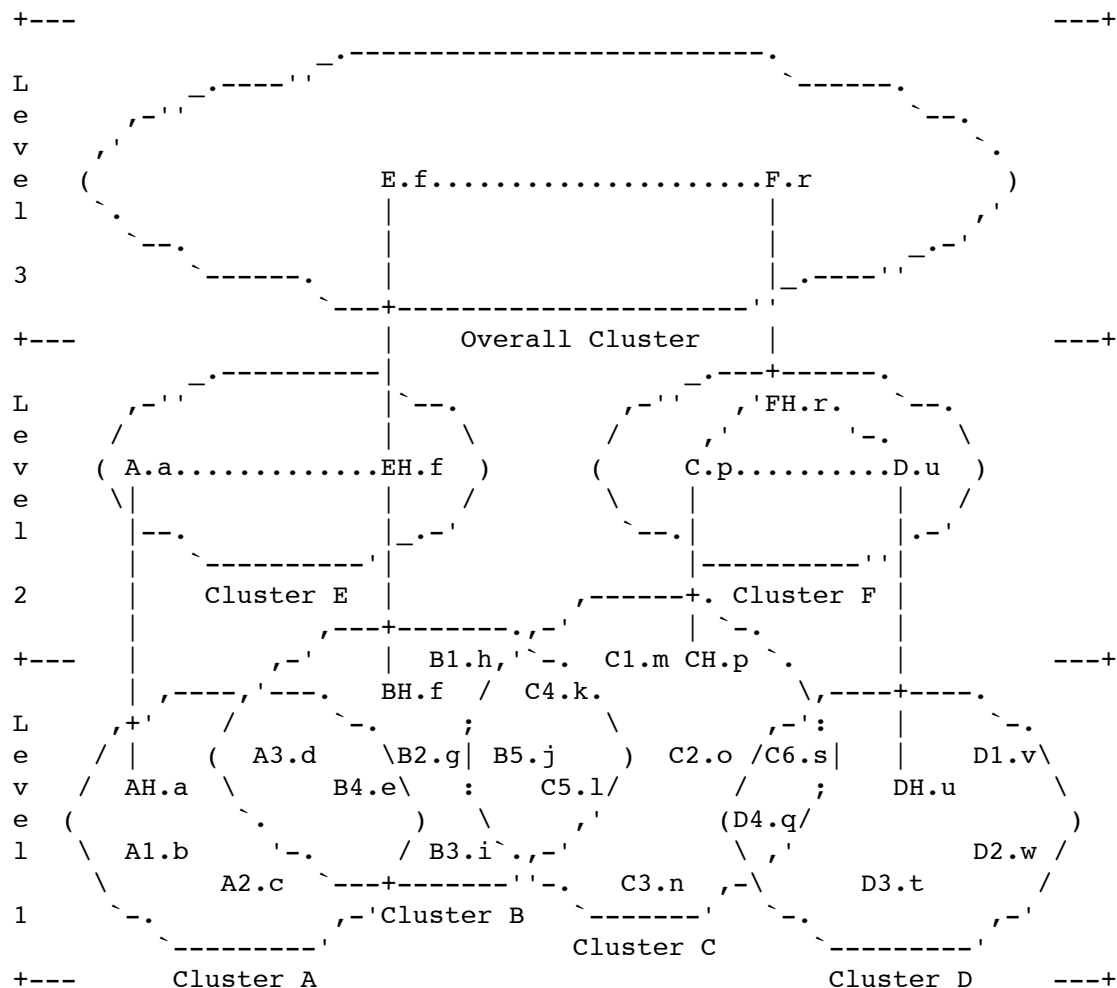


Figure 4 - Data Transfer using Hierarchical Structure

Routing information is updated by all the routers in a given cluster. Thus, communication between routers in a cluster is achieved via the routing information in the routing tables. For data transmissions outside the cluster, the cluster heads are used as gateways. To describe the communication between clusters, let's look at an example network with 3 levels.

First, let's specify the terminology of the example. Routers are represented by a "." followed by a lowercase letter. Figure 4 contains 23 routers labeled from ".a" to ".w". Prefixed to the router id is a capital character followed by a number, or an H. It defines the cluster membership of a router and its cluster head status. In the case where a router is involved in many clusters, on many levels (because it contains multiple interfaces), the capital

Internet-Draft

HOLSR

July 2009

letter references the cluster membership from the lowest level. For example "AH.a" denotes router "a" which is the cluster head "H" of cluster "A". Similarly, "D3.t" denotes router "t", the 3rd router member of cluster D, and not a cluster head. On level 2, "C.p" is member of cluster F, but it is also the same router as level 1 "CH.p" which is the cluster head of cluster C. A bit more complex, router "BH.f" is the cluster head of cluster B at level 1, it is the same router as "EH.f" which is a cluster head of cluster E at level 2, and it is also the same router as "E.f" which is a member of the overall cluster at level 3. Hence for the purpose of our example, router .f has 3 interfaces each one participating at a different level of the hierarchy. Additionally, in the text description we denote levels with the prefix <l#->, e.g. <l1-> applies to level 1.

Now, the example: Starting from the base topology level 1, when a router .b (l1-A1.b) of cluster A wants to transmit data to router .t (l1-D3.t) from another cluster D; it will route through the hierarchy levels instead of communicating directly at level 1. Here is the actual path taken by the data. Router l1-A1.b first sends its traffic to its cluster head router AH.a (l1-AH.a) -- On the next topology level 2, cluster head l1-AH.a corresponds to l2-A.a and is a member of Cluster E -- Router l2-A.a transmits in turn to its cluster head router l2-EH.f -- which corresponds to l3-E.f on level 3 -- Router l3-E.f transmits to its neighbor l3-F.r -- equivalent to l2-FH.r on level 2 -- l2-FH.r forwards to l2-D.u -- cluster head of level one l1-DH.u -- Finally l1-DH.u transmits data to its destination l1-D3.t. In summary, the data path can be represented by: l1-D3.t --> l1-AH.a == l2-A.a --> l2-EH.f == l3-E.f --> l3-F.r == l2-FH.r --> l2-D.u == l1-DH.u --> l1-D3.t. Or, abstracting the level/cluster information, traffic traversed 5 hops: .t --> .a --> .f --> .r --> .u --> .t.

As we trace the path traveled by the data packets, we see that the cluster head is always used as the gateway by its member routers at lower hierarchical levels, for destinations outside the cluster. The rules governing the three levels network can be scaled to a network with additional topology levels.

Normally, data transmission between clusters follows clustering hierarchy described above. However, a different transmission path is used when transmitting data between neighboring routers that belong to different clusters at the same topology level; in this situation, the cluster head is not used as a gateway to relay the information to the destination router in a different cluster. Instead, data is sent to the neighbor router directly since TC messages are propagated one hop further on cluster edges to neighbor routers in other clusters. The depth of the routing infiltration at the edge of clusters is configurable via the C_INTER_DEPTH parameter. With this strategy the

Internet-Draft

HOLSR

July 2009

HOLSR makes efficient use of the high-capacity routers without overloading them.

Appendix B. Example of Cluster Decommission Handling

This appendix gives an example of cluster decommission handling in a hierarchical network with three levels. It reuses the figure presented in the Protocol Overview (Section 4). In this example, node .f has an interface at level 3 that is set to C_CRITICAL_PATH = '1' (true).

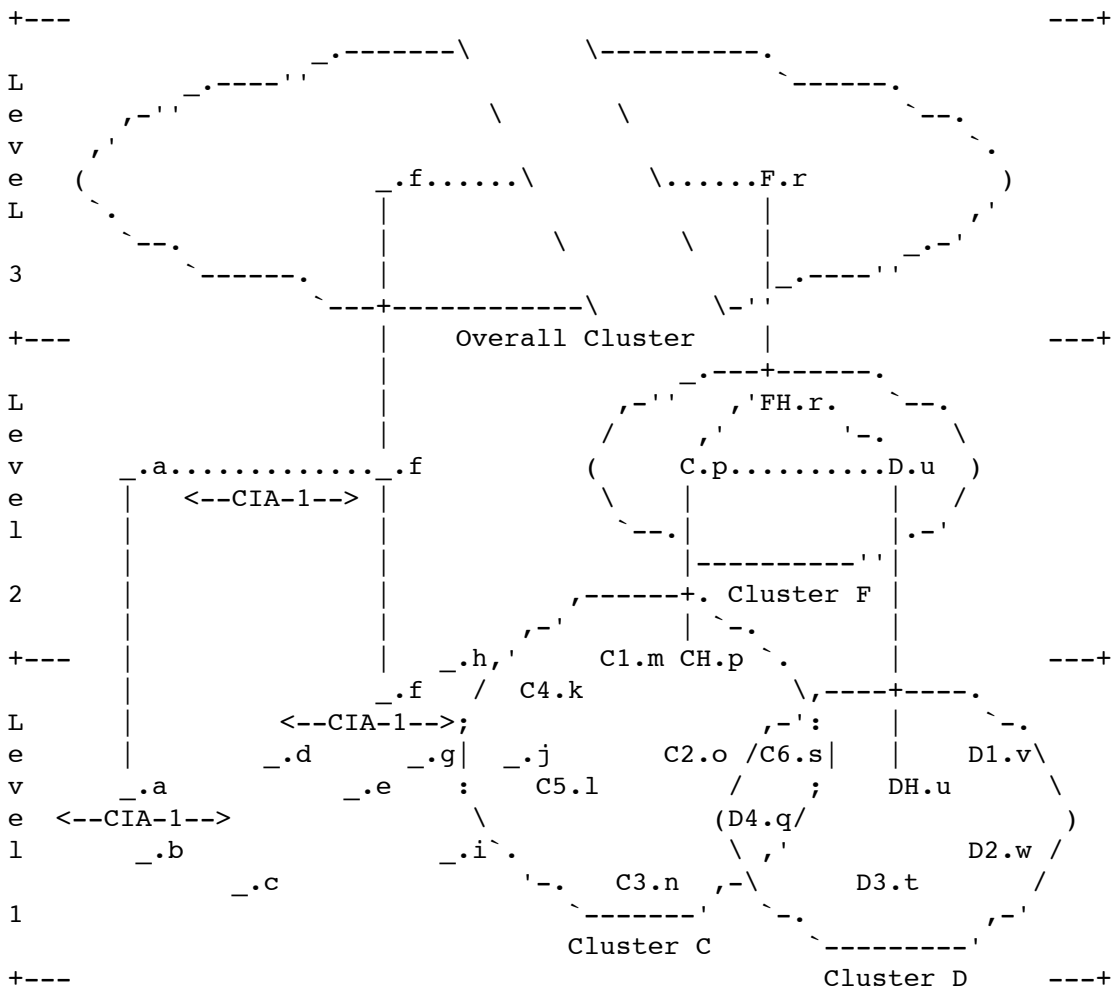


Figure 5 - Cluster Head Decommissioning

When the link on layer 3 in-between router .f and .r is broken, the decommission requirement on router .f is met. Hence, router .f stops

Internet-Draft

HOLSR

July 2009

being a cluster head on layer 2 and layer 1, and starts to generate CIA decommission message (CIA-1 in the example, as shown in Figure 5). The immediate impact, is that we see clusters A, B and E de-associate completely. De-associated routers will now process and generate CIA messages from cluster C.

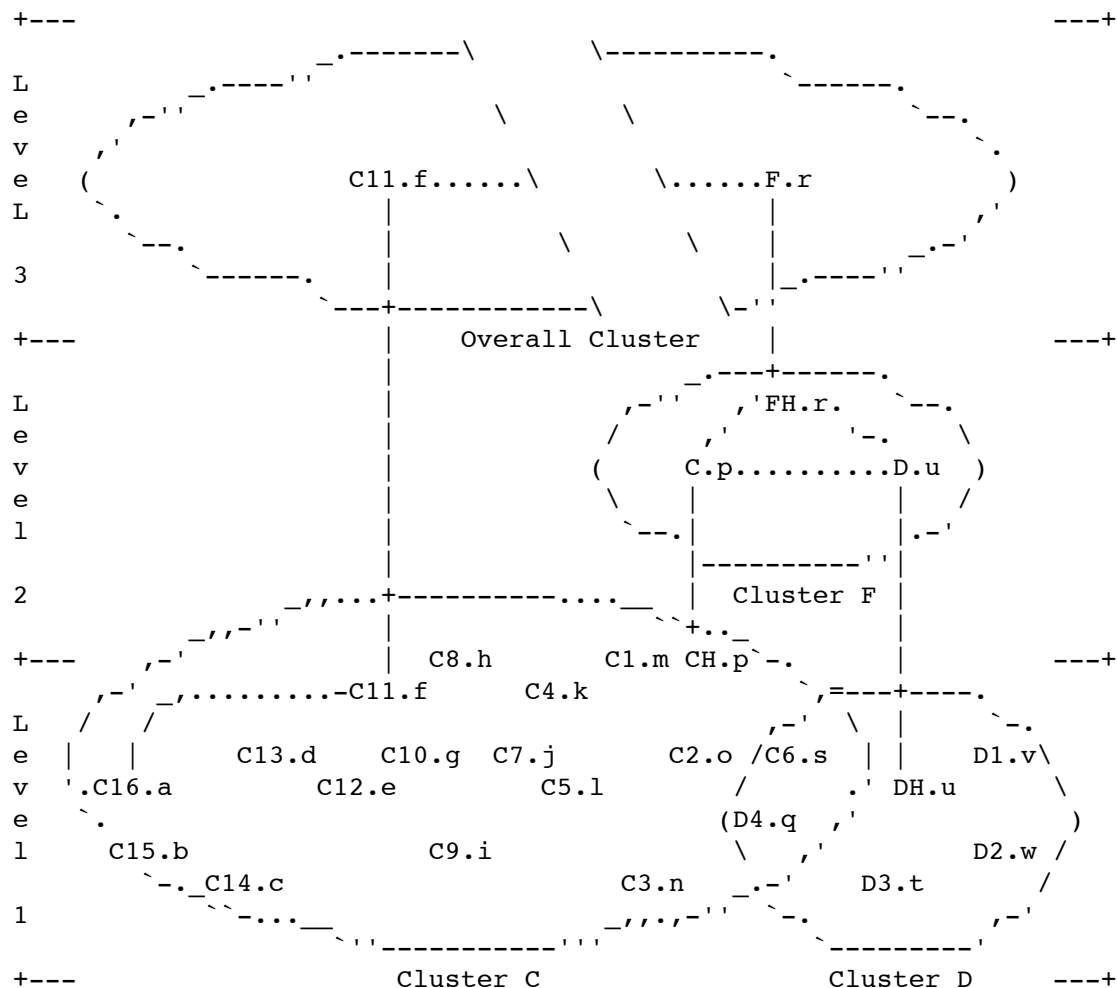


Figure 6 - Updated Architecture

As shown Figure 6, de-associated routers end up joining cluster C, which now contains all routers from previous clusters A, B, and E. Routers .a and .f still are still linked via their second interface (higher capability), but their multiple interfaces are now handled by OLSRv2 (no more HTC messaging required from routers .a and .f since they are no longer cluster heads).

Note. From this example, it can also be shown that routers .f and .r

Internet-Draft

HOLSR

July 2009

at level 1 shall not be both set with parameter C_CRITICAL_PATH = 1. Doing so, if the link in-between these 2 routers fails, all the network cluster heads would be decommissioned. It is the responsibility of the network administrator to ensure that the C_CRITICAL_PATH parameter gets set properly on strategically selected routers.

Appendix C. Packet and Message Layout

This appendix illustrates the translation from the abstract descriptions of packets employed in the protocol specification, and the bit-layout packets actually exchanged between MANET routers.

C.1. Packet and Message Options

The basic layout of an HOLSR packet, message body, address block, TLV block, and TLV is as described in [RFC5444], allowing all options.

C.2. Example CIA Message

An example CIA message, using IPv4 (four octet) addresses is as follows. The overall message length is 37 octets.

The message has the mhasorig, mhashoplimit, mhashopcount, and mhasseqnum flags of its four-bit flags field set (value 15), and hence includes an Originator Address, a Hop Limit, a Hop Count, and a Message Sequence Number (which is type independent). Its four-bit message Address Length field has value 3, indicating that the message uses IPv4 address. Thus, the overall message flags octet is of value 243. The message has a Message TLV block with content length 16 octets, containing four message TLVs. These TLVs represent message validity time, message interval time, CIA Cluster Head Id, and CIA Cluster Head Distance. Each message TLV has the thasvalue flag of its flag octet set (value 16), and hence contains a Value field, with preceding Value Length 1 octet.

The message has one Address Block. It contains the interface address of the cluster head. The first octet indicates that it contains one address. The flags octet has the ahashead flag set indicating it has a head. Its head length is 4, this is equal to the address length, it thus has no mid section. This Address Block has no TLVs (TLV block content length is 0 octets).

Internet-Draft

HOLSR

July 2009

```

      0             1             2             3
0 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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|      CIA      |1 1 1 1 0 0 1 1|0 0 0 0 0 0 0 0 0 0 1 0 0 1 0 1|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|      Originator Address      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  Hop Limit  |  Hop Count  |  Message Sequence Number  |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0 0 0 0 0 0 0 0 0 0 0 1 0 0 0 0| VALIDITY_TIME |0 0 0 1 0 0 0 0|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0 0 0 0 0 0 0 1|      Value      | INTERVAL_TIME |0 0 0 1 0 0 0 0|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0 0 0 0 0 0 0 1|      Value      | CLUSTER_HD_ID |0 0 0 1 0 0 0 0|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0 0 0 0 0 0 0 1|      Value      | CLUSTER_HD_DST|0 0 0 1 0 0 0 0|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0 0 0 0 0 0 0 1|      Value      |0 0 0 0 0 0 0 1|1 0 0 0 0 0 0 0|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0 0 0 0 0 1 0 0|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  Head (cont)  |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

C.3. Example HTC Message

An example HTC message, using IPv4 (four octet) addresses, is as follows. The overall message length is 48 octets.

The message has the mhasorig, mhashoplimit, mhashopcount, and mhasseqnum flags of its four-bit flags field set (value 15), and hence includes an Originator Address, a Hop Limit, a Hop Count, and a Message Sequence Number (which is type independent). Its four-bit message Address Length field has value 3, indicating that the message uses IPv4 address. Thus, the overall message flags octet is of value 243. The message has a Message TLV block with content length 17 octets containing four TLVs. The first two TLVs are validity and interval times as for the CIA message above. The third TLV is the HTC message type. The fourth TLV is the HTC sequence number TLV used to carry the 2 octet HTC_SEQ_NUM. Each message TLV has the thasvalue flag of its flag octet set (value 16), and hence contains a Value field. First three TLVs have a preceding Value Length 1 octet, and the last TLV has a preceding Value Length 2 octets.

The message has one Address Block. It contains 6 cluster members' addresses. The first octet indicates that it contains 6 addresses. The flag octet has the ahashead flag set indicating it has a head. Its head length is 2 octets, hence mid sections length is two octets.

Internet-Draft

HOLSR

July 2009

This Address Block has no TLVs (TLV block content length 0 octets).

```

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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          HTC          |1 1 1 1 0 0 1 1|0 0 0 0 0 0 0 0 0 0 1 1 0 0 0 0|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Originator Address          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  Hop Limit  |  Hop Count  |  Message Sequence Number  |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0 0 0 0 0 0 0 0 0 0 0 0 1 0 0 0 1| VALIDITY_TIME |0 0 0 1 0 0 0 0|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0 0 0 0 0 0 0 0 1|  Value  | INTERVAL_TIME |0 0 0 1 0 0 0 0|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0 0 0 0 0 0 0 0 1|  Value  |  HTC_MSG_TYPE |0 0 0 1 0 0 0 0|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0 0 0 0 0 0 0 0 1|  Value  |  HTC_SEQ_NUM  |0 0 0 1 0 0 0 0|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0 0 0 0 0 0 1 0|          Value          |0 0 0 0 0 1 1 0|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|1 0 0 0 0 0 0 0 0|0 0 0 0 0 0 1 0|          Head          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Mid          |          Mid          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Mid          |          Mid          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Mid          |          Mid          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

Authors' Addresses

Yannick Lacharite
Communications Research Centre Canada
3701 Carling Avenue
Ottawa, Ontario K2H 8S2
Canada

Phone: +1 613 998 1207
Fax:
Email: yannick.lacharite@crc.gc.ca
URI:

Internet-Draft

HOLSR

July 2009

Maoyu Wang
Communications Research Centre Canada
3701 Carling Avenue
Ottawa, Ontario K2H 8S2
Canada

Phone: +1 613 991 1671
Fax:
Email: maoyu.wang@crc.gc.ca
URI:

Pascale Minet
INRIA
Rocquencourt, 78153 Le Chesnay Cedex
France

Phone: +33 1 3963 5233
Fax:
Email: pascale.minet@inria.fr
URI: <http://hipercom.inria.fr/~minet>

Thomas Clausen
LIX, Ecole polytechnique
92128 Palaiseau Cedex
France

Phone: +33 6 6058 9349
Fax:
Email: t.clausen@computer.org
URI: <http://www.thomasclausen.org>

