

Mobile Ad hoc Networking (MANET)
Internet-Draft
Intended status: Standards Track
Expires: September 11, 2009

T. Clausen
LIX, Ecole Polytechnique
C. Dearlove
BAE Systems ATC
P. Jacquet
Project Hipercom, INRIA
The OLSRV2 Design Team
MANET Working Group
March 10, 2009

The Optimized Link State Routing Protocol version 2
draft-ietf-manet-olsrv2-08

Status of This Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79. This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on September 11, 2009.

Copyright Notice

Internet-Draft

OLSRv2

March 2009

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

This document describes version 2 of the Optimized Link State Routing (OLSRv2) protocol. The protocol embodies an optimization of the classical link state algorithm tailored to the requirements of a Mobile Ad hoc NETWORK (MANET).

Table of Contents

1. Introduction	5
2. Terminology	6
3. Applicability Statement	7
4. Protocol Overview and Functioning	8
5. Protocol Parameters and Constants	10
5.1. Local History Times	11
5.2. Message Intervals	11
5.3. Advertised Information Validity Times	12
5.4. Received Message Validity Times	13
5.5. Jitter	13
5.6. Hop Limit Parameter	14
5.7. Willingness	14
5.8. Parameter Change Constraints	15
6. Information Bases	16
6.1. Local Information Base	17
6.1.1. Originator Set	17
6.1.2. Local Attached Network Set	17
6.2. Neighbor Information Base	18
6.3. Topology Information Base	18
6.3.1. Advertised Neighbor Set	19
6.3.2. Advertising Remote Router Set	19
6.3.3. Topology Set	20
6.3.4. Attached Network Set	20
6.3.5. Routing Set	21
6.4. Processing and Forwarding Information Base	21
6.4.1. Received Set	22
6.4.2. Processed Set	22
6.4.3. Forwarded Set	22
6.4.4. Relay Set	23
7. Message Processing and Forwarding	23

Internet-Draft

OLSRv2

March 2009

7.1. Actions when Receiving a Message	24
7.2. Message Considered for Processing	25
7.3. Message Considered for Forwarding	25
8. Packets and Messages	27
8.1. HELLO Messages	28
8.1.1. HELLO Message TLVs	28
8.1.2. HELLO Message Address Block TLVs	29
8.2. TC Messages	29
8.2.1. TC Message TLVs	30
8.2.2. TC Message Address Block TLVs	31
9. HELLO Message Generation	31
9.1. HELLO Message: Transmission	32
10. HELLO Message Processing	32
10.1. Updating Willingness	32
10.2. Updating MPR Selectors	33
10.3. Symmetric 1-Hop and 2-Hop Neighborhood Changes	33
11. TC Message Generation	34
11.1. TC Message: Transmission	35
12. TC Message Processing	36
12.1. Invalid Message	36
12.2. Initial TC Message Processing	37
12.3. Initial TC Message Processing	38
12.3.1. Populating the Advertising Remote Router Set	38
12.3.2. Populating the Topology Set	39
12.3.3. Populating the Attached Network Set	40
12.4. Completing TC Message Processing	40
12.4.1. Purging the Topology Set	40
12.4.2. Purging the Attached Network Set	41
13. Information Base Changes	41
14. Selecting MPRs	42
15. Populating Derived Sets	43
15.1. Populating the Relay Set	43
15.2. Populating the Advertised Neighbor Set	44
16. Routing Set Calculation	44
16.1. Network Topology Graph	44
16.2. Populating the Routing Set	46
16.3. Routing Set Updates	46
17. Proposed Values for Parameters and Constants	47
17.1. Local History Time Parameters	47
17.2. Message Interval Parameters	47
17.3. Advertised Information Validity Time Parameters	47
17.4. Received Message Validity Time Parameters	47
17.5. Jitter Time Parameters	48
17.6. Hop Limit Parameter	48
17.7. Willingness Parameter and Constants	48
18. Sequence Numbers	48
19. IANA Considerations	49
19.1. Message Types	49

Internet-Draft

OLSRv2

March 2009

19.2. Message TLV Types	49
19.3. Address Block TLV Types	50
20. Security Considerations	51
20.1. Confidentiality	51
20.2. Integrity	51
20.3. Interaction with External Routing Domains	53
21. Contributors	53
22. Acknowledgments	54
23. References	54
23.1. Normative References	54
23.2. Informative References	55
Appendix A. Router Configuration	55
Appendix B. Example Algorithm for Calculating MPRs	56
B.1. Terminology	56
B.2. MPR Selection Algorithm for each OLSRv2 Interface	57
Appendix C. Example Algorithm for Calculating the Routing Set	58
C.1. Add Local Symmetric Links	58
C.2. Add Remote Symmetric Links	59
C.3. Add Attached Networks	60
Appendix D. Example Message Layout	61
Appendix E. Constraints	62
Appendix F. Flow and Congestion Control	66

Internet-Draft

OLSRv2

March 2009

1. Introduction

The Optimized Link State Routing protocol version 2 (OLSRv2) is an update to OLSRv1 as published in [RFC3626]. Compared to [RFC3626], OLSRv2 retains the same basic mechanisms and algorithms, while providing a more flexible signaling framework and some simplification of the messages being exchanged. Also, OLSRv2 accommodates either IPv4 and IPv6 addresses in a compact manner.

OLSRv2 is developed for mobile ad hoc networks. It operates as a table driven, proactive protocol, i.e. it exchanges topology information with other routers in the network regularly. Each router selects a set of its neighbor routers as "MultiPoint Relays" (MPRs). Control traffic may be flooded through the network using hop by hop forwarding, but where a router only needs to forward control traffic directly received from its MPR selectors (routers which have selected it as an MPR). This mechanism, denoted "MPR flooding", provides an efficient mechanism for information distribution within the MANET by reducing the number of transmissions required.

Routers selected as MPRs also have a special responsibility when declaring link state information in the network. A sufficient requirement for OLSRv2 to provide shortest (lowest hop count) path routes to all destinations is that routers declare link state information for their MPR selectors, if any. Additional available link state information may be transmitted, e.g. for redundancy. Thus, as well as being used to facilitate MPR flooding, use of MPRs allows the reduction of the number and size of link state messages, and MPRs are used as intermediate routers in multi-hop routes.

A router selects MPRs from among its one hop neighbors connected by "symmetric", i.e. bi-directional, links. Therefore, selecting routes through MPRs automatically avoids the problems associated with data packet transfer over uni-directional links (such as the problem of not getting link layer acknowledgments at each hop, for link layers employing this technique).

OLSRv2 is developed to work independently from other protocols. (Parts of OLSRv2 have been published separately as [RFC5444], [timetlv], [RFC5148] and [NHDP] for wider use.) Likewise, OLSRv2 makes no assumptions about the underlying link layer. However, OLSRv2 may use link layer information and notifications when available and applicable, as described in [NHDP].

OLSRv2, as OLSRv1, inherits its concept of forwarding and relaying from HIPERLAN (a MAC layer protocol) which is standardized by ETSI [HIPERLAN], [HIPERLAN2].

Internet-Draft

OLSRv2

March 2009

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

MANET specific terminology is to be interpreted as described in [RFC5444] and [NHDP].

Additionally, this document uses the following terminology:

Router - A MANET router which implements the Optimized Link State Routing protocol version 2 as specified in this document.

OLSRv2 interface - A MANET interface, running OLSRv2. Note that all references to MANET interfaces in [NHDP] refer to OLSRv2 interfaces when using [NHDP] as part of OLSRv2.

Address - An address, as recorded in the Information Bases specified by this protocol, and included in HELLO and TC messages generated by this protocol, may be either an address or an address prefix. These can be represented as a single address object in a HELLO or TC message, as defined by [RFC5444]. An address so represented is considered to have a prefix length equal to its length (in bits) when considered as an address object, and a similar convention is used in the Information Bases specified by this protocol. Two addresses (address objects) are considered equal only if their prefix lengths are also equal.

Willingness - The willingness of a router is a numerical value between WILL_NEVER and WILL_ALWAYS (both inclusive), which represents the router's willingness to be selected as an MPR.

Willing symmetric 1-hop neighbor - A symmetric 1-hop neighbor of this router which has willingness not equal to WILL_NEVER.

Symmetric strict 2-hop neighbor - A symmetric 2-hop neighbor of this router which is not a symmetric 1-hop neighbor of this router, and is a symmetric 1-hop neighbor of a willing symmetric 1-hop neighbor of this router.

Symmetric strict 2-hop neighbor through OLSRv2 interface I - A symmetric strict 2-hop neighbor of this router which is a symmetric 1-hop neighbor of a willing symmetric 1-hop neighbor of this router via a symmetric link including OLSRv2 interface I. This router MAY elect to consider only information received over OLSRv2 interface I in making this determination.

Internet-Draft

OLSRv2

March 2009

Symmetric strict 2-hop neighborhood - The set of the symmetric strict 2-hop neighbors of a router.

Multipoint relay (MPR) - A router which is selected by its symmetric 1-hop neighbor, router X, to "re-transmit" all the broadcast messages that it receives from router X, provided that the message is not a duplicate, and that the hop limit field of the message is greater than one.

MPR selector - A router which has selected its symmetric 1-hop neighbor, router X, as one of its MPRs is an MPR selector of router X.

MPR flooding - The optimized MANET-wide information distribution mechanism, employed by this protocol, in which a message is relayed by only a reduced subset of the routers in the network.

This document employs the same notational conventions as in [RFC5444] and [NHDP].

3. Applicability Statement

OLSRv2 is a proactive routing protocol for mobile ad hoc networks (MANETs) [RFC2501]. The larger and more dense a network, the more optimization can be achieved by using MPRs compared to the classic link state algorithm. OLSRV2 enables hop-by-hop routing, i.e. each router using its local information provided by OLSRV2 to route packets.

As OLSRV2 continuously maintains routes to all destinations in the network, the protocol is beneficial for traffic patterns where the traffic is random and sporadic between a large subset of routers, and where the (source, destination) pairs are changing over time. No additional control traffic need be generated in this case since routes are maintained for all known destinations at all times. Also, since routes are maintained continuously, traffic is subject to no delays due to buffering or to route discovery, and consequently no data traffic buffering is imposed.

OLSRv2 supports routers which have multiple interfaces which participate in the MANET using OLSRV2. As described in [NHDP], each OLSRV2 interface may have one or more network addresses (which may have prefix lengths). OLSRV2, additionally, supports routers which have non-OLSRv2 interfaces which may be local to a router or which can serve as gateways towards other networks.

OLSRv2 uses the format specified in [RFC5444] for all messages and packets. OLSRV2 is thereby able to allow for extensions via

Internet-Draft

OLSRv2

March 2009

"external" and "internal" extensibility. External extensibility allows a protocol extension to specify and exchange new Message Types, which can be forwarded and delivered correctly even by routers which do not support that extension. Internal extensibility allows a protocol extension to define additional attributes to be carried embedded in the standard OLSRv2 control messages detailed in this specification (or any new Message Types defined by other protocol extensions) using the TLV mechanism specified in [RFC5444], while still allowing routers not supporting that extension to forward messages including the extension, and to process messages ignoring the extension.

The OLSRv2 neighborhood discovery protocol using HELLO messages is specified in [NHDP]. This neighborhood discovery protocol serves to ensure that each OLSRv2 router has available continuously updated Information Bases describing the router's 1-hop and symmetric 2-hop neighbors. This neighborhood discovery protocol, which also uses [RFC5444], is extended in this document by the addition of MPR information.

OLSRv2 does not make any assumption about router addresses, other than that each router is assumed to have at least one unique and routable IP address for each interface that it has which participates in the MANET.

OLSRv2 can, as does [NHDP], use the link local multicast address "LL-MANET-Routers", and either the "manet" UDP port or the "manet" IP protocol number, all as specified in [manet-iana].

4. Protocol Overview and Functioning

OLSRv2 is a proactive routing protocol for mobile ad hoc networks. The protocol inherits the stability of a link state algorithm and has the advantage of having routes immediately available when needed due to its proactive nature. OLSRv2 is an optimization of the classical link state protocol, tailored for mobile ad hoc networks. The main tailoring and optimizations of OLSRv2 are:

- o Unacknowledged transmission of all control messages; control messages are sent periodically, but may also be sent in response to changes in the local neighborhood.
- o MPR flooding for MANET-wide link state information distribution.
- o Partial topology maintenance - each router knows only a subset of the links in the network, sufficient for a minimum hop route to all destinations.

Internet-Draft

OLSRv2

March 2009

The MPR flooding and partial topology maintenance are based on the concept of MultiPoint Relays (MPRs), selected independently by routers based on the symmetric 1-hop and 2-hop neighbor information maintained using [NHDP].

Using the message exchange format [RFC5444] and the neighborhood discovery protocol [NHDP], OLSRv2 also contains the following main components:

- o A TLV, to be included within the HELLO messages of [NHDP], allowing a router to signal MPR selection.
- o The optimized mechanism for MANET-wide information distribution, denoted "MPR flooding".
- o A specification of MANET-wide signaling, denoted TC (Topology Control) messages. TC messages in OLSRv2 serve to:
 - * inject link state information into the entire MANET;
 - * inject addresses of hosts and networks for which they may serve as a gateway into the entire network.

TC messages are emitted periodically, thereby allowing routers to continuously track changes in the network. Incomplete TC messages may be used to report additions to advertised information without repeating unchanged information. Some TC messages may be MPR flooded over only part of the network, allowing a router to ensure that nearer routers are kept more up to date than distant routers, such as is used in Fisheye State Routing [FSR] and Fuzzy Sighted Link State routing [FSLs].

Each router in the network selects a set of MPRs. The MPRs of a router X may be any subset of router X's willing symmetric 1-hop neighbors, such that every router in the symmetric strict 2-hop neighborhood of router X has a symmetric link to at least one of router X's MPRs. The MPRs of a router may thus be said to "cover" the router's symmetric strict 2-hop neighborhood. Each router also maintains information about the set of symmetric 1-hop neighbors that have selected it as an MPR, its MPR selectors.

As long as the condition above is satisfied, any algorithm selecting MPRs is acceptable in terms of implementation interoperability. However if smaller sets of MPRs are selected then the greater the efficiency gains that are possible. An analysis and examples of MPR selection algorithms is given in [MPR].

A router may independently determine and advertise its willingness to

Internet-Draft

OLSRv2

March 2009

be selected as an MPR. A router may advertise that it always should be selected as an MPR or that it should never be selected as an MPR. In the latter case, the router will neither relay control messages, nor will that router be included as an intermediate router in any routing table calculations. Use of variable willingness is most effective in dense networks.

In OLSRv2, actual efficiency gains are based on the sizes of each router's Relay Set, the set of symmetric 1-hop neighbors for which it is to relay broadcast traffic, and its Advertised Neighbor Set, the set of symmetric 1-hop neighbors for which it is to advertise link state information into the network in TC messages. Each of these sets MUST contain all MPR selectors, and MAY contain additional routers. If the Advertised Neighbor Set is empty, TC messages are not generated by that router, unless needed for gateway reporting, or for a short period to accelerate the removal of outdated link state information.

OLSRv2 is designed to work in a completely distributed manner and does not depend on any central entity. The protocol does not require reliable transmission of control messages; each router sends control messages periodically, and can therefore sustain a reasonable loss of some such messages. Such losses may occur frequently in radio networks due to collisions or other transmission problems. OLSRv2 MAY use "jitter", randomized adjustments to message transmission times, to reduce the incidence of collisions [RFC5148].

OLSRv2 does not require sequenced delivery of messages. Each TC message contains a sequence number which is incremented for each message. Thus the recipient of a TC message can, if required, easily identify which information is more recent - even if messages have been re-ordered while in transmission.

OLSRv2 only interacts with IP through routing table management. OLSRv2 sends its control messages as described in [RFC5444] and [NHDP].

5. Protocol Parameters and Constants

The parameters and constants used in this specification are those defined in [NHDP] plus those defined in this section. The separation in [NHDP] into interface parameters, router parameters and constants is also used in OLSRv2, however all but one (RX_HOLD_TIME) of the parameters added by OLSRv2 are router parameters. Parameters may be classified into the following categories:

Internet-Draft

OLSRv2

March 2009

- o Local history times
- o Message intervals
- o Advertised information validity times
- o Received message validity times
- o Jitter times
- o Hop limits
- o Willingness

In addition, constants for particular cases of a router's willingness to be an MPR are defined. These parameters and constants are detailed in the following sections. As for the parameters in [NHDP], parameters defined in this document may be changed dynamically by a router, and need not be the same on different routers, even in the same MANET, or on different interfaces of the same router (for interface parameters).

5.1. Local History Times

The following router parameter manages the time for which local information is retained:

`O_HOLD_TIME` - is used to define the time for which a recently used and replaced originator address is used to recognize the router's own messages.

The following constraint applies to this parameter:

- o `O_HOLD_TIME` ≥ 0

5.2. Message Intervals

The following router parameters regulate TC message transmissions by a router. TC messages are usually sent periodically, but MAY also be sent in response to changes in the router's Advertised Neighbor Set and Local Attached Network Set. With a larger value of the parameter `TC_INTERVAL`, and a smaller value of the parameter `TC_MIN_INTERVAL`, TC messages may more often be transmitted in response to changes in a highly dynamic network. However because a router has no knowledge of, for example, routers remote to it (i.e. beyond 2 hops away) joining the network, TC messages MUST NOT be sent purely responsively.

Internet-Draft

OLSRv2

March 2009

TC_INTERVAL - is the maximum time between the transmission of two successive TC messages by this router. When no TC messages are sent in response to local network changes (by design, or because the local network is not changing) then TC messages SHOULD be sent at a regular interval TC_INTERVAL, possibly modified by jitter as specified in [RFC5148].

TC_MIN_INTERVAL - is the minimum interval between transmission of two successive TC messages by this router. (This minimum interval MAY be modified by jitter, as specified in [RFC5148].)

The following constraints apply to these parameters:

- o TC_INTERVAL > 0
- o TC_MIN_INTERVAL >= 0
- o TC_INTERVAL >= TC_MIN_INTERVAL
- o If INTERVAL_TIME TLVs as defined in [timetlv] are included in TC messages, then TC_INTERVAL MUST be representable as described in [timetlv].

5.3. Advertised Information Validity Times

The following router parameters manage the validity time of information advertised in TC messages:

T_HOLD_TIME - is used to define the minimum value in the VALIDITY_TIME TLV included in all TC messages sent by this router. If a single value of parameter TC_HOP_LIMIT (see Section 5.6) is used then this will be the only value in that TLV.

A_HOLD_TIME - is the period during which TC messages are sent after they no longer have any advertised information to report, but are sent in order to accelerate outdated information removal by other routers.

The following constraints apply to these parameters:

- o T_HOLD_TIME > 0
- o A_HOLD_TIME >= 0
- o T_HOLD_TIME >= TC_INTERVAL
- o If TC messages can be lost, then both T_HOLD_TIME and A_HOLD_TIME SHOULD be significantly greater than TC_INTERVAL; a value >= 3 x

Internet-Draft

OLSRv2

March 2009

TC_INTERVAL is RECOMMENDED.

- o T_HOLD_TIME MUST be representable as described in [timetlv].

5.4. Received Message Validity Times

The following parameters manage the validity time of recorded received message information:

RX_HOLD_TIME - is an interface parameter, and is the period after receipt of a message by the appropriate OLSRv2 interface of this router for which that information is recorded, in order that the message is recognized as having been previously received on this OLSRv2 interface.

P_HOLD_TIME - is a router parameter, and is the period after receipt of a message which is processed by this router for which that information is recorded, in order that the message is not processed again if received again.

F_HOLD_TIME - is a router parameter, and is the period after receipt of a message which is forwarded by this router for which that information is recorded, in order that the message is not forwarded again if received again.

The following constraints apply to these parameters:

- o RX_HOLD_TIME > 0
- o P_HOLD_TIME > 0
- o F_HOLD_TIME > 0
- o All of these parameters SHOULD be greater than the maximum difference in time that a message may take to traverse the MANET, taking into account any message forwarding jitter as well as propagation, queuing, and processing delays.

5.5. Jitter

If jitter, as defined in [RFC5148], is used then these parameters are as follows:

TP_MAXJITTER - represents the value of MAXJITTER used in [RFC5148] for periodically generated TC messages sent by this router.

Internet-Draft

OLSRv2

March 2009

TT_MAXJITTER - represents the value of MAXJITTER used in [RFC5148] for externally triggered TC messages sent by this router.

F_MAXJITTER - represents the default value of MAXJITTER used in [RFC5148] for messages forwarded by this router. However before using F_MAXJITTER a router MAY attempt to deduce a more appropriate value of MAXJITTER, for example based on any INTERVAL_TIME or VALIDITY_TIME TLVs contained in the message to be forwarded.

For constraints on these parameters see [RFC5148].

5.6. Hop Limit Parameter

The parameter TC_HOP_LIMIT is the hop limit set in each TC message. TC_HOP_LIMIT MAY be a single fixed value, or MAY be different in TC messages sent by the same router. However each other router, at any hop count distance, SHOULD see a regular pattern of TC messages, in order that meaningful values of INTERVAL_TIME and VALIDITY_TIME TLVs at each hop count distance can be included as defined in [timetlv]. Thus the pattern of TC_HOP_LIMIT SHOULD be defined to have this property. For example the repeating pattern (255 4 4) satisfies this property (having period TC_INTERVAL at hop counts up to 4, inclusive, and 3 x TC_INTERVAL at hop counts greater than 4), but the repeating pattern (255 255 4 4) does not satisfy this property because at hop counts greater than 4, message intervals are alternately TC_INTERVAL and 3 x TC_INTERVAL.

The following constraints apply to this parameter:

- o The maximum value of TC_HOP_LIMIT \geq the network diameter in hops, a value of 255 is RECOMMENDED.
- o All values of TC_HOP_LIMIT \geq 2.

5.7. Willingness

Each router has a WILLINGNESS parameter, which MUST be in the range WILL_NEVER to WILL_ALWAYS, inclusive, and represents its willingness to be an MPR, and hence its willingness to forward messages and be an intermediate router on routes. If a router has WILLINGNESS = WILL_NEVER it does not perform these tasks. A MANET using OLSRv2 with too many routers with WILLINGNESS = WILL_NEVER will not function; it MUST be ensured, by administrative or other means, that this does not happen.

Routers MAY have different WILLINGNESS values; however the three constants WILL_NEVER, WILL_DEFAULT and WILL_ALWAYS MUST have the

Internet-Draft

OLSRv2

March 2009

values defined in Section 5.7. (Use of WILLINGNESS = WILL_DEFAULT allows a router to avoid including an MPR_WILLING TLV in its TC messages, use of WILLINGNESS = WILL_ALWAYS means that a router will always be selected as an MPR by all symmetric 1-hop neighbors.)

The following constraints apply to this parameter:

- o WILLINGNESS >= WILL_NEVER
- o WILLINGNESS <= WILL_ALWAYS

5.8. Parameter Change Constraints

This section presents guidelines, applicable if protocol parameters are changed dynamically.

O_HOLD_TIME

- * If O_HOLD_TIME for a router changes, then O_time for all Originator Tuples MAY be changed.

TC_INTERVAL

- * If the TC_INTERVAL for a router increases, then the next TC message generated by this router MUST be generated according to the previous, shorter, TC_INTERVAL. Additional subsequent TC messages MAY be generated according to the previous, shorter, TC_INTERVAL.
- * If the TC_INTERVAL for a router decreases, then the following TC messages from this router MUST be generated according to the current, shorter, TC_INTERVAL.

RX_HOLD_TIME

- * If RX_HOLD_TIME for an OLSRv2 interface changes, then RX_time for all Received Tuples for that OLSRv2 interface MAY be changed.

P_HOLD_TIME

- * If P_HOLD_TIME changes, then P_time for all Processed Tuples MAY be changed.

Internet-Draft

OLSRv2

March 2009

F_HOLD_TIME

- * If F_HOLD_TIME changes, then F_time for all Forwarded Tuples MAY be changed.

TP_MAXJITTER

- * If TP_MAXJITTER changes, then the periodic TC message schedule on this router MAY be changed immediately.

TT_MAXJITTER

- * If TT_MAXJITTER changes, then externally triggered TC messages on this router MAY be rescheduled.

F_MAXJITTER

- * If F_MAXJITTER changes, then TC messages waiting to be forwarded with a delay based on this parameter MAY be rescheduled.

TC_HOP_LIMIT

- * If TC_HOP_LIMIT changes, and the router uses multiple values after the change, then message intervals and validity times included in TC messages MUST be respected. The simplest way to do this is to start any new repeating pattern of TC_HOP_LIMIT values with its largest value.

6. Information Bases

Each router maintains the Information Bases described in the following sections. These are used for describing the protocol in this document. An implementation of this protocol MAY maintain this information in the indicated form, or in any other organization which offers access to this information. In particular note that it is not necessary to remove Tuples from Sets at the exact time indicated, only to behave as if the Tuples were removed at that time.

The purpose of OLSRV2 is to determine the Routing Set, which may be used to update IP's Routing Table, providing "next hop" routing information for IP datagrams. OLSRV2 maintains the following Information Bases:

Internet-Draft

OLSRv2

March 2009

Local Information Base - as defined in [NHDP], extended by the addition of an Originator Set, defined in Section 6.1.1 and a Local Attached Network Set, defined in Section 6.1.2.

Interface Information Bases - as defined in [NHDP], one Interface Information Base for each OLSRv2 interface.

Neighbor Information Base - as defined in [NHDP], extended by the addition of three elements to each Neighbor Tuple, as defined in Section 6.2.

Topology Information Base - this Information Base is specific to OLSRv2, and is defined in Section 6.3.

Processing and Forwarding Information Base - this Information Base is specific to OLSRv2, and is defined in Section 6.4.

The ordering of sequence numbers, when considering which is the greater, is as defined in Section 18.

6.1. Local Information Base

The Local Information Base as defined in [NHDP] is extended by the addition of an Originator Set, defined in Section 6.1.1, and a Local Attached Network Set, defined in Section 6.1.2.

6.1.1. Originator Set

A router's Originator Set records addresses that were recently used as originator addresses by this router. If a router's originator address is immutable then this set is always empty and MAY be omitted. It consists of Originator Tuples:

(O_orig_addr, O_time)

where:

O_orig_addr is a recently used originator address;

O_time specifies the time at which this Tuple expires and MUST be removed.

6.1.2. Local Attached Network Set

A router's Local Attached Network Set records its local non-OLSRv2 interfaces via which it can act as gateways to other networks. The Local Attached Network Set is not modified by this protocol. This protocol MAY respond to changes to the Local Attached Network Set,

Internet-Draft

OLSRv2

March 2009

which MUST reflect corresponding changes in the router's status. It consists of Local Attached Network Tuples:

(AL_net_addr, AL_dist)

where:

AL_net_addr is the network address of an attached network which can be reached via this router.

AL_dist is the number of hops to the network with address AL_net_addr from this router.

Attached networks local to this router SHOULD be treated as local non-MANET interfaces, and added to the Local Interface Set, as specified in [NHDP], rather than being added to the Local Attached Network Set.

An attached network MAY also be attached to other routers.

It is not the responsibility of OLSRv2 to maintain routes from this router to networks recorded in the Local Attached Network Set.

Local Attached Neighbor Tuples are removed from the Local Attached Network Set only when the routers' local attached network configuration changes, i.e. they are not subject to timer-based expiration or changes due to received messages.

6.2. Neighbor Information Base

Each Neighbor Tuple in the Neighbor Set, defined in [NHDP], has these additional elements:

N_willingness is the router's willingness to be selected as an MPR, in the range from WILL_NEVER to WILL_ALWAYS, both inclusive;

N_mpr is a boolean flag, describing if this neighbor is selected as an MPR by this router;

N_mpr_selector is a boolean flag, describing if this neighbor has selected this router as an MPR, i.e. is an MPR selector of this router.

6.3. Topology Information Base

The Topology Information Base stores information required for the generation and processing of TC messages, and information received in TC messages. The Advertised Neighbor Set contains interface

Internet-Draft

OLSRv2

March 2009

addresses of symmetric 1-hop neighbors which are to be reported in TC messages. The Advertising Remote Router Set, the Topology Set and the Attached Network Set record information received in TC messages.

Additionally, a Routing Set is maintained, derived from the information recorded in the Neighborhood Information Base, Topology Set, Attached Network Set and Advertising Remote Router Set.

6.3.1. Advertised Neighbor Set

A router's Advertised Neighbor Set contains interface addresses of symmetric 1-hop neighbors which are to be advertised through TC messages. It consists of Advertised Neighbor Tuples:

(A_neighbor_iface_addr)

In addition, an Advertised Neighbor Set Sequence Number (ANSN) is maintained. Each time the Advertised Neighbor Set is updated, the ANSN MUST be incremented. The ANSN MUST also be incremented if there is a change to the set of Local Attached Network Tuples that are to be advertised in the router's TC messages.

The Advertised Neighbor Set for a router is derived from the Neighbor Set of that same router, and so Advertised Neighbor Tuples are removed when, for example, the corresponding Neighbor Tuples in the Neighbor Set are removed. Advertised Neighbor Tuples are not subject to timer-based expiration.

6.3.2. Advertising Remote Router Set

A router's Advertising Remote Router Set records information describing each remote router in the network that transmits TC messages. It consists of Advertising Remote Router Tuples:

(AR_orig_addr, AR_seq_number, AR_iface_addr_list, AR_time)

where:

AR_orig_addr is the originator address of a received TC message, note that this does not include a prefix length;

AR_seq_number is the greatest ANSN in any TC message received which originated from the router with originator address AR_orig_addr (i.e. which contributed to the information contained in this Tuple);

Internet-Draft

OLSRv2

March 2009

AR_iface_addr_list is an unordered list of the interface addresses of the router with originator address AR_orig_addr;

AR_time is the time at which this Tuple expires and MUST be removed.

6.3.3. Topology Set

A router's Topology Set records topology information about the network. It consists of Topology Tuples:

(T_dest_iface_addr, T_orig_addr, T_seq_number, T_time)

where:

T_dest_iface_addr is an interface address of a destination router, which may be reached in one hop from the router with originator address T_orig_addr;

T_orig_addr is the originator address of a router which is the last hop on a path towards the router with interface address T_dest_iface_addr, note that this does not include a prefix length;

T_seq_number is the greatest ANSN in any TC message received which originated from the router with originator address T_orig_addr (i.e. which contributed to the information contained in this Tuple);

T_time specifies the time at which this Tuple expires and MUST be removed.

6.3.4. Attached Network Set

A router's Attached Network Set records information about networks attached to other routers. It consists of Attached Network Tuples:

(AN_net_addr, AN_orig_addr, AN_dist, AN_seq_number, AN_time)

where:

AN_net_addr is the network address of an attached network, which may be reached via the router with originator address AN_orig_addr;

AN_orig_addr is the originator address of a router which can act as gateway to the network with address AN_net_addr, note that this does not include a prefix length;

Internet-Draft

OLSRv2

March 2009

AN_dist is the number of hops to the network with address AN_net_addr from the router with originator address AN_orig_addr;

AN_seq_number is the greatest ANSN in any TC message received which originated from the router with originator address AN_orig_addr (i.e. which contributed to the information contained in this Tuple);

AN_time specifies the time at which this Tuple expires and MUST be removed.

6.3.5. Routing Set

A router's Routing Set records the selected path to each destination for which a route is known. It consists of Routing Tuples:

(R_dest_addr, R_next_iface_addr, R_dist, R_local_iface_addr)

where:

R_dest_addr is the address of the destination, either the address of an interface of a destination router, or the network address of an attached network;

R_next_iface_addr is the OLSRv2 interface address of the "next hop" on the selected path to the destination;

R_dist is the number of hops on the selected path to the destination;

R_local_iface_addr is the address of the local OLSRv2 interface over which a packet MUST be sent to reach the destination by the selected path.

The Routing Set for a router is derived from the contents of the other sets of the router, and is updated (Routing Tuples added or removed) when routing paths are calculated. Routing Tuples are not subject to timer-based expiration.

6.4. Processing and Forwarding Information Base

The Processing and Forwarding Information Base records information required to ensure that a message is processed at most once and is forwarded at most once per OLSRv2 interface of a router, using MPR flooding.

Internet-Draft

OLSRv2

March 2009

6.4.1. Received Set

A router has a Received Set per local OLSRV2 interface. Each Received Set records the signatures of messages which have been received over that OLSRV2 interface. Each consists of Received Tuples:

(RX_type, RX_orig_addr, RX_seq_number, RX_time)

where:

RX_type is the received Message Type;

RX_orig_addr is the originator address of the received message;

RX_seq_number is the message sequence number of the received message;

RX_time specifies the time at which this Tuple expires and MUST be removed.

6.4.2. Processed Set

A router's Processed Set records signatures of messages which have been processed by the router. It consists of Processed Tuples:

(P_type, P_orig_addr, P_seq_number, P_time)

where:

P_type is the processed Message Type;

P_orig_addr is the originator address of the processed message;

P_seq_number is the message sequence number of the processed message;

P_time specifies the time at which this Tuple expires and MUST be removed.

6.4.3. Forwarded Set

A router's Forwarded Set records signatures of messages which have been processed by the router. It consists of Forwarded Tuples:

(F_type, F_orig_addr, F_seq_number, F_time)

where:

Internet-Draft

OLSRv2

March 2009

F_type is the forwarded Message Type;

F_orig_addr is the originator address of the forwarded message;

F_seq_number is the message sequence number of the forwarded message;

F_time specifies the time at which this Tuple expires and MUST be removed.

6.4.4. Relay Set

A router has a Relay Set per local OLSRv2 interface. Each Relay Set records the OLSRv2 interface addresses of symmetric 1-hop neighbors, such that the router is to forward messages received from those neighbors' OLSRv2 interfaces, on that local OLSRv2 interface, if not otherwise excluded from forwarding that message (e.g. by it having been previously forwarded). It consists of Relay Tuples:

(RY_neighbor_iface_addr)

The Relay Set for an interface is derived from the Link Set for the same interface, and so Relay Tuples are removed when the corresponding Link Tuples in the Link Set of this interface are removed, or when processing otherwise suggests their removal. Relay Tuples are not subject to timer-based expiration.

7. Message Processing and Forwarding

On receiving a packet, as defined in [RFC5444], a router divides the packet into the Packet Header and messages. OLSRv2 defines, and hence owns, the TC Message Type, and hence receives all TC messages. OLSRv2 is responsible for determining whether a TC message is to be processed (updating Information Bases) and/or forwarded.

OLSRv2 also receives HELLO messages, which are defined, and hence owned, by [NHDP]. Received HELLO messages MUST be made available to OLSRv2 when received on an OLSRv2 interface and after NHDP has completed its processing thereof. OLSRv2 also processes HELLO messages, OLSRv2 does not forward HELLO messages.

Extensions to OLSRv2 which define, and hence own, other Messages Types, MAY manage the processing and/or forwarding of these messages using the same mechanism as for TC messages. These mechanisms contain elements (P_type, RX_type, F_type) required only for such usage.

The processing selection and forwarding mechanisms are designed to

Internet-Draft

OLSRv2

March 2009

only need to parse the Message Header in order to determine whether a message is to be processed and/or forwarded, and not to have to parse the Message Body even if the message is forwarded (but not processed). An implementation MAY either only parse the Message Body if necessary, or MAY always parse the Message Body. An implementation MUST discard the message silently if it is unable to parse the Message Header or (if attempted) the Message Body.

OLSRv2 does not require any part of the Packet Header.

7.1. Actions when Receiving a Message

If the router receives a HELLO message from NHDP, then the message is processed according to Section 10.

A router MUST perform the following tasks for each received TC message or other Message Type defined by an extension to OLSRV2 and specified to use this process:

1. If the router recognizes from the originator address of the message that the message is one which the receiving router itself originated (i.e. is the current originator address of the router, or is an O_orig_addr in an Originator Tuple) then the message MUST be silently discarded.
2. Otherwise:
 1. Otherwise:
 1. If the message is of a type which may be processed, including being a TC message, then the message is considered for processing according to Section 7.2, AND;
 2. If for the message is of a type which may be forwarded, including being a TC message, AND:
 - <msg-hop-limit> is present and <msg-hop-limit> > 1, AND;
 - <msg-hop-count> is not present or <msg-hop-count> < 255

then the message is considered for forwarding according to Section 7.3.

Internet-Draft

OLSRv2

March 2009

7.2. Message Considered for Processing

If a message (the "current message") is considered for processing, then the following tasks MUST be performed:

1. If a Processed Tuple exists with:

- * P_type = the Message Type of the current message, AND;
- * P_orig_addr = the originator address of the current message, AND;
- * P_seq_number = the message sequence number of the current message;

then the current message MUST NOT be processed.

2. Otherwise:

1. Create a Processed Tuple with:

- + P_type := the Message Type of the current message;
- + P_orig_addr := the originator address of the current message;
- + P_seq_number := the sequence number of the current message;
- + P_time := current time + P_HOLD_TIME.

2. Process the current message according to its type. For a TC message this is as defined in Section 12.

7.3. Message Considered for Forwarding

If a message (the "current message") is considered for forwarding, then the following tasks MUST be performed:

1. If the sending interface address (the source address of the IP datagram containing the current message) does not match (taking into account any address prefix) an OLSRv2 interface address in an L_neighbor_iface_addr_list of a Link Tuple, with L_status = SYMMETRIC, in the Link Set for the OLSRv2 interface on which the current message was received (the "receiving interface") then the current message MUST be silently discarded.

Internet-Draft

OLSRv2

March 2009

2. Otherwise:

1. If a Received Tuple exists in the Received Set for the receiving interface, with:

- + RX_type = the Message Type of the current message, AND;
- + RX_orig_addr = the originator address of the current message, AND;
- + RX_seq_number = the sequence number of the current message;

then the current message MUST be silently discarded.

2. Otherwise:

1. Create a Received Tuple in the Received Set for the receiving interface with:

- RX_type := the Message Type of the current message;
- RX_orig_addr := originator address of the current message;
- RX_seq_number := sequence number of the current message;
- RX_time := current time + RX_HOLD_TIME.

2. If a Forwarded Tuple exists with:

- F_type = the Message Type of the current message, AND;
- F_orig_addr = the originator address of the current message, AND;
- F_seq_number = the sequence number of the current message.

then the current message MUST be silently discarded.

3. Otherwise if the sending interface address matches (taking account of any address prefix) an RY_neighbor_iface_addr in the Relay Set for the receiving interface, then:

Internet-Draft

OLSRv2

March 2009

1. Create a Forwarded Tuple with:
 - o F_type := the Message Type of the current message;
 - o F_orig_addr := originator address of the current message;
 - o F_seq_number := sequence number of the current message;
 - o F_time := current time + F_HOLD_TIME.
2. The Message Header of the current message is modified by:
 - o if present, decrement <msg-hop-limit> in the Message Header by 1, AND;
 - o if present, increment <msg-hop-count> in the Message Header by 1.
3. For each OLSRv2 interface of the router, include the message in a packet to be transmitted on that OLSRv2 interface, as described in Section 8. This packet MAY contain other forwarded messages and/or messages generated by this router, including by other protocols using [RFC5444]. Forwarded messages MAY be jittered as described in [RFC5148]. The value of MAXJITTER used in jittering a forwarded message MAY be based on information in that message (in particular any INTERVAL_TIME or VALIDITY_TIME TLVs in that message) or otherwise SHOULD be with a maximum delay of F_MAXJITTER. A router MAY modify the jitter applied to a message in order to more efficiently combine messages in packets, as long as the maximum jitter is not exceeded.

8. Packets and Messages

The packet and message format used by OLSRv2 is defined in [RFC5444]. Except as otherwise noted, options defined in [RFC5444] may be freely used, in particular alternative formats defined by packet, message, Address Block and TLV flags.

OLSRv2 defines and owns the TC Message Type. OLSRv2 also modifies HELLO messages (owned by [NHDP]) by adding TLVs to these messages when sent over OLSRv2 interfaces, and processes these HELLO messages, subsequent to their processing by NHDP. Extensions to OLSRv2 MAY

Internet-Draft

OLSRv2

March 2009

define additional Message Types to be handled similarly to TC messages.

Routers using OLSRV2 exchange information through messages. One or more messages sent by a router at the same time SHOULD be combined into a single packet. These messages may have originated at the sending router, or have originated at another router and are forwarded by the sending router. Messages with different originating routers MAY be combined for transmission within the same packet. Messages from other protocols defined using [RFC5444] MAY be combined for transmission within the same packet.

The remainder of this section defines, within the framework of [RFC5444], Message Types and TLVs specific to OLSRV2. All references in this specification to TLVs that do not indicate a type extension, assume Type Extension = 0. TLVs in processed messages with a type extension which is neither zero as so assumed, nor a specifically indicated non-zero type extension, are ignored.

8.1. HELLO Messages

A HELLO message in OLSRV2 is generated as specified in [NHDP]. In addition, an OLSRV2 router MUST be able to modify such messages, prior to these being sent on an OLSRV2 interface, so that such HELLO messages:

- o MUST include TLV(s) with Type := MPR associated with all OLSRV2 interface addresses that:
 - * are included in the HELLO message associated with a TLV with Type = LINK_STATUS and Value = SYMMETRIC; AND
 - * are included in a Neighbor Tuple with N_mpr = true.
- o MUST NOT include any TLVs with Type = MPR associated with any other addresses.
- o MAY include a message TLV with Type := MPR_WILLING, indicating the router's willingness to be selected as an MPR.

An OLSRV2 router MUST also be able to process any HELLO message received on an OLSRV2 interface, subsequent to the processing specified in [NHDP].

8.1.1. HELLO Message TLVs

In a HELLO message, a router MUST include an MPR_WILLING Message TLV as specified in Table 1, unless WILLINGNESS = WILL_DEFAULT (in which

Internet-Draft

OLSRv2

March 2009

case it MAY be included). A router MUST NOT include more than one MPR_WILLING Message TLV.

Type	Value Length	Value
MPR_WILLING	1 octet	Router parameter WILLINGNESS; unused bits (based on the maximum willingness value WILL_ALWAYS) are RESERVED and SHOULD be set to zero.

Table 1

If a router does not advertise an MPR_WILLING TLV in a HELLO message, then the router MUST be assumed to have WILLINGNESS equal to WILL_DEFAULT.

8.1.2. HELLO Message Address Block TLVs

In a HELLO message, a router MAY include MPR Address Block TLV(s) as specified in Table 2.

Type	Value Length	Value
MPR	0 octets	None.

Table 2

8.2. TC Messages

A TC message MUST contain:

- o <msg-orig-addr>, <msg-seq-num> and <msg-hop-limit> elements in its Message Header, as specified in [RFC5444].
- o A <msg-hop-count> element in its Message Header if the message contains a TLV with either Type = VALIDITY_TIME or Type = INTERVAL_TIME indicating more than one time value according to distance.
- o A single Message TLV with Type := CONT_SEQ_NUM, and Type Extension := COMPLETE or Type Extension := INCOMPLETE, as specified in Section 8.2.1 (for complete and incomplete TC messages, respectively).

Internet-Draft

OLSRv2

March 2009

- o A Message TLV with Type := VALIDITY_TIME, as specified in [timetlv]. The options included in [timetlv] for representing zero and infinite times MUST NOT be used.
- o All of the router's interface addresses. These MUST be included in the message's Address Blocks, unless:
 - * the router has a single interface, with a single interface address with maximum prefix length; AND
 - * that address is the router's originator address.

In this exceptional case, the address will be included as the message's originator address, and MAY be omitted from the message's Address Blocks.

- o TLV(s) with Type := LOCAL_IF and Value := UNSPEC_IF associated with all of the router's interface addresses.
- o If the TC message is complete, all addresses in the Advertised Address Set and all addresses in the Local Attached Network Set, the latter (only) with associated GATEWAY Address Block TLV(s), as specified in Section 8.2.2.

A TC message MAY contain:

- o If the TC message is incomplete, any addresses in the Advertised Address Set and any addresses in the Local Attached Network Set, the latter (only) with associated GATEWAY Address Block TLV(s), as specified in Section 8.2.2.
- o A Message TLV with Type := INTERVAL_TIME, as specified in [timetlv]. The options included in [timetlv] for representing zero and infinite times MUST NOT be used.

8.2.1. TC Message TLVs

In a TC message, a router MUST include a single CONT_SEQ_NUM Message TLV, as specified in Table 3, and with Type Extension = COMPLETE or Type Extension = INCOMPLETE, according to whether the TC message is complete or incomplete.

Type	Value Length	Value
CONT_SEQ_NUM	2 octets	The ANSN contained in the Advertised Neighbor Set.

Internet-Draft

OLSRv2

March 2009

Table 3

8.2.2. TC Message Address Block TLVs

In a TC message, a router MAY include GATEWAY Address Block TLV(s) as specified in Table 4.

Type	Value Length	Value
GATEWAY	1 octet	Number of hops to attached network.

Table 4

GATEWAY Address Block TLV(s) MUST be associated with all attached network addresses, and MUST NOT be associated with any other addresses.

9. HELLO Message Generation

An OLSRv2 HELLO message is composed and generated as defined in [NHDP], with the following additions:

- o A Message TLV with Type := MPR_WILLING and Value := WILLINGNESS MUST be included, unless WILLINGNESS = WILL_DEFAULT (in which case it MAY be included).
- o For each address which is included in the message with an associated TLV with Type = LINK_STATUS and Value = SYMMETRIC, and is of an MPR (i.e. the address is in the N_neighbor_iface_addr_list of a Neighbor Tuple with N_mpr = true), that address (including a different copy of that address, in the same or a different Address Block) MUST be associated with an Address Block TLV with Type := MPR.
- o For each address which is included in the message and is not associated with a TLV with Type = LINK_STATUS and Value = SYMMETRIC, or is not of an MPR (i.e. the address is not in the N_neighbor_iface_addr_list of a Neighbor Tuple with N_mpr = true), that address (including different copies of that address, in the same or different Address Blocks) MUST NOT be associated with an Address Block TLV with Type := MPR.
- o An additional HELLO message MAY be sent when the router's set of MPRs changes, in addition to the cases specified in [NHDP], and subject to the same constraints.

Internet-Draft

OLSRv2

March 2009

9.1. HELLO Message: Transmission

HELLO messages are included in packets as specified in [RFC5444]. These packets may contain other messages, including TC messages.

10. HELLO Message Processing

All HELLO message processing, including determination of whether a message is invalid, considers only TLVs with Type Extension = 0. TLVs with any other type extension are ignored. All references to, for example, a TLV with Type = MPR_WILLING refer to a TLV with Type = MPR_WILLING and Type Extension = 0.

In addition to the reasons specified in [NHDP], a HELLO message MUST NOT:

- o Have more than one TLV with Type = MPR_WILLING in its Message TLV Block, where TLVs have different Values.
- o Contain any address associated with a TLV with Type = MPR, where that address (including a different copy of that address, in the same or a different Address Block) which is not also associated with the single value SYMMETRIC by a TLV with Type = LINK_STATUS or Type = OTHER_NEIGHB.

Such a HELLO message MAY be discarded before processing. If it is not then all TLVs with the type(s) for which an error was indicated MUST be ignored (treated as not present) in the following processing.

HELLO messages are first processed as specified in [NHDP]. The router MUST identify the Neighbor Tuple corresponding to the originator of the HELLO message (the "current Neighbor Tuple") and update its N_willingness as described in Section 10.1 and its N_mpr_selector as described in Section 10.2. Following these, the router MUST also perform the processing defined in Section 10.3.

10.1. Updating Willingness

N_willingness in the current Neighbor Tuple is updated as follows:

1. If the HELLO message contains a Message TLV with Type = MPR_WILLING then N_willingness := the value of that TLV;
2. Otherwise, N_willingness := WILL_DEFAULT.

Internet-Draft

OLSRv2

March 2009

10.2. Updating MPR Selectors

N_mpr_selector is updated as follows:

1. If a router finds any of its local OLSRv2 interface addresses with an associated TLV with Type = MPR in the HELLO message (indicating that the originator router has selected the receiving router as an MPR) then, for the current Neighbor Tuple:

- * N_mpr_selector := true

2. Otherwise, if a router finds any of its own interface addresses with an associated TLV with Type = LINK_STATUS and Value = SYMMETRIC in the HELLO message, then for the current Neighbor Tuple:

- * N_mpr_selector := false

10.3. Symmetric 1-Hop and 2-Hop Neighborhood Changes

A router MUST also perform the following:

1. If N_symmetric of a Neighbor Tuple changes from true to false, for that Neighbor Tuple:

- * N_mpr_selector := false

2. The set of MPRs of a router MUST be recalculated if:

- * a Link Tuple is added with L_status = SYMMETRIC, OR;

- * a Link Tuple with L_status = SYMMETRIC is removed, OR;

- * a Link Tuple with L_status = SYMMETRIC changes to having L_status = HEARD or L_status = LOST, OR;

- * a Link Tuple with L_status = HEARD or L_status = LOST changes to having L_status = SYMMETRIC, OR;

- * a 2-Hop Tuple is added or removed, OR;

- * the N_willingness of a Neighbor Tuple with N_symmetric = true changes from WILL_NEVER to any other value, OR;

- * the N_willingness of a Neighbor Tuple with N_symmetric = true and N_mpr = true changes to WILL_NEVER from any other value, OR;

Internet-Draft

OLSRv2

March 2009

* the N_willingness of a Neighbor Tuple with N_symmetric = true and N_mpr = false changes to WILL_ALWAYS from any other value.

3. Otherwise the set of MPRs of a router MAY be recalculated if the N_willingness of a Neighbor Tuple with N_symmetric = true changes in any other way; it SHOULD be recalculated if N_mpr = false and this is an increase in N_willingness or if N_mpr = true and this is a decrease in N_willingness.

If the set of MPRs of a router is recalculated, this MUST be as described in Section 14. Before that calculation, the N_mpr of all Neighbor Tuples are set false. After that calculation the N_mpr of all Neighbor Tuples representing symmetric 1-hop neighbors which are chosen as MPRs, are set true.

11. TC Message Generation

A router with one or more OLSRv2 interfaces, and with a non-empty Advertised Neighbor Set or a non-empty Local Attached Network Set MUST generate TC messages. A router with an empty Advertised Neighbor Set and empty Local Attached Network Set SHOULD also generate "empty" TC messages for a period A_HOLD_TIME after it last generated a non-empty TC message. TC messages (non-empty and empty) are generated according to the following:

1. The message hop count, if included, MUST be set to zero.
2. The message hop limit MUST be set to a value greater than 1. A router MAY use the same hop limit TC_HOP_LIMIT in all TC messages, or use different values of the hop limit TC_HOP_LIMIT in TC messages, see Section 5.6.
3. The message MUST contain a Message TLV with Type := CONT_SEQ_NUM and Value := ANSN from the Advertised Neighbor Set. If the TC message is complete then this Message TLV MUST have Type Extension := COMPLETE, otherwise it MUST have Type Extension := INCOMPLETE.
4. The message MUST contain a Message TLV with Type := VALIDITY_TIME, as specified in [timetlv]. If all TC messages are sent with the same hop limit then this TLV MUST have Value := T_HOLD_TIME. If TC messages are sent with different hop limits (more than one value of TC_HOP_LIMIT) then this TLV MUST specify times which vary with the number of hops distance appropriate to the chosen pattern of TC message hop limits, as specified in [timetlv], these times SHOULD be appropriate multiples of T_HOLD_TIME.

Internet-Draft

OLSRv2

March 2009

5. The message MAY contain a Message TLV with Type := INTERVAL_TIME, as specified in [timetlv]. If all TC messages are sent with the same hop limit then this TLV MUST have Value := TC_INTERVAL. If TC messages are sent with different hop limits, then this TLV MUST specify times which vary with the number of hops distance appropriate to the chosen pattern of TC message hop limits, as specified in [timetlv], these times SHOULD be appropriate multiples of TC_INTERVAL.
6. Unless the router has a single interface, with a single interface address with maximum prefix length, and that address is the router's originator address, the message MUST contain all of the router's interface addresses (i.e. all addresses in an I_local_iface_addr_list) in its Address Blocks.
7. All addresses of the router's interfaces that are included in an Address Block MUST each be associated with a TLV with Type := LOCAL_IF and Value := UNSPEC_IF.
8. A complete message MUST include, and an incomplete message MAY include, in its Address Blocks:
 1. Each A_neighbor_iface_addr from the Advertised Neighbor Set;
 2. AL_net_addr from each Local Attached Neighbor Tuple, each associated with a TLV with Type := GATEWAY and Value := AL_dist.

11.1. TC Message: Transmission

Complete TC messages are generated and transmitted periodically on all OLSRV2 interfaces, with a default interval between two consecutive TC transmissions by the same router of TC_INTERVAL.

TC messages MAY be generated in response to a change of contents, indicated by a change in ANSN. In this case a router MAY send a complete TC message, and if so MAY re-start its TC message schedule. Alternatively a router MAY send an incomplete TC message with at least the new content in its Address Blocks. Note that a router cannot report removal of advertised content using an incomplete TC message.

When sending a TC message in response to a change of contents, a router must respect a minimum interval of TC_MIN_INTERVAL between generated TC messages. Sending an incomplete TC message MUST NOT cause the interval between complete TC messages to be increased, and thus a router MUST NOT send an incomplete TC message if within TC_MIN_INTERVAL of the next scheduled complete TC message.

Internet-Draft

OLSRv2

March 2009

The generation of TC messages, whether scheduled or triggered by a change of contents MAY be jittered as described in [RFC5148]. The values of MAXJITTER used SHOULD be:

- o TP_MAXJITTER for periodic TC message generation;
- o TT_MAXJITTER for responsive TC message generation.

TC messages are included in packets as specified in [RFC5444]. These packets MAY contain other messages, including HELLO messages and TC messages with different originator addresses. TC messages are forwarded according to the specification in Section 7.3.

12. TC Message Processing

On receiving a TC message, a router MUST first check if the message is invalid for processing by this router, as defined in Section 12.1. Otherwise the receiving router MUST update its appropriate Interface Information Base and its Router Information Base as specified in Section 12.2.

All TC message processing, including determination of whether a message is invalid, unless otherwise noted considers only TLVs with Type Extension = 0. TLVs with any other type extension (or any unmentioned type extension when other type extensions are considered) are ignored. All references to, for example, a TLV with Type = VALIDITY_TIME refer to a TLV with Type = VALIDITY_TIME and Type Extension = 0.

12.1. Invalid Message

A received TC message is invalid for processing by this router if any of the following conditions are true.

- o The Message Header does not include an originator address, a message sequence number, and at least one of a hop limit and a hop count.
- o The message does not have a TLV with Type = VALIDITY_TIME in its Message TLV Block.
- o The message has more than one TLV with Type = VALIDITY_TIME in its Message TLV Block, and these TLVs indicate different validity times, as specified by [timetlv].
- o The message has more than one TLV with Type = INTERVAL_TIME in its Message TLV Block, and these TLVs indicate different interval times, as specified by [timetlv].

Internet-Draft

OLSRv2

March 2009

- o The message does not have a TLV with Type = CONT_SEQ_NUM and Type Extension = COMPLETE or Type Extension = INCOMPLETE in its Message TLV Block.
- o The message has more than one TLV with Type = CONT_SEQ_NUM and Type Extension = COMPLETE or Type Extension = INCOMPLETE in its Message TLV Block, and these do not have the same type extension and the same Value.
- o The message has any Address Block TLV(s) with Type = LOCAL_IF and any single value(s) which are not equal to UNSPEC_IF.
- o Any address associated with a TLV with Type = LOCAL_IF is one of the receiving router's current or recently used interface addresses (i.e. is in any I_local_iface_addr_list in the Local Interface Set or is equal to any IR_local_iface_addr in the Removed Interface Address Set).
- o Any address (including different copies of an address, in the same or different Address Blocks) is associated with more than one single value by one or more TLV(s) with Type = GATEWAY.

An invalid message MUST be silently discarded, without updating the router's Information Bases. A router MAY recognize additional reasons for identifying that a message is badly formed and discard such messages.

12.2. Initial TC Message Processing

When, according to Section 7.2, a TC message is to be "processed according to its type", this means that:

- o If the TC message contains a Message TLV with Type = CONT_SEQ_NUM and Type Extension = COMPLETE, then processing according to Section 12.3 and then according to Section 12.4 is carried out.
- o If the TC message contains a Message TLV with Type = CONT_SEQ_NUM and Type Extension = INCOMPLETE, then only processing according to Section 12.3 is carried out.

For the purposes of this section:

- o "originator address" refers to the originator address in the TC Message Header.
- o "validity time" is calculated from a VALIDITY_TIME Message TLV in the TC message according to the specification in [timetlv]. All information in the TC message has the same validity time.

Internet-Draft

OLSRv2

March 2009

- o "ANSN" is defined as being the Value of a Message TLV with Type = CONT_SEQ_NUM.
- o "sending address list" refers to the list of addresses in all Address Blocks which have associated TLV(s) with Type = LOCAL_IF and Value = UNSPEC_IF. If the sending address list is otherwise empty, then the message's originator address is added to the sending address list, with maximum prefix length.
- o Comparisons of sequence numbers are carried out as specified in Section 18.

12.3. Initial TC Message Processing

The TC message is processed as follows:

1. The Advertising Remote Router Set is updated according to Section 12.3.1; if the TC message is indicated as discarded in that processing then the following steps are not carried out.
2. The Topology Set is updated according to Section 12.3.2.
3. The Attached Network Set is updated according to Section 12.3.3.

12.3.1. Populating the Advertising Remote Router Set

The router MUST update its Advertising Remote Router Set as follows:

1. If there is an Advertising Remote Router Tuple with:
 - * AR_orig_addr = originator address; AND
 - * AR_seq_number > ANSNthen the TC message MUST be discarded.
2. Otherwise:
 1. If there is no Advertising Remote Router Tuple such that:
 - + AR_orig_addr = originator address;then create an Advertising Remote Router Tuple with:
 - + AR_orig_addr := originator address.

Internet-Draft

OLSRv2

March 2009

2. This Advertising Remote Router Tuple (existing or new, the "current tuple") is then modified as follows:
 - + AR_seq_number := ANSN;
 - + AR_time := current time + validity time.
 - + AR_iface_addr_list := sending address list
3. For each other Advertising Remote Router Tuple (with a different AR_orig_addr, the "other tuple") whose AR_iface_addr_list contains any address in the AR_iface_addr_list of the current tuple:
 1. remove all Topology Tuples with T_orig_addr = AR_orig_addr of the other tuple;
 2. remove all Attached Network Tuples with AN_orig_addr = AR_orig_addr of the other tuple;
 3. remove the other tuple.

12.3.2. Populating the Topology Set

The router MUST update its Topology Set as follows:

1. For each address (henceforth advertised address) in an Address Block that does not have an associated TLV with Type = LOCAL_IF, or an associated TLV with Type = GATEWAY:
 1. If there is no Topology Tuple such that:
 - + T_dest_iface_addr = advertised address; AND
 - + T_orig_addr = originator addressthen create a new Topology Tuple with:
 - + T_dest_iface_addr := advertised address;
 - + T_orig_addr := originator address.
 2. This Topology Tuple (existing or new) is then modified as follows:
 - + T_seq_number := ANSN;

Internet-Draft

OLSRv2

March 2009

+ T_time := current time + validity time.

12.3.3. Populating the Attached Network Set

The router MUST update its Attached Network Set as follows:

1. For each address (henceforth network address) in an Address Block that does not have an associated TLV with Type = LOCAL_IF, and does have an associated TLV with Type = GATEWAY:

1. If there is no Attached Network Tuple such that:

+ AN_net_addr = network address; AND

+ AN_orig_addr = originator address

then create a new Attached Network Tuple with:

+ AN_net_addr := network address;

+ AN_orig_addr := originator address

2. This Attached Network Tuple (existing or new) is then modified as follows:

+ AN_dist := the value of the associated GATEWAY TLV;

+ AN_seq_number := ANSN;

+ AN_time := current time + validity time.

12.4. Completing TC Message Processing

The TC message is processed as follows:

1. The Topology Set is updated according to Section 12.4.1.
2. The Attached Network Set is updated according to Section 12.4.2.

12.4.1. Purging the Topology Set

The Topology Set MUST be updated as follows:

1. Any Topology Tuples with:

* T_orig_addr = originator address; AND

Internet-Draft

OLSRv2

March 2009

* T_seq_number < ANSN

MUST be removed.

12.4.2. Purging the Attached Network Set

The Attached Network Set MUST be updated as follows:

1. Any Attached Network Tuples with:

* AN_orig_addr = originator address; AND

* AN_seq_number < ANSN

MUST be removed.

13. Information Base Changes

1. The Originator Set in the Local Information Base MUST be updated when the router changes originator address. If there is no Originator Tuple with:

* O_orig_addr = old originator address

then create an Originator Tuple with:

* O_orig_addr := old originator address

This Originator Tuple (existing or new) is then modified as follows:

* O_time := current time + O_HOLD_TIME

2. The Topology Information Base MUST be changed when an Advertising Remote Router Tuple expires (AR_time is reached). The following changes are required before the Advertising Remote Router Tuple is removed:

1. All Topology Tuples with:

+ T_orig_addr = AR_orig_addr of the Advertising Remote Router Tuple

are removed.

2. All Attached Network Tuples with:

Internet-Draft

OLSRv2

March 2009

+ AN_orig_addr = AR_orig_addr of the Advertising Remote Router Tuple

are removed.

14. Selecting MPRs

Each router MUST select, from among its willing symmetric 1-hop neighbors, a subset of routers as MPRs. MPRs are used to flood control messages from a router into the network, while reducing the number of retransmissions that will occur in a region. Thus, the concept of MPR flooding is an optimization of a classical flooding mechanism. MPRs MAY also be used to reduce the shared topology information in the network. Consequently, while it is not essential that the set of MPRs is minimal, keeping the number of MPRs small ensures that the overhead of OLSRv2 is kept at a minimum.

A router MUST select MPRs for each of its OLSRv2 interfaces, but then forms the union of those sets as its single set of MPRs. This union MUST include all symmetric 1-hop neighbors with willingness WILL_ALWAYS. Only this overall set of MPRs is relevant, the recorded and used MPR relationship is one of routers, not interfaces. Routers MAY select their MPRs by any process which satisfies the conditions which follow. Routers can freely interoperate whether they use the same or different MPR selection algorithms.

For each OLSRv2 interface a router MUST select a set of MPRs. This set MUST have the properties that:

- o All of the selected MPRs are willing symmetric 1-hop neighbors, AND;
- o If the selecting router sends a message on that OLSRv2 interface, and that message is successfully forwarded by all of the selected MPRs for that interface, then all symmetric strict 2-hop neighbors of the selecting router through that OLSRv2 interface will receive that message on a symmetric link.

Note that it is always possible to select a valid set of MPRs. The set of all willing symmetric 1-hop neighbors of a router is a (maximal) valid set of MPRs for that router. However a router SHOULD NOT select a symmetric 1-hop neighbor with Willingness != WILL_ALWAYS as an MPR if there are no symmetric strict 2-hop neighbors with a symmetric link to that symmetric 1-hop neighbor. Thus a router with no symmetric 1-hop neighbors with willingness WILL_ALWAYS and with no symmetric strict 2-hop neighbors SHOULD NOT select any MPRs.

A router MAY select its MPRs for each OLSRv2 interface independently,

Internet-Draft

OLSRv2

March 2009

or it MAY coordinate its MPR selections across its OLSRV2 interfaces, as long as the required condition is satisfied for each OLSRV2 interface. Each router MAY select its MPRs independently from the MPR selection by other routers, or it MAY, for example, give preference to routers that either are, or are not, already selected as MPRs by other routers.

When selecting MPRs for each OLSRV2 interface independently, this MAY be done using information from the Link Set and 2-Hop Set of that OLSRV2 interface, and the Neighbor Set of the router (specifically the N_willingness elements).

The selection of MPRs (overall, not per OLSRV2 interface) is recorded in the Neighbor Set of the router (using the N_mpr elements). A selected MPR MUST be a willing symmetric 1-hop neighbor (i.e. the corresponding N_symmetric = true, and the corresponding N_willingness != WILL_NEVER).

A router MUST recalculate its MPRs whenever the currently selected set of MPRs does not still satisfy the required conditions. It MAY recalculate its MPRs if the current set of MPRs is still valid, but could be more efficient. It is sufficient to recalculate a router's MPRs when there is a change to any of the router's Link Sets affecting the symmetry of any link (addition or removal of a Link Tuple with L_status = SYMMETRIC, or change of any L_status to or from SYMMETRIC), any change to any of the router's 2-Hop Sets, or a change of the N_willingness (to or from WILL_NEVER or to WILL_ALWAYS is sufficient) of any Neighbor Tuple with N_symmetric = true.

An algorithm that creates a set of MPRs that satisfies the required conditions is given in Appendix B.

15. Populating Derived Sets

The Relay Sets and the Advertised Neighbor Set of a router are denoted derived sets, since updates to these sets are not directly a function of message exchanges, but rather are derived from updates to other sets, in particular to the MPR selector status of other routers recorded in the Neighbor Set.

15.1. Populating the Relay Set

The Relay Set for an OLSRV2 interface contains the set of OLSRV2 interface addresses of those symmetric 1-hop neighbors for which this OLSRV2 interface is to relay broadcast traffic. This set MUST contain only addresses of OLSRV2 interfaces with which this OLSRV2 interface has a symmetric link. This set MUST include all such addresses of all such OLSRV2 interfaces of routers which are MPR

Internet-Draft

OLSRv2

March 2009

selectors of this router.

The Relay Set for an OLSRV2 interface of this router is thus created by:

1. For each Link Tuple in the Link Set for this OLSRV2 interface with `L_status = SYMMETRIC`, and the corresponding Neighbor Tuple with `N_neighbor_iface_addr_list` containing `L_neighbor_iface_addr_list`:
 1. All addresses from `L_neighbor_iface_addr_list` MUST be included in the Relay Set of this OLSRV2 interface if `N_mpr_selector = true`, and otherwise MAY be so included.

15.2. Populating the Advertised Neighbor Set

The Advertised Neighbor Set of a router contains all interface addresses of those symmetric 1-hop neighbors to which the router advertises a link in its TC messages. This set MUST include all addresses in all MPR selector of this router.

The Advertised Neighbor Set for this router is thus created by:

1. For each Neighbor Tuple with `N_symmetric = true`:
 1. All addresses from `N_neighbor_iface_addr_list` MUST be included in the Advertised Neighbor Set if `N_mpr_selector = true`, and otherwise MAY be so included.

Whenever address(es) are added to or removed from the Advertised Neighbor Set, its ANSN MUST be incremented.

16. Routing Set Calculation

The Routing Set of a router is populated with Routing Tuples that represent paths from that router to all destinations in the network. These paths are calculated based on the Network Topology Graph, which is constructed from information in the Information Bases, obtained via HELLO and TC message exchange.

16.1. Network Topology Graph

The Network Topology Graph is formed from information from the router's Link Sets, Neighbor Set, Topology Set and Attached Network Set. The Network Topology Graph SHOULD also use information from the router's 2-Hop Sets. The Network Topology Graph forms that router's topological view of the network in form of a directed graph, containing the following arcs:

Internet-Draft

OLSRv2

March 2009

- o Local symmetric links - all arcs X -> Y such that:
 - * X is an address in the I_local_iface_addr_list of a Local Interface Tuple of this router, AND;
 - * Y is an address in the L_neighbor_iface_addr_list of a Link Tuple in the corresponding (to the OLSRv2 interface of that I_local_iface_addr_list) Link Set which has L_status = SYMMETRIC.
- o 2-hop symmetric links - all arcs Y -> Z such that:
 - * Y is an address in the L_neighbor_iface_addr_list of a Link Tuple, in any of the router's Link Sets, which has L_status = SYMMETRIC, AND;
 - * the Neighbor Tuple with Y in its N_neighbor_iface_addr_list has N_willingness not equal to WILL_NEVER, AND;
 - * Z is the N2_2hop_iface_addr of a 2-Hop Tuple in the 2-Hop Set corresponding to the OLSRv2 interface of the chosen Link Set.
- o Advertised symmetric links - all arcs U -> V such that there exists a Topology Tuple and a corresponding Advertising Remote Router Tuple (i.e. with AR_orig_addr = T_orig_addr) with:
 - * U is in the AR_iface_addr_list of the Advertising Remote Router Tuple, AND;
 - * V is the T_dest_iface_addr of the Topology Tuple.
- o Symmetric 1-hop neighbor addresses - all arcs Y -> W such that:
 - * Y is, and W is not, an address in the L_neighbor_iface_addr_list of a Link Tuple, in any of the router's Link Sets, which has L_status = SYMMETRIC, AND;
 - * W and Y are included in the same N_neighbor_iface_addr_list (i.e. the one in the Neighbor Tuple whose N_neighbor_iface_addr_list contains the L_neighbor_iface_addr_list that includes Y).
- o Attached network addresses - all arcs U -> T such that there exists an Attached Network Tuple and a corresponding Advertising Remote Router Tuple (i.e. with AR_orig_addr = AN_orig_addr) with:
 - * U is in the AR_iface_addr_list of the Advertising Remote Router Tuple, AND;

Internet-Draft

OLSRv2

March 2009

* T is the AN_net_addr of the Attached Network Tuple.

All links in the first three cases above have a hop count of one, the symmetric 1-hop neighbor addresses have a hop count of zero, and the attached network addresses have a hop count given by the appropriate value of AN_dist.

16.2. Populating the Routing Set

The Routing Set MUST contain the shortest paths for all destinations from all local OLSRV2 interfaces using the Network Topology Graph. This calculation MAY use any algorithm, including any means of choosing between paths of equal length.

Using the notation of Section 16.1, each path will have as its first arc a local symmetric link X -> Y. There will be a path for each terminating Y, Z, V, W and T which can be connected to local OLSRV2 interface address X using the indicated arcs. The corresponding Routing Tuple for this path will have:

- o R_dest_addr := the terminating Y, Z, V, W or T;
- o R_next_iface_addr := the first arc's Y;
- o R_dist := the total hop count of the path;
- o R_local_iface_addr := the first arc's X.

An example algorithm for calculating the Routing Set of a router is given in Appendix C.

16.3. Routing Set Updates

The Routing Set MUST be updated when changes in the Neighborhood Information Base or the Topology Information Base indicate a change of the known symmetric links and/or attached networks in the MANET. It is sufficient to consider only changes which affect at least one of:

- o The Link Set of any OLSRV2 interface, and to consider only Link Tuples which have, or just had, L_status = SYMMETRIC (including removal of such Link Tuples).
- o The Neighbor Set of the router, and to consider only Neighbor Tuples that have, or just had, N_symmetric = true.
- o The 2-Hop Set of any OLSRV2 interface.

Internet-Draft

OLSRv2

March 2009

- o The Advertising Remote Router Set of the router.
- o The Topology Set of the router.
- o The Attached Network Set of the router.

Updates to the Routing Set do not generate or trigger any messages to be transmitted. The state of the Routing Set SHOULD, however, be reflected in the IP routing table by adding and removing entries from the IP routing table as appropriate.

17. Proposed Values for Parameters and Constants

OLSRv2 uses all parameters and constants defined in [NHDP] and additional parameters and constants defined in this document. All but one (RX_HOLD_TIME) of these additional parameters are router parameters as defined in [NHDP]. These proposed values of the additional parameters are appropriate to the case where all parameters (including those defined in [NHDP]) have a single value. Proposed values for parameters defined in [NHDP] are given in that document.

17.1. Local History Time Parameters

- o O_HOLD_TIME := 30 seconds

17.2. Message Interval Parameters

- o TC_INTERVAL := 5 seconds
- o TC_MIN_INTERVAL := TC_INTERVAL/4

17.3. Advertised Information Validity Time Parameters

- o T_HOLD_TIME := 3 x TC_INTERVAL
- o A_HOLD_TIME := T_HOLD_TIME

17.4. Received Message Validity Time Parameters

- o RX_HOLD_TIME := 30 seconds
- o P_HOLD_TIME := 30 seconds
- o F_HOLD_TIME := 30 seconds

Internet-Draft

OLSRv2

March 2009

17.5. Jitter Time Parameters

- o TP_MAXJITTER := HP_MAXJITTER
- o TT_MAXJITTER := HT_MAXJITTER
- o F_MAXJITTER := TT_MAXJITTER

17.6. Hop Limit Parameter

- o TC_HOP_LIMIT := 255

17.7. Willingness Parameter and Constants

- o WILLINGNESS := WILL_DEFAULT
- o WILL_NEVER := 0
- o WILL_DEFAULT := 3
- o WILL_ALWAYS := 7

18. Sequence Numbers

Sequence numbers are used in OLSRv2 with the purpose of discarding "old" information, i.e. messages received out of order. However with a limited number of bits for representing sequence numbers, wrap-around (that the sequence number is incremented from the maximum possible value to zero) will occur. To prevent this from interfering with the operation of OLSRv2, the following MUST be observed when determining the ordering of sequence numbers.

The term MAXVALUE designates in the following one more than the largest possible value for a sequence number. For a 16 bit sequence number (as are those defined in this specification) MAXVALUE is 65536.

The sequence number S1 is said to be "greater than" the sequence number S2 if:

- o $S1 > S2$ AND $S1 - S2 < MAXVALUE/2$ OR
- o $S2 > S1$ AND $S2 - S1 > MAXVALUE/2$

When sequence numbers S1 and S2 differ by MAXVALUE/2 their ordering cannot be determined. In this case, which should not occur, either ordering may be assumed.

Internet-Draft

OLSRv2

March 2009

Thus when comparing two messages, it is possible - even in the presence of wrap-around - to determine which message contains the most recent information.

19. IANA Considerations

19.1. Message Types

This specification defines one Message Type, to be allocated from the 0-223 range of the "Message Types" namespace defined in [RFC5444], as specified in Table 5.

Name	Type	Description
TC	TBD1	Topology Control (MANET-wide signaling)

Table 5

19.2. Message TLV Types

This specification defines two 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 two new type extension registries with assignments as specified in Table 6 and Table 7. Specifications of these TLVs are in Section 8.1.1 and Section 8.2.1.

Name	Type	Type extension	Description
MPR_WILLING	TBD2	0	Specifies the originating router's willingness to act as a relay and to partake in network formation
		1-255	Expert Review

Table 6

Internet-Draft

OLSRv2

March 2009

Name	Type	Type extension	Description
CONT_SEQ_NUM	TBD3	0 (COMPLETE)	Specifies a content sequence number for this complete message
		1 (INCOMPLETE)	Specifies a content sequence number for this incomplete message
		2-255	Expert Review

Table 7

Type extensions indicated as Expert Review SHOULD be allocated as described in [RFC5444], based on Expert Review as defined in [RFC5226].

19.3. Address Block TLV Types

This specification defines two Address Block TLV Types, which must be allocated from the "Address Block TLV Types" namespace defined in [RFC5444]. IANA are requested to make allocations in the 8-127 range for these types. This will create two new type extension registries with assignments as specified in Table 8 and Table 9. Specifications of these TLVs are in Section 8.1.2 and Section 8.2.2.

Name	Type	Type extension	Description
MPR	TBD4	0	Specifies that a given address is of a router selected as an MPR
		1-255	Expert Review

Table 8

Name	Type	Type extension	Description
GATEWAY	TBD5	0	Specifies that a given address is reached via a gateway on the originating router
		1-255	Expert Review

Internet-Draft

OLSRv2

March 2009

Table 9

Type extensions indicated as Expert Review SHOULD be allocated as described in [RFC5444], based on Expert Review as defined in [RFC5226].

The Address Block TLV with Type = LOCAL_IF defined in [NHDP] is extended to also permit inclusion of the value UNSPEC_IF = 2, representing a local interface address which may or may not be that on which this message is transmitted.

20. Security Considerations

Currently, OLSRV2 does not specify any special security measures. As a proactive routing protocol, OLSRV2 makes a target for various attacks. The various possible vulnerabilities are discussed in this section.

20.1. Confidentiality

Being a proactive protocol, OLSRV2 periodically MPR floods topological information to all routers in the network. Hence, if used in an unprotected wireless network, the network topology is revealed to anyone who listens to OLSRV2 control messages.

In situations where the confidentiality of the network topology is of importance, regular cryptographic techniques, such as exchange of OLSRV2 control traffic messages encrypted by PGP [RFC4880] or encrypted by some shared secret key, can be applied to ensure that control traffic can be read and interpreted by only those authorized to do so.

20.2. Integrity

In OLSRV2, each router is injecting topological information into the network through transmitting HELLO messages and, for some routers, TC messages. If some routers for some reason, malicious or malfunction, inject invalid control traffic, network integrity may be compromised. Therefore, message authentication is recommended.

Different such situations may occur, for instance:

1. a router generates TC messages, advertising links to non-neighbor routers;
2. a router generates TC messages, pretending to be another router;

Internet-Draft

OLSRv2

March 2009

3. a router generates HELLO messages, advertising non-neighbor routers;
4. a router generates HELLO messages, pretending to be another router;
5. a router forwards altered control messages;
6. a router does not forward control messages;
7. a router does not select multipoint relays correctly;
8. a router forwards broadcast control messages unaltered, but does not forward unicast data traffic;
9. a router "replays" previously recorded control traffic from another router.

Authentication of the originator router for control messages (for situations 2, 4 and 5) and on the individual links announced in the control messages (for situations 1 and 3) may be used as a countermeasure. However to prevent routers from repeating old (and correctly authenticated) information (situation 9) temporal information is required, allowing a router to positively identify such delayed messages.

In general, digital signatures and other required security information may be transmitted as a separate OLSRv2 Message Type, or signatures and security information may be transmitted within the OLSRv2 HELLO and TC messages, using the TLV mechanism. Either option permits that "secured" and "unsecured" routers can coexist in the same network, if desired,

Specifically, the authenticity of entire OLSRv2 control packets can be established through employing IPsec authentication headers, whereas authenticity of individual links (situations 1 and 3) require additional security information to be distributed.

An important consideration is that all control messages in OLSRv2 are transmitted either to all routers in the neighborhood (HELLO messages) or broadcast to all routers in the network (TC messages).

For example, a control message in OLSRv2 is always a point-to-multipoint transmission. It is therefore important that the authentication mechanism employed permits that any receiving router can validate the authenticity of a message. As an analogy, given a block of text, signed by a PGP private key, then anyone with the corresponding public key can verify the authenticity of the text.

Internet-Draft

OLSRv2

March 2009

20.3. Interaction with External Routing Domains

OLSRv2 does, through the use of TC messages, provide a basic mechanism for injecting external routing information to the OLSRV2 domain. Appendix A also specifies that routing information can be extracted from the topology table or the routing table of OLSRV2 and, potentially, injected into an external domain if the routing protocol governing that domain permits.

Other than as described in Appendix A, when operating routers connecting OLSRV2 to an external routing domain, care MUST be taken not to allow potentially insecure and untrustworthy information to be injected from the OLSRV2 domain to external routing domains. Care MUST be taken to validate the correctness of information prior to it being injected as to avoid polluting routing tables with invalid information.

A recommended way of extending connectivity from an existing routing domain to an OLSRV2 routed MANET is to assign an IP prefix (under the authority of the routers/gateways connecting the MANET with the exiting routing domain) exclusively to the OLSRV2 MANET area, and to configure the gateways statically to advertise routes to that IP sequence to routers in the existing routing domain.

21. Contributors

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

- o Cedric Adjih, INRIA, France, <Cedric.Adjih@inria.fr>
- o Emmanuel Baccelli, INRIA , France, <Emmanuel.Baccelli@inria.fr>
- o Thomas Heide Clausen, LIX, France, <T.Clausen@computer.org>
- o Justin Dean, NRL, USA, <jdean@itd.nrl.navy.mil>
- o Christopher Dearlove, BAE Systems, UK, <chris.dearlove@baesystems.com>
- o Satoh Hiroki, Hitachi SDL, Japan, <hiroki.satoh.yj@hitachi.com>
- o Philippe Jacquet, INRIA, France, <Philippe.Jacquet@inria.fr>
- o Monden Kazuya, Hitachi SDL, Japan, <kazuya.monden.vw@hitachi.com>
- o Kenichi Mase, Niigata University, Japan, <mase@ie.niigata-u.ac.jp>

Internet-Draft

OLSRv2

March 2009

o Ryuji Wakikawa, KEIO University, Japan, <ryuji@sfc.wide.ad.jp>

22. Acknowledgments

The authors would like to acknowledge the team behind OLSRv1, specified in RFC3626, including Anis Laouiti (INT, Paris), Pascale Minet (INRIA, France), Laurent Viennot (INRIA, France), and Amir Qayyum (M.A. Jinnah University, Islamabad) for their contributions.

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): Khaldoun Al Agha (LRI), Song-Yean Cho (LIX), Alan Cullen (BAE Systems), Louise Lamont (CRC), Li Li (CRC), Joe Macker (NRL), Richard Ogier (SRI), Charles E. Perkins (WiChorus), Shubhranshu Singh (Samsung AIT), and the entire IETF MANET working group.

23. References

23.1. Normative References

- [RFC5444] Clausen, T., Dean, J., Dearlove, C., and C. Adjih, "Generalized MANET Packet/Message Format", RFC 5444, February 2009.
- [timetlv] Clausen, T. and C. Dearlove, "Representing multi-value time in MANETs", Work In Progress draft-ietf-manet-timetlv-08.txt, September 2008.
- [RFC5148] Clausen, T., Dearlove, C., and B. Adamson, "Jitter considerations in MANETs", RFC 5148, February 2008.
- [NHDP] Clausen, T., Dean, J., and C. Dearlove, "MANET Neighborhood Discovery Protocol (NHDP)", work in progress draft-ietf-manet-nhdp-08.txt, February 2009.
- [manet-iana] Chakeres, I., "IANA Allocations for MANET Protocols", Work In Progress draft-ietf-manet-iana-07.txt, November 2007.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, BCP 14, March 1997.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 5226, BCP 26, May 2008.

Internet-Draft

OLSRv2

March 2009

23.2. Informative References

- [RFC2501] Macker, J. and S. Corson, "Mobile Ad hoc Networking (MANET): Routing Protocol Performance Issues and Evaluation Considerations", RFC 2501, January 1999.
- [RFC3626] Clausen, T. and P. Jacquet, "The Optimized Link State Routing Protocol", RFC 3626, October 2003.
- [RFC4880] Callas, J., Donnerhacke, L., Finney, H., and R. Thayer, "OpenPGP message format", RFC 4880, November 2007.
- [HIPERLAN] ETSI, "ETSI STC-RES10 Committee. Radio equipment and systems: HIPERLAN type 1, functional specifications ETS 300-652", June 1996.
- [HIPERLAN2] Jacquet, P., Minet, P., Muhlethaler, P., and N. Rivierre, "Increasing reliability in cable free radio LANs: Low level forwarding in HIPERLAN.", 1996.
- [MPR] Qayyum, A., Viennot, L., and A. Laouiti, "Multipoint relaying: An efficient technique for flooding in mobile wireless networks.", 2001.
- [FSR] Pei, G., Gerla, M., and T. Chen, "Fisheye state routing in mobile ad hoc networks", 2000.
- [FSLs] Santivanez, C., Ramanathan, R., and I. Stavrakakis, "Making link-state routing scale for ad hoc networks", 2000.

Appendix A. Router Configuration

OLSRv2 does not make any assumption about router addresses, other than that each router is assumed to have at least one unique and routable IP address for each interface that it has which participates in the MANET.

When applicable, a recommended way of connecting an OLSRV2 network to an existing IP routing domain is to assign an IP prefix (under the authority of the routers/gateways connecting the MANET with the routing domain) exclusively to the OLSRV2 area, and to configure the gateways statically to advertise routes to that IP sequence to routers in the existing routing domain.

Internet-Draft

OLSRv2

March 2009

Appendix B. Example Algorithm for Calculating MPRs

The following specifies an algorithm which MAY be used to select MPRs. MPRs are calculated per OLSRV2 interface, but then a single set of MPRs is formed from the union of the MPRs for all OLSRV2 interfaces. (As noted in Section 14 a router MAY improve on this, by coordination between OLSRV2 interfaces.) A router's MPRs are recorded using the element `N_mpr` in Neighbor Tuples.

If using this algorithm then the following steps MUST be executed in order for a router to select its MPRs:

1. Set `N_mpr := false` in all Neighbor Tuples;
2. For each Neighbor Tuple with `N_symmetric = true` and `N_willingness = WILL_ALWAYS`, set `N_mpr := true`;
3. For each OLSRV2 interface of the router, use the algorithm in Appendix B.2. Note that this sets `N_mpr := true` for some Neighbor Tuples, these routers are already selected as MPRs when using the algorithm for following OLSRV2 interfaces.
4. OPTIONALLY, consider each selected MPR in turn, and if the set of selected MPRs without that router still satisfies the necessary conditions, for all OLSRV2 interfaces, then that router MAY be removed from the set of MPRs. This process MAY be repeated until no MPRs are removed. Routers MAY be considered in order of increasing `N_willingness`.

Symmetric 1-hop neighbor routers with `N_willingness = WILL_NEVER` MUST NOT be selected as MPRs, and MUST be ignored in the following algorithm, as MUST be symmetric 2-hop neighbor routers which are also symmetric 1-hop neighbor routers (i.e. when considering 2-Hop Tuples, ignore any 2-Hop Tuples whose `N2_2hop_iface_addr` is in the `N_neighbor_iface_addr_list` of any Neighbor Tuple, or whose `N2_neighbor_iface_addr_list` is included in the `N_neighbor_iface_addr_list` of any Neighbor Tuple with `N_willingness = WILL_NEVER`).

B.1. Terminology

The following terminology will be used when selecting MPRs for the OLSRV2 interface `I`:

`N(I)` - The set of symmetric 1-hop neighbors which have a symmetric link to `I`.

Internet-Draft

OLSRv2

March 2009

$N2(I)$ - The set of addresses of interfaces of a router with a symmetric link to a router in $N(I)$; this MAY be restricted to considering only information received over I (in which case $N2(I)$ is the set of $N2_2hop_iface_addr$ in 2-Hop Tuples in the 2-Hop Set for OLSRV2 interface I).

Connected to I via Y - An address A in $N2(I)$ is connected to I via a router Y in $N(I)$ if A is an address of an interface of a symmetric 1-hop neighbor of Y (i.e. A is the $N2_2hop_iface_addr$ in a 2-Hop Tuple in the 2-Hop Set for OLSRV2 interface I , and whose $N2_neighbor_iface_addr_list$ is contained in the set of interface addresses of Y).

$D(Y, I)$ - For a router Y in $N(I)$, the number of addresses in $N2(I)$ which are connected to I via Y .

$R(Y, I)$: - For a router Y in $N(I)$, the number of addresses in $N2(I)$ which are connected to I via Y , but are not connected to I via any router which has already been selected as an MPR.

B.2. MPR Selection Algorithm for each OLSRV2 Interface

When selecting MPRs for the OLSRV2 interface I :

1. For each address A in $N2(I)$ for which there is only one router Y in $N(I)$ such that A is connected to I via Y , select that router Y as an MPR (i.e. set $N_mpr := true$ in the Neighbor Tuple corresponding to Y).
2. While there exists any router Y in $N(I)$ with $R(Y, I) > 0$:
 1. Select a router Y in $N(I)$ with $R(Y, I) > 0$ in the following order of priority:
 - + greatest $N_willingness$ in the Neighbor Tuple corresponding to Y , THEN;
 - + greatest $R(Y, I)$, THEN;
 - + greatest $D(Y, I)$, THEN;
 - + $N_mpr_selector$ is equal to true, if possible, THEN;
 - + any choice.
 2. Select Y as an MPR (i.e. set $N_mpr := true$ in the Neighbor Tuple corresponding to Y).

Internet-Draft

OLSRv2

March 2009

Appendix C. Example Algorithm for Calculating the Routing Set

The following procedure is given as an example for calculating the Routing Set using a variation of Dijkstra's algorithm. First all Routing Tuples are removed, and then the procedures in the following sections are applied in turn.

C.1. Add Local Symmetric Links

1. For each Local Interface Tuple:
 1. Select an address (the "local address") in `I_local_iface_addr_list`.
 2. For each Link Tuple for this local interface with `L_status = SYMMETRIC`:
 1. For each address (the "current address") in `L_neighbor_iface_addr_list`, if there is no Routing Tuple with `R_dest_addr = current address`, then add a Routing Tuple with:
 - `R_dest_addr := current address`;
 - `R_next_iface_addr := current address`;
 - `R_dist := 1`;
 - `R_local_iface_addr := local address`.
 2. For each Neighbor Tuple whose `N_neighbor_iface_addr_list` contains the `R_dest_addr` of a Routing Tuple (the "previous Tuple"):
 1. For each address (the "current address") in `N_neighbor_iface_addr_list`, if there is no Routing Tuple with `R_dest_addr = current address`, then add a Routing Tuple with:
 - + `R_dest_addr := current address`;
 - + `R_next_iface_addr := R_dest_addr of the previous Tuple`;
 - + `R_dist := 1`;
 - + `R_local_iface_addr := R_local_iface_addr of the previous Tuple`.

Internet-Draft

OLSRv2

March 2009

C.2. Add Remote Symmetric Links

The following procedure, which adds Routing Tuples for destination routers $h+1$ hops away, MUST be executed for each value of h , starting with $h := 1$ and incrementing by 1 for each iteration. The execution MUST stop if no new Routing Tuples are added in an iteration.

1. For each Topology Tuple, if:

- * $T_dest_iface_addr$ is not equal to R_dest_addr of any Routing Tuple, AND;
- * for the Advertising Remote Router Tuple with $AR_orig_addr = T_orig_addr$, there is an address in the $AR_iface_addr_list$ which is equal to the R_dest_addr of a Routing Tuple (the "previous Routing Tuple") whose $R_dist = h$

then add a new Routing Tuple, with:

- * $R_dest_addr := T_dest_iface_addr$;
- * $R_next_iface_addr := R_next_iface_addr$ of the previous Routing Tuple;
- * $R_dist := h+1$;
- * $R_local_iface_addr := R_local_iface_addr$ of the previous Routing Tuple.

More than one Topology Tuple may be usable to select the next hop $R_next_iface_addr$ for reaching the address R_dest_addr . Ties should be broken such that routers with greater willingness are preferred, and between routers of equal willingness, MPR selectors are preferred over non-MPR selectors.

2. After the above iteration has completed, if $h = 1$, for each 2-Hop Neighbor Tuple where:

- * $N2_2hop_iface_addr$ is not equal to R_dest_addr of any Routing Tuple, AND;
- * The Neighbor Tuple whose $N_neighbor_iface_addr_list$ contains $N2_neighbor_iface_addr_list$ has $N_willingness$ not equal to $WILL_NEVER$

select a Routing Tuple (the "previous Routing Tuple") whose R_dest_addr is contained in $N2_neighbor_iface_addr_list$, and add a new Routing Tuple with:

Internet-Draft

OLSRv2

March 2009

- * R_dest_addr := N2_2hop_iface_addr;
- * R_next_iface_addr := R_next_iface_addr of the previous Routing Tuple;
- * R_dist := 2;
- * R_local_iface_addr := R_local_iface_addr of the previous Routing Tuple.

More than one 2-Hop Neighbor Tuple may be usable to select the next hop R_next_iface_addr for reaching the address R_dest_addr. Ties should be broken such that routers with greater willingness are preferred, and between routers of equal willingness, MPR selectors are preferred over non-MPR selectors.

C.3. Add Attached Networks

1. For each Attached Network Tuple, if for the Advertising Remote Router Tuple with AR_orig_addr = AN_orig_addr, there is an address in the AR_iface_addr_list which is equal to the R_dest_addr of a Routing Tuple (the "previous Routing Tuple"), then:
 1. If there is no Routing Tuple with R_dest_addr = AN_net_addr, then add a new Routing Tuple with:
 - + R_dest_addr := AN_net_addr;
 - + R_next_iface_addr := R_next_iface_addr of the previous Routing Tuple;
 - + R_dist := (R_dist of the previous Routing Tuple) + AN_dist;
 - + R_local_iface_addr := R_local_iface_addr of the previous Routing Tuple.
 2. Otherwise if the Routing Tuple with R_dest_addr = AN_net_addr (the "current Routing Tuple") has R_dist > (R_dist of the previous Routing Tuple) + AN_dist, then modify the current Routing Tuple by:
 - + R_next_iface_addr := R_next_iface_addr of the previous Routing Tuple;
 - + R_dist := (R_dist of the previous Routing Tuple) + AN_dist;

Internet-Draft

OLSRv2

March 2009

- + R_local_iface_addr := R_local_iface_addr of the previous Routing Tuple.

Appendix D. Example Message Layout

An example TC message is as follows. The message has full Message Header (four bit flags field value is 15). Its four bit Message Address Length field has value 3 and hence addresses in the message have length four octets, here being IPv4 addresses. The overall message length is 65 octets.

The message has a Message TLV Block with content length 13 octets containing three TLVs. The first two TLVs are interval and validity times for the message. The third TLV is the content sequence number TLV used to carry the 2 octet ANSN, and (with default type extension zero, i.e. COMPLETE) indicating that the TC message is complete. Each TLV uses a TLV with flags octet value 16, indicating that it has a value, but no type extension or start and stop indexes. The first two TLVs have a value length of 1 octet, the last has a value length of 2 octets.

The message has two Address Blocks. The first Address Block contains 6 addresses, with flags octet value 128, hence with a Head section, (with length 2 octets) but no Tail section, and hence Mid sections with length two octets. The following TLV Block (content length 6 octets) contains a single LOCAL_IF TLV (flags octet value 48) indicating that the first three addresses (indexes 0 to 2) are associated with the value (length 1 octet) UNSPEC_IF, i.e. they are the originating router's local interface addresses. The remaining three addresses have no associated TLV, they are the interface addresses of advertised neighbors.

The second Address Block contains 1 address, with flags octet 176 indicating that there is a Head section (with length 2 octets), that the Tail section (length 2 octets) consists of zero valued octets (not included), and that there is a single prefix length, which is 16. The network address is thus Head.0.0/16. The following TLV Block (content length 8 octets) includes one TLV that indicates that the originating router is a gateway to this network, at a given number of hops distance (value length 1 octet). The TLV flags octet value of 16 indicates that no indexes are needed.

Internet-Draft

OLSRv2

March 2009

```

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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          TC          |1 1 1 1 0 0 1 1|0 0 0 0 0 0 0 0 0 1 0 0 0 0 0 1|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Originator Address          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Hop Limit | Hop Count | Message Sequence Number |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0 0 0 0 0 0 0 0 0 0 0 0 1 1 0 1| INTERVAL_TIME |0 0 0 1 0 0 0 0|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0 0 0 0 0 0 0 1| Value | VALIDITY_TIME |0 0 0 1 0 0 0 0|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0 0 0 0 0 0 0 1| Value | CONT_SEQ_NUM |0 0 0 1 0 0 0 0|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0 0 0 0 0 0 1 0| Value (ANSN) |0 0 0 0 0 1 1 0|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|1 0 0 0 0 0 0 0|0 0 0 0 0 0 1 0|          Head          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Mid          |          Mid          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Mid          |          Mid          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Mid          |          Mid          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0 0 0 0 0 0 0 0 0 0 0 0 0 1 1 0| LOCAL_IF |0 0 1 1 0 0 0 0|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0 0 0 0 0 0 0 0|0 0 0 0 0 0 1 0|0 0 0 0 0 0 0 1| UNSPEC_IF |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0 0 0 0 0 0 0 1|1 0 1 1 0 0 0 0|0 0 0 0 0 0 1 0|          Head          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Head (cont) |0 0 0 0 0 0 1 0|0 0 0 1 0 0 0 0|0 0 0 0 0 0 0 0|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0 0 0 0 0 1 0 0| GATEWAY |0 0 0 1 0 0 0 0|0 0 0 0 0 0 0 1|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Number Hops |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

Appendix E. Constraints

Any process which updates the Local Information Base, the Neighborhood Information Base or the Topology Information Base MUST ensure that all constraints specified in this appendix are maintained, as well as those specified in [NHDP].

In each Originator Tuple:

Internet-Draft

OLSRv2

March 2009

- o O_orig_addr MUST NOT equal any other O_orig_addr.
- o O_orig_addr MUST NOT equal this router's originator address.

In each Local Attached Network Tuple:

- o AL_net_addr MUST NOT equal any other AL_net_addr.
- o AL_net_addr MUST NOT be in the I_local_iface_addr_list of any Local Interface Tuple or be equal to the IR_local_iface_addr of any Removed Interface Address Tuple.
- o AL_dist MUST NOT be less than zero.

In each Link Tuple:

- o L_neighbor_iface_addr_list MUST NOT contain the AL_net_addr of any Local Attached Network Tuple.
- o If L_status = SYMMETRIC and the Neighbor Tuple whose N_neighbor_iface_addr_list contains L_neighbor_iface_addr_list has N_mpr_selector = true, then, for each address in this L_neighbor_iface_addr_list, there MUST be an equal RY_neighbor_iface_addr in the Relay Set associated with the same OLSRv2 interface.

In each Neighbor Tuple:

- o N_neighbor_iface_addr_list MUST NOT contain the AL_net_addr of any Local Attached Network Tuple.
- o If N_willingness MUST be in the range from WILL_NEVER to WILL_ALWAYS, inclusive.
- o If N_mpr = true, then N_symmetric MUST be true and N_willingness MUST NOT equal WILL_NEVER.
- o If N_symmetric = true and N_mpr = false, then N_willingness MUST NOT equal WILL_ALWAYS.
- o If N_mpr_selector = true, then N_symmetric MUST be true.
- o If N_mpr_selector = true, then, for each address in this N_neighbor_iface_addr_list, there MUST be an equal A_neighbor_iface_addr in the Advertised Neighbor Set.

In each Lost Neighbor Tuple:

Internet-Draft

OLSRv2

March 2009

- o NL_neighbor_iface_addr MUST NOT equal the AL_net_addr of any Local Attached Network Tuple.

In each 2-Hop Tuple:

- o N2_2hop_iface_addr MUST NOT equal the AL_net_addr of any Local Attached Network Tuple.

In each Received Tuple:

- o RX_orig_addr MUST NOT equal this router's originator address or any O_orig_addr.
- o Each ordered triple (RX_type, RX_orig_addr, RX_seq_number) MUST NOT equal the corresponding triple in any other Received Tuple in the same Received Set.

In each Processed Tuple:

- o P_orig_addr MUST NOT equal this router's originator address or any O_orig_addr.
- o Each ordered triple (P_type, P_orig_addr, P_seq_number) MUST NOT equal the corresponding triple in any other Processed Tuple.

In each Forwarded Tuple:

- o F_orig_addr MUST NOT equal this router's originator address or any O_orig_addr.
- o Each ordered triple (F_type, F_orig_addr, F_seq_number) MUST NOT equal the corresponding triple in any other Forwarded Tuple.

In each Relay Tuple:

- o RY_neighbor_iface_addr MUST NOT equal the RY_neighbor_iface_addr in any other Relay Tuple in the same Relay Set.
- o RY_neighbor_iface_addr MUST be in the L_neighbor_iface_addr_list of a Link Tuple with L_status = SYMMETRIC.

In the Advertised Neighbor Set:

- o Each A_neighbor_iface_addr MUST NOT equal any other A_neighbor_iface_addr.
- o Each A_neighbor_iface_addr MUST be in the N_neighbor_iface_addr_list of a Neighbor Tuple with N_symmetric =

Internet-Draft

OLSRv2

March 2009

true.

In each Advertising Remote Router Tuple:

- o AR_orig_addr MUST NOT equal this router's originator address or any O_orig_addr.
- o AR_orig_addr MUST NOT equal the AR_orig_addr in any other ANSN History Tuple.
- o AR_iface_addr_list MUST NOT be empty.
- o AR_iface_addr_list MUST NOT contain any duplicated addresses.
- o AR_iface_addr_list MUST NOT contain any address which is in the I_local_iface_addr_list of any Local Interface Tuple or be equal to the IR_local_iface_addr of any Removed Interface Address Tuple.
- o AR_iface_addr_list MUST NOT contain any address which is the AL_net_addr of any Local Attached Network Tuple.

In each Topology Tuple:

- o T_dest_iface_addr MUST NOT be in the I_local_iface_addr_list of any Local Interface Tuple or be equal to the IR_local_iface_addr of any Removed Interface Address Tuple.
- o T_dest_iface_addr MUST NOT equal the AL_net_addr of any Local Attached Network Tuple.
- o There MUST be an Advertising Remote Router Tuple with AR_orig_addr = T_orig_addr.
- o T_dest_iface_addr MUST NOT be in the AR_iface_addr_list of the Advertising Remote Router Tuple with AR_orig_addr = T_orig_addr.
- o T_seq_number MUST NOT be greater than AR_seq_number of the Advertising Remote Router Tuple with AR_orig_addr = T_orig_addr.
- o The ordered pair (T_dest_iface_addr, T_orig_addr) MUST NOT equal the corresponding pair in any other Topology Tuple.

In each Attached Network Tuple:

- o AN_net_addr MUST NOT be in the I_local_iface_addr_list of any Local Interface Tuple or be equal to the IR_local_iface_addr of any Removed Interface Address Tuple.

Internet-Draft

OLSRv2

March 2009

- o AN_net_addr MUST NOT equal the AL_net_addr of any Local Attached Network Tuple.
- o There MUST be an Advertising Remote Router Tuple with AR_orig_addr = AN_orig_addr.
- o AN_seq_number MUST NOT be greater than AR_seq_number of the Advertising Remote Router Tuple with AR_orig_addr = AN_orig_addr.
- o AN_dist MUST NOT be less than zero.
- o The ordered pair (AN_net_addr, AN_orig_addr) MUST NOT equal the corresponding pair in any other Attached Network Tuple.

Appendix F. Flow and Congestion Control

Due to its proactive nature, the OLSRv2 protocol has a natural control over the flow of its control traffic. Routers transmit control messages at predetermined rates specified and bounded by message intervals.

OLSRv2 employs [NHDP] for local signaling, embedding MPR selection advertisement through a simple Address Block TLV, and router willingness advertisement (if any) as a single Message TLV. OLSRv2 local signaling, therefore, shares the characteristics and constraints of [NHDP].

Furthermore, MPR flooding greatly reduces signaling overhead from from link state information dissemination in two ways. First, the amount of link state information for a router to declare is reduced to only contain that router's MPR selectors. This reduces the size of a link state declaration as compared to declaring full link state information. In particular some routers may not need to declare any such information. Second, using MPR flooding, the cost of distributing link state information throughout the network is greatly reduced, as compared to when using classic flooding, since only MPRs need to forward link state declaration messages. In dense networks, the reduction of control traffic can be of several orders of magnitude compared to routing protocols using classical flooding [MPR]. This feature naturally provides more bandwidth for useful data traffic and pushes further the frontier of congestion.

Since the control traffic is continuous and periodic, it keeps the quality of the links used in routing more stable. However, using certain OLSRv2 options, some control messages (HELLO messages or TC messages) may be intentionally sent in advance of their deadline in order to increase the responsiveness of the protocol to topology changes. This may cause a small, temporary, and local increase of

Internet-Draft

OLSRv2

March 2009

control traffic, however this is at all times bounded by the use of minimum message intervals.

Authors' Addresses

Thomas Heide Clausen
LIX, Ecole Polytechnique

Phone: +33 6 6058 9349
EMail: T.Clausen@computer.org
URI: <http://www.ThomasClausen.org/>

Christopher Dearlove
BAE Systems ATC

Phone: +44 1245 242194
EMail: chris.dearlove@baesystems.com
URI: <http://www.baesystems.com/>

Philippe Jacquet
Project Hipercom, INRIA

Phone: +33 1 3963 5263
EMail: philippe.jacquet@inria.fr

The OLSRV2 Design Team
MANET Working Group

