

Internet Engineering Task Force (IETF)
Request for Comments: 7185
Category: Informational
ISSN: 2070-1721

C. Dearlove
BAE Systems ATC
T. Clausen
LIX, Ecole Polytechnique
P. Jacquet
Alcatel-Lucent Bell Labs
April 2014

Rationale for the Use of Link Metrics
in the Optimized Link State Routing Protocol Version 2 (OLSRv2)

Abstract

The Optimized Link State Routing Protocol version 2 (OLSRv2) includes the ability to assign metrics to links and to use those metrics to allow routing by other than minimum hop count routes. This document provides a historic record of the rationale for, and design considerations behind, how link metrics were included in OLSRv2.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7185>.

Copyright Notice

Copyright (c) 2014 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 (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must

include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

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.

Table of Contents

1. Introduction	3
2. Terminology	5
3. Applicability	5
4. Motivational Scenarios	5
5. Link Metrics	7
5.1. Link Metric Properties	7
5.2. Link Metric Types	8
5.3. Directional Link Metrics	10
5.4. Reporting Link and Neighbor Metrics	10
5.5. Defining Incoming Link Metrics	12
5.6. Link Metric Values	12
6. MPRs with Link Metrics	14
6.1. Flooding MPRs	14
6.2. Routing MPRs	16
6.3. Relationship between MPR Sets	19
7. Security Considerations	21
8. Acknowledgements	21
9. Informative References	21
Appendix A. MPR Routing Property	23

1. Introduction

The Optimized Link State Routing Protocol version 1 (OLSRv1) [RFC3626] is a proactive routing protocol for mobile ad hoc networks (MANETs) [RFC2501]. OLSRv1 finds the shortest, defined as minimum number of hops, routes from a router to all possible destinations.

Using only minimum hop routes may result in what are, in practice, inferior routes. Some examples are given in Section 4. Thus, one of the distinguishing features of the Optimized Link State Routing Protocol version 2 (OLSRv2) [RFC7181] is the introduction of the ability to select routes using link metrics other than the number of hops.

During the development of OLSRv2, the working group and authors repeatedly discussed how and why some choices were made in the protocol specification, particularly at the metric integration level. Some of the issues may be non-intuitive, and this document is presented as a record of the considerations and decisions to provide informational discussion about motivation and historic design choices. This document is intended to be useful as a reference if those questions arise again.

Use of the extensible message format [RFC5444] by OLSRv2 has allowed the addition, by OLSRv2, of link metric information to the HELLO messages defined in the MANET Neighborhood Discovery Protocol (NHDP) [RFC6130] as well as inclusion in the Topology Control (TC) messages defined in [RFC7181].

OLSRv2 essentially first determines local link metrics from 1-hop neighbors, these being defined by a process outside OLSRv2, then distributes required link metric values in HELLO messages and TC messages, and then finally forms routes with minimum total link metric. Using a definition of route metric other than number of hops is a natural extension that is commonly used in link state protocols.

A metric-based route selection process for OLSRv2 could have been handled as an extension to OLSRv2. However, were this to have been done, OLSRv2 routers that did not implement this extension would not recognize any link metric information and would attempt to use minimum hop-count routes. This would have meant that, in effect, routers that did implement and routers that did not implement this extension would differ over their valuation of links and routes. This would have led to the fundamental routing problem of "looping". Thus, if metric-based route selection were to have been considered only as an extension to OLSRv2, then routers that did implement and routers that did not implement this extension would not have been

able to interoperate. This would have been a significant limitation of such an extension. Link metrics were therefore included as standard in OLSRv2.

This document discusses the motivation and design rationale behind how link metrics were included in OLSRv2. The principal issues involved when including link metrics in OLSRv2 were:

- o Assigning metrics to links involved considering separate metrics for the two directions of a link, with the receiving router determining the metric from transmitter to receiver. A metric used by OLSRv2 may be either of:
 - * A link metric, the metric of a specific link from an OLSRv2 interface of the transmitting router to an OLSRv2 interface of the receiving router.
 - * A neighbor metric, the minimum of the link metrics between two OLSRv2 routers, in the indicated direction.

These metrics are necessarily the same when these routers each have a single OLSRv2 interface but may differ when either has more. HELLO messages may include both link metrics and neighbor metrics. TC messages include only neighbor metrics.

- o Metrics as used in OLSRv2 are defined to be dimensionless and additive. The assignment of metrics, including their relationship to real parameters such as data rate, loss rate, and delay, and the management of the choice of metric, is outside the scope of [RFC7181], which simply uses these metrics in a consistent manner. Within a single MANET, including all components of a temporarily fragmented MANET, a single choice of link metric is used. By use of a registry of metric types (employing extended types of a single Address Block TLV type), routers can be configured to use only a subset of the available metric types.
- o Node metrics were not included in OLSRv2. Node metrics can be implemented by the addition of the corresponding value to all incoming link metrics by the corresponding router.
- o The separation of the two functions performed by multipoint relays (MPRs) in OLSRv1, optimized flooding and reduced topology advertisement for routing, into separate sets of MPRs in OLSRv2 [RFC7181], denoted "flooding MPRs" and "routing MPRs". Flooding MPRs can be calculated as in [RFC3626], but the use of link metrics in OLSRv2 can improve the MPR selection. Routing MPRs need a metric-aware selection algorithm. The selection of routing MPRs guarantees the use of minimum distance routes using the

chosen metric, while using only symmetric 2-hop neighborhood information from HELLO messages and routing MPR selector information from TC messages.

- o The protocol Information Bases defined in OLSRv2 include required metric values. This has included additions to the protocol Information Bases defined in NHDP [RFC6130] when used by OLSRv2.

2. Terminology

All terms introduced in [RFC5444], including "message" and "TLV" (type-length-value), are to be interpreted as described there.

All terms introduced in [RFC6130], including "MANET interface", "HELLO message", "heard", "link", "symmetric link", "1-hop neighbor", "symmetric 1-hop neighbor", "2-hop neighbor", "symmetric 2-hop neighbor", "symmetric 2-hop neighborhood", and the symbolic constants SYMMETRIC and HEARD, are to be interpreted as described there.

All terms introduced in [RFC7181], including "router", "OLSRv2 interface", "willingness", "multipoint relay (MPR)", "MPR selector", "MPR flooding", and the TLV type LINK_METRIC, are to be interpreted as described there.

3. Applicability

The objective of this document is to retain the design considerations behind how link metrics were included in [RFC7181]. This document does not prescribe any behavior but explains some aspects of the operation of OLSRv2.

4. Motivational Scenarios

The basic situation that suggests the desirability of use of routes other than minimum hop routes is shown in Figure 1.

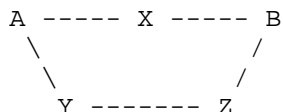


Figure 1

The minimum hop route from A to B is via X. However, if the links A to X and X to B are poor (e.g., have low data rate or are unreliable) but the links A to Y, Y to Z, and Z to B are better (e.g., have reliable high data rate), then the route A to B via Y and Z may be preferred to that via X.

There are other situations where the use of some links should be discouraged, even if the avoidance of them does not show immediately obvious benefits to users. Consider a network with many short-range links and a few long-range links. Use of minimum hop routes will immediately lead to heavy use of the long-range links. This will be particularly undesirable if those links achieve their longer range through reduced data rate or through being less reliable. However, even if the long-range links have the same characteristics as the short-range links, it may be better to reserve usage of the long-range links for when this usage is particularly valuable -- for example, when the use of one long-range link saves several short-range links, rather than the single link saving that is needed for a minimum hop route.

A related case is that of a privileged relay. An example is an aerial router in an otherwise ground-based network. The aerial router may have a link to many, or even all, other routers. That would lead to all routers attempting to send all their traffic (other than to symmetric 1-hop neighbors and some symmetric 2-hop neighbors) via the aerial router. It may, however, be important to reserve that capacity for cases where the aerial router is actually essential, such as if the ground-based portion of the network is not connected.

Link metrics provide a possible solution to these scenarios. For example, in Figure 1, the route A to Y to Z to B could be preferred to A to X to B by making the metrics on the former path 1 and those on the latter path 2. The aerial privileged relay could be used only when necessary by giving its links maximal metric values, with much smaller other metric values or, if the aerial link is to be preferred to N ground links, by giving the ground links metric values of 1 while making the sum of the aerial node uplink and downlink metrics equal to N.

Other cases may involve attempts to avoid areas of congestion, attempts to route around insecure routers, and attempts by routers to discourage being used as relays due to, for example, limited battery power. OLSRv2 does have another mechanism to aid in this: a router's willingness to act as an MPR. However, there are cases where that cannot help but where use of non-minimum hop routes could.

Similarly, note that OLSRv2's optional use of link quality (through its use of [RFC6130]) is not a solution to these problems. Use of link quality as specified in [RFC6130] allows a router to decline to use a link, not only on its own, but on all routers' behalf. It does not, for example, allow the use of a link otherwise determined to be too low quality to be generally useful as part of a route where no better links exist. These mechanisms (link quality and link metrics) solve distinctly different problems.

It should also be noted that the loop-free property of OLSRv2 applies strictly only in the static state. When the network topology is changing and when messages can be lost, it is possible for transient loops to form. However, with update rates appropriate to the rate of topology change, such loops will be sufficiently rare. Changing link metrics is a form of network topology change and should be limited to a rate slower than the message information update rate (defined by the parameters HELLO_INTERVAL, HELLO_MIN_INTERVAL, REFRESH_INTERVAL, TC_INTERVAL, and TC_MIN_INTERVAL).

5. Link Metrics

This section describes the required and selected properties of the link metrics used in OLSRv2, followed by implementation details achieving those properties.

5.1. Link Metric Properties

Link metrics in OLSRv2 are:

- o Dimensionless. While they may, directly or indirectly, correspond to specific physical information (such as delay, loss rate, or data rate), this knowledge is not used by OLSRv2. Instead, generating the metric value is the responsibility of a mechanism external to OLSRv2.
- o Additive, so that the metric of a route is the sum of the metrics of the links forming that route. Note that this requires a metric where a low value of a link metric indicates a "good" link and a high value of a link metric indicates a "bad" link, and the former will be preferred to the latter.
- o Directional, the metric from router A to router B need not be the same as the metric from router B to router A, even when using the same OLSRv2 interfaces. At router A, a link metric from router B to router A is referred to as an incoming link metric, while a link metric from router A to router B is referred to as an outgoing link metric. (These are, of course, reversed at router B.)
- o Specific to a pair of OLSRv2 interfaces, so that if there is more than one link from router A to router B, each has its own link metric in that direction. There is also an overall metric, a "neighbor metric", from router A to router B (its 1-hop neighbor). This is the minimum value of the link metrics from router A to router B, considering symmetric links only; it is undefined if there are no such symmetric links. A neighbor metric from one router to another is always equal to a link metric in the same

direction between OLSRv2 interfaces of those routers. When referring to a specific OLSRv2 interface (for example, in a Link Tuple or a HELLO message sent on that OLSRv2 interface), a link metric always refers to a link on that OLSRv2 interface to or from the indicated 1-hop neighbor OLSRv2 interface, while a neighbor metric may be equal to a link metric to and/or from another OLSRv2 interface.

5.2. Link Metric Types

There are various physical characteristics that may be used to define a link metric. Some examples, which also illustrate some characteristics of metrics that result, are:

- o Delay is a straightforward metric; as it is naturally additive, the delay of a multi-link route is the sum of the delays of the links. This does not directly take into account delays due to routers (such as due to router queues or transition of packets between router interfaces) rather than links, but these delays can be divided among incoming and outgoing links.
- o Probability of loss on a link is, as long as probabilities of loss are small and independent, approximately additive. (A slightly more accurate approach is using a negatively scaled logarithm of the probability of not losing a packet.) If losses are not independent, then this will be pessimistic.
- o Data rates are not additive. They even have the wrong characteristic of being good when high and bad when low; thus, a mapping that inverts the ordering must be applied. Such a mapping can, at best, only produce a metric that is acceptable to treat as additive. Consider, for example, a preference for a route that maximizes the minimum data rate link on the route and then prefers a route with the fewest links of each data rate from the lowest. If links may be of three discrete data rates, "high", "medium", and "low", then this preference can be achieved, on the assumption that no route will have more than 10 links, with metric values of 1, 10, and 100 for the three data rates.

If routes can have more than 10 links, the range of metrics must be increased; this was one reason for a preference for a wide "dynamic range" of link metric values. Depending on the ratios of the numerical values of the three data rates, the same effect may be achieved by using a scaling of an inverse power of the numerical values of the data rates. For example, if the three data rates were 2, 5, and 10 Mbit/s, then a possible mapping would be the fourth power of 10 Mbit/s divided by the data rate, giving metric values of 625, 16, and 1 (good for up to 16 links in a

route). This mapping can be extended to a system with more data rate values, for example, giving a 4 Mbit/s data rate a metric value of about 39. This may lose the capability to produce an absolutely maximal minimum data rate route but will usually produce either that, or something close (and at times maybe better, is a route of three 5 Mbit/s links really better than one of a single 4 Mbit/s link?). Specific metrics will need to define the mapping.

There are also many other possible metrics, including using physical-layer information (such as signal-to-noise ratio and error-control statistics) and information such as packet-queuing statistics.

In a well-designed network, all routers will use the same metric type. It will not produce good routes if, for example, some link metrics are based on data rate and some on path loss (except to the extent that these may be correlated). How to achieve this is an administrative matter, outside the scope of OLSRv2. In fact, even the actual physical meanings of the metrics is outside the scope of OLSRv2. This is because new metrics may be added in the future, for example, as data rates increase, and may be based on new, possibly non-physical, considerations, for example, financial cost. Each such type will have a metric type number. Initially, a single link metric type zero is defined as indicating a dimensionless metric with no predefined physical meaning.

An OLSRv2 router is instructed which single link metric type to use and recognize, without knowing whether it represents delay, probability of loss, data rate, cost, or any other quantity. This recognized link metric type number is a router parameter and subject to change in case of reconfiguration or possibly the use of a protocol (outside the scope of OLSRv2) permitting a process of link metric type agreement between routers.

The use of link metric type numbers also suggests the possibility of use of multiple link metric types and multiple network topologies. This is a possible future extension to OLSRv2. To allow for that future possibility, the sending of more than one metric of different physical types, which should otherwise not be done for reasons of efficiency, is not prohibited, but types other than that configured will be ignored.

The following three sections assume a chosen single link metric type, of unspecified physical nature.

5.3. Directional Link Metrics

OLSRv2 uses only "symmetric" (bidirectional) links, which may carry traffic in either direction. A key decision was whether these links should each be assigned a single metric, used in both directions, or a metric in each direction, noting that:

- o Links can have different characteristics in each direction. Use of directional link metrics recognizes this.
- o In many (possibly most) cases, the two ends of a link will naturally form different views as to what the link metric should be. To use a single link metric requires a coordination between the two that can be avoided if using directional metrics. Note that if using a single metric, it would be essential that the two ends agree as to its value; otherwise, it is possible for looping to occur. This problem does not occur for directional metrics.

Based on these considerations, directional metrics are used in OLSRV2. Each router must thus be responsible for defining the metric in one direction only. This could have been in either direction, i.e., a router is responsible for either incoming or outgoing link metrics, as long as the choice is universal. The former (incoming) case is used in OLSRV2 because, in general, receiving routers have more information available to determine link metrics (for example, received signal strength, interference levels, and error-control coding statistics).

Note that, using directional metrics, if router A defines the metric of the link from router B to router A, then router B must use router A's definition of that metric on that link in that direction. (Router B could, if appropriate, use a bad mismatch between directional metrics as a reason to discontinue use of this link, using the link quality mechanism defined in [RFC6130]; note that this is a distinct mechanism from the use of link metrics.)

5.4. Reporting Link and Neighbor Metrics

Links, and hence link metrics, are reported in HELLO messages. A router must report incoming link metrics in its HELLO messages in order for these link metrics to be available at the other end of the link. This means that, for a symmetric link, both ends of the link will know both of the incoming and outgoing link metrics.

As well as advertising incoming link metrics, HELLO messages also advertise incoming neighbor metrics. These are used for routing MPR selection (see Section 6.2), which requires use of the lowest metric

link between two routers when more than one link exists. This neighbor metric may be using another OLSRV2 interface, and hence, the link metric alone is insufficient.

Metrics are also reported in TC messages. It can be shown that these need to be outgoing metrics:

- o Router A must be responsible for advertising a metric from router A to router B in TC messages. This can be seen by considering a route connecting single OLSRV2 interface routers P to Q to R to S. Router P receives its only information about the link from R to S in the TC messages transmitted by router R, which is an MPR of router S (assuming that only MPR selectors are reported in TC messages). Router S may not even transmit TC messages (if no routers have selected it as an MPR and it has no attached networks to report). So any information about the metric of the link from R to S must also be included in the TC messages sent by router R; hence, router R is responsible for reporting the metric for the link from R to S.
- o In a more general case, where there may be more than one link from R to S, the TC message must, so that minimum metric routes can be constructed (e.g., by router P), report the minimum of these outgoing link metrics, i.e., the outgoing neighbor metric from R to S.

In this example, router P also receives information about the existence of a link between Q and R in the HELLO messages sent by router Q. Without the use of metrics, this link could be used by OLSRV2 for 2-hop routing to router R, using just HELLO messages sent by router Q. For this property (which accelerates local route formation) to be retained (from OLSRV1), router P must receive the metric from Q to R in HELLO messages sent by router Q. This indicates that router Q must be responsible for reporting the metric for the outgoing link from Q to R. This is in addition to the incoming link metric information that a HELLO message must report. Again, in general, this must be the outgoing neighbor metric, rather than the outgoing link metric.

In addition, Section 6.1 offers an additional reason for reporting outgoing neighbor metrics in HELLO messages, without which metrics can properly affect only routing, not flooding.

Note that there is no need to report an outgoing link metric in a HELLO message. The corresponding 1-hop neighbor knows that value; it specified it. Furthermore, for 2-hop neighborhood use, neighbor metrics are required (as these will, in general, not use the same OLSRV2 interface).

5.5. Defining Incoming Link Metrics

When a router reports a 1-hop neighbor in a HELLO message, it may do so for the first time with link status HEARD. As the router is responsible for defining and reporting incoming link metrics, it must evaluate that metric and attach that link metric to the appropriate address (which will have link status HEARD) in the next HELLO message reporting that address on that OLSRv2 interface. There will, at this time, be no outgoing link metric available to report, but a router must be able to immediately decide on an incoming link metric once it has heard a 1-hop neighbor on an OLSRv2 interface for the first time.

This is because, when receiving a HELLO message from this router, the 1-hop neighbor seeing its own address listed with link status HEARD will (unless the separate link quality mechanism indicates otherwise) immediately consider that link to be SYMMETRIC, advertise it with that link status in future HELLO messages, and use it (for MPR selection and data traffic forwarding).

It may, depending on the physical nature of the link metric, be too early for an ideal decision as to that metric; however, a choice must be made. The metric value may later be refined based on further observation of HELLO messages, other message transmissions between the routers, or other observations of the environment. It will probably be best to over-estimate the metric if initially uncertain as to its value, to discourage, rather than over-encourage, its use. If no information other than the receipt of the HELLO message is available, then a conservative maximum link metric value, denoted MAXIMUM_METRIC in [RFC7181], should be used.

5.6. Link Metric Values

Link metric values are recorded in LINK_METRIC TLVs, defined in [RFC7181], using a compressed (lossy) representation that occupies 12 bits. The use of 12 bits is convenient because, when combined with 4 flag bits of additional information, described below, this results in a 2-octet value field. However, the use of 12 bits, and thus the availability of 4 flag bits, was a consequence of a design to use a modified exponent/mantissa form with the following characteristics:

- o The values represented are to be positive integers starting 1, 2, ...
- o The maximum value represented should be close to, but less than 2^{24} (^ denotes exponentiation in this section). This is so that with a route limited to no more than 255 hops, the maximum route metric is less than 2^{32} , i.e., can be stored in 32 bits. (The link metric value can be stored in 24 bits.)

A representation that is modified from an exponent/mantissa form with e bits of exponent and m bits of mantissa and that has the first of these properties is one that starts at 1, then is incremented by 1 up to 2^m , then has a further 2^m increments by 2, then a further 2^m increments by 4, and so on for 2^e sets of increments. This means that the represented value is never in error by more than a half (if rounding) or one (if truncating) part in 2^m , usually less.

The position in the increment sequence, from 0 to 2^m-1 , is considered as a form of mantissa and denoted a . The increment sequence number, from 0 to 2^e-1 , is considered as a form of exponent and denoted b .

The value represented by (b,a) can then be shown to be equal to $(2^{m+a+1})2^b-2^m$. To verify this, note that:

- o With fixed b , the difference between two values with consecutive values of a is 2^b , as expected.
- o The value represented by $(b,2^m-1)$ is $(2^{m+2^m})2^b-2^m$. The value represented by $(b+1,0)$ is $(2^{m+1})(2^{b+1})-2^m$. The difference between these two values is 2^{b+1} , as expected.

The maximum represented value has $b = 2^e-1$ and $a = 2^m-1$ and is $(2^{m+2^m})(2^{(2^e-1)})-2^m = 2^{(2^e+m)}-2^m$. This is slightly less than $2^{(2^e+m)}$. The required 24-bit limit can be achieved if $2^e+m = 24$. Of the possible (e,m) pairs that satisfy this equation, the pair $e = 4$, $m = 8$ was selected as most appropriate and is that used by OLSRv2. It uses the previously indicated $e+m = 12$ bits. An algorithm for converting from a 24-bit value v to a 12-bit pair (b,a) is given in Section 6.2 of [RFC7181].

As noted above, the 12-bit representation then shares two octets with 4 flag bits. Putting the flag bits first, it is then natural to put the exponent bits in the last four bits of the first octet and to put the mantissa bits in the second octet. The 12 consecutive bits, using network byte order (most significant octet first), then represent $256b+a$. Note that the ordering of these 12-bit representation values is the same as the ordering of the 24-bit metric values. In other words, two 12-bit metrics fields can be compared for equality/ordering as if they were unsigned integers.

The four flag bits each represent one kind of metric, defined by its direction (incoming or outgoing) and whether the metric is a link metric or a neighbor metric. As indicated by the flag bits set, a metric value may be of any combination of these four kinds of metric.

6. MPRs with Link Metrics

MPRs are used for two purposes in OLSRV2. In both cases, it is MPR selectors that are actually used, MPR selectors being determined from MPRs advertised in HELLO messages.

- o Optimized Flooding. This uses the MPR selector status of symmetric 1-hop neighbor routers from which messages are received in order to determine if these messages are to be forwarded. MPR selector status is recorded in the Neighbor Set (defined in [RFC6130] and extended in [RFC7181]) and determined from received HELLO messages.
- o Routing. Non-local link information is based on information recorded in this router's Topology Information Base. That information is based on received TC messages. The neighbor information in these TC messages consists of addresses of the originating router's advertised (1-hop) neighbors, as recorded in that router's Neighbor Set (defined in [RFC6130] and extended in [RFC7181]). These advertised neighbors include all of the MPR selectors of the originating router.

Metrics interact with these two uses of MPRs differently, as described in the following two sections. This leads to the requirement for two separate sets of MPRs for these two uses when using metrics. The relationship between these two sets of MPRs is considered in Section 6.3.

6.1. Flooding MPRs

The essential detail of the "flooding MPR" selection specification is that a router must select a set of MPRs such that a message transmitted by a router and retransmitted by all its flooding MPRs will reach all of the selecting router's symmetric 2-hop neighbors.

Flooding MPR selection can ignore metrics and produce a solution that meets the required specification. However, that does not mean that metrics cannot be usefully considered in selecting flooding MPRs. Consider the network in Figure 2, where numbers are metrics of links in the direction away from router A, towards router D.

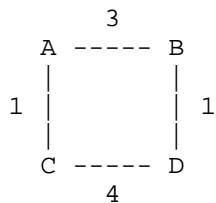


Figure 2

Which is the better flooding MPR selection by router A: B or C? If the metric represents probability of message loss, then clearly choosing B maximizes the probability of a message sent by A reaching D. This is despite C having a lower metric in its connection to A than B does. (Similar arguments about a preference for B can be made if, for example, the metric represents data rate or delay rather than probability of loss.)

However, neither should only the second hop be considered. If this example is modified to that in Figure 3, where the numbers still are metrics of links in the direction away from router A, towards router D, then it is possible that, when A is selecting flooding MPRs, selecting C is preferable to selecting B.

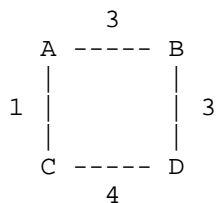


Figure 3

If the metrics represent scaled values of delay or the probability of loss, then selecting C is clearly better. This indicates that the sum of metrics is an appropriate measure to use to choose between B and C.

However, this is a particularly simple example. Usually, it is not a simple choice between two routers as a flooding MPR, each only adding one router coverage. When considering which router to next add as a flooding MPR, a more general process should incorporate the metric to that router and the metric from that router to each symmetric 2-hop neighbor as well as the number of newly covered symmetric 2-hop neighbors. Other factors may also be included.

The required specification for flooding MPR selection is in Section 18.4 (also using Section 18.3) of [RFC7181], which may use the example MPR selection algorithm in Appendix B of [RFC7181]. However, note that (as in [RFC3626]) each router can make its own independent choice of flooding MPRs, and flooding MPR selection algorithm, and still interoperate.

Also note that the references above to the direction of the metrics is correct: for flooding, directional metrics outward from a router are appropriate, i.e., metrics in the direction of the flooding. This is an additional reason for including outward metrics in HELLO messages, as otherwise a metric-aware MPR selection for flooding is not possible. The second-hop metrics are outgoing neighbor metrics because the OLSRV2 interface used for a second-hop transmission may not be the same as that used for the first-hop reception.

6.2. Routing MPRs

The essential detail of the "routing MPR" selection specification is that a router must, per OLSRV2 interface, select a set of MPRs such that there is a 2-hop route from each symmetric 2-hop neighbor of the selecting router to the selecting router, with the intermediate router on each such route being a routing MPR of the selecting router.

It is sufficient, when using an additive link metric rather than a hop count, to require that these routing MPRs provide not just a 2-hop route but a minimum distance 2-hop route. In addition, a router is a symmetric 2-hop neighbor even if it is a symmetric 1-hop neighbor, as long as there is a 2-hop route from it that is shorter than the 1-hop link from it. (The property that no routes go through routers with willingness WILL_NEVER is retained. Examples below assume that all routers are equally willing, with none having willingness WILL_NEVER.)

For example, consider the network in Figure 4. Numbers are metrics of links in the direction towards router A, away from router D. Router A must pick router B as a routing MPR, whereas for minimum hop count routing, it could alternatively pick router C. Note that the use of incoming neighbor metrics in this case follows the same reasoning as for the directionality of metrics in TC messages, as described in Section 5.4.

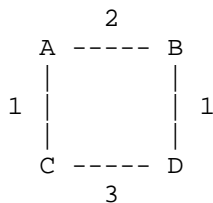


Figure 4

In Figure 5, where numbers are metrics of links in the direction towards router A and away from router C, router A must pick router B as a routing MPR, but for minimum hop count routing, it would not need to pick any MPRs.

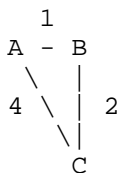


Figure 5

In Figure 6, where numbers are metrics of links in the direction towards router A and away from routers D and E, router A must pick both routers B and C as routing MPRs, but for minimum hop count routing, it could pick either.

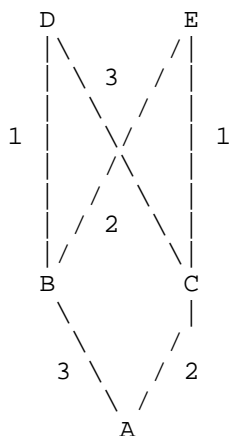


Figure 6

It is shown in Appendix A that selecting routing MPRs according to this definition and advertising only such links (plus knowledge of local links from HELLO messages) will result in selection of lowest total metric routes, even if all links (advertised or not) are considered in the definition of a shortest route.

However, the definition noted above as sufficient for routing MPR selection is not necessary. For example, consider the network in Figure 7, where numbers are metrics of links in the direction towards router A, away from other routers; the metrics from B to C and C to B are both assumed to be 2.

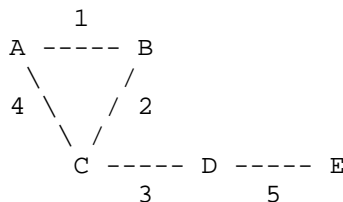


Figure 7

Using the above definition, A must pick both B and C as routing MPRs, in order to cover the symmetric 2-hop neighbors C and D, respectively. (C is a symmetric 2-hop neighbor because the route length via B is shorter than the 1-hop link.)

However, A only needs to pick B as a routing MPR, because the only reason to pick C as a routing MPR would be so that C can advertise the link to A for routing -- to be used by, for example, E. However, A knows that no other router should use the link C to A in a shortest route because routing via B is shorter. So, if there is no need to advertise the link from C to A, then there is no reason for A to select C as a routing MPR.

This process of "thinning out" the routing MPR selection uses only local information from HELLO messages. Using any minimum distance algorithm, the router identifies shortest routes, whether one, two, or more hops, from all routers in its symmetric 2-hop neighborhood. It then selects as MPRs all symmetric 1-hop neighbors that are the last router (before the selecting router itself) on any such route. Where there is more than one shortest distance route from a router, only one such route is required. Alternative routes may be selected so as to minimize the number of last routers -- this is the equivalent to the selection of a minimal set of MPRs in the non-metric case.

Note that this only removes routing MPRs whose selection can be directly seen to be unnecessary. Consequently, if (as is shown in Appendix A) the first approach creates minimum distance routes, then so does this process.

The examples in Figures 5 and 6 show that use of link metrics may require a router to select more routing MPRs than when not using metrics and even require a router to select routing MPRs when, without metrics, it would not need any routing MPRs. This may result in more, and larger, messages being generated and forwarded more often. Thus, the use of link metrics is not without cost, even excluding the cost of link metric signaling.

These examples consider only single OLSRv2 interface routers. However, if routers have more than one OLSRv2 interface, then the process is unchanged; other than that, if there is more than one known metric between two routers (on different OLSRv2 interfaces), then, considering symmetric links only (as only these are used for routing) the smallest link metric, i.e., the neighbor metric, is used. There is no need to calculate routing MPRs per OLSRv2 interface. That requirement results from the consideration of flooding and the need to avoid certain "race" conditions, which are not relevant to routing, only to flooding.

The required specification for routing MPR selection is in Section 18.5 (also using Section 18.3) of [RFC7181], which may use the example MPR selection algorithm in Appendix B of [RFC7181]. However, note that (as in [RFC3626]) each router can make its own independent choice of routing MPRs, and routing MPR selection algorithm, and still interoperate.

6.3. Relationship between MPR Sets

It would be convenient if the two sets of flooding and routing MPRs were the same. This can be the case if all metrics are equal, but in general, for "good" sets of MPRs, they are not. (A reasonable definition of this is that there is no common minimal set of MPRs.) If metrics are asymmetrically valued (the two sets of MPRs use opposite direction metrics) or routers have multiple OLSRv2 interfaces (where routing MPRs can ignore this but flooding MPRs cannot), this is particularly unlikely. However, even using a symmetrically valued metric with a single OLSRv2 interface on each router, the ideal sets need not be equal, nor is one always a subset of the other. To show this, consider these examples, where all lettered routers are assumed equally willing to be MPRs, and numbers are bidirectional metrics for links.

In Figure 8, A does not require any flooding MPRs. However, A must select B as a routing MPR.

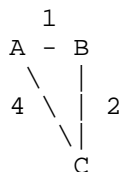


Figure 8

In Figure 9, A must select C and D as routing MPRs. However, A's minimal set of flooding MPRs is just B. In this example, the set of routing MPRs serves as a set of flooding MPRs, but a non-minimal one (although one that might be better, depending on the relative importance of number of MPRs and flooding link metrics).

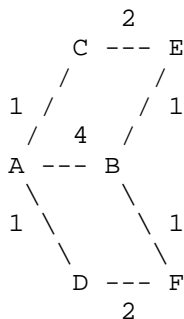


Figure 9

However, this is not always the case. In Figure 10, A's set of routing MPRs must contain B but need not contain C. A's set of flooding MPRs need not contain B but must contain C. (In this case, flooding with A selecting B rather than C as a flooding MPR will reach D but in three hops rather than the minimum two that MPR flooding guarantees.)

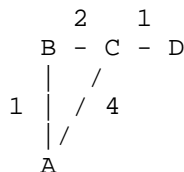


Figure 10

7. Security Considerations

An attacker can have an adverse impact on an OLSRv2 network by creating apparently valid messages that contain incorrect link metrics. This could take the form of influencing the choice of routes or, in some cases, producing routing loops. This is a more subtle, and likely to be less effective, attack than other forms of invalid message injection. These can add and remove other and more basic forms of network information, such as the existence of some routers and links.

As such, no significantly new security issues arose from the inclusion of metrics in OLSRv2. Defenses to the injection of invalid link metrics are the same as to other forms of invalid message injection, as discussed in the Security Considerations section of [RFC7181].

There are possible uses for link metrics in the creation of security countermeasures to prefer the use of links that have better security properties, including better availability, to those with poorer security properties. This, however, is beyond the scope of both this document and [RFC7181].

8. Acknowledgements

The authors would like to gratefully acknowledge the following people (listed alphabetically) for intense technical discussions, early reviews, and comments on the documents and its components: Brian Adamson (NRL), Alan Cullen (BAE Systems), Justin Dean (NRL), Ulrich Herberg (Fujitsu), Charles Perkins (Huawei), Stan Ratliff (Cisco), and Henning Rogge (FGAN).

Finally, the authors would like to express their gratitude to (listed alphabetically) Benoit Claise, Adrian Farrel, Stephen Farrell, and Suresh Krishnan for their reviews and comments on the later draft versions of this document.

9. Informative References

- [RFC2501] Corson, S. and J. Macker, "Mobile Ad hoc Networking (MANET): Routing Protocol Performance Issues and Evaluation Considerations", RFC 2501, January 1999.
- [RFC3626] Clausen, T. and P. Jacquet, "Optimized Link State Routing Protocol (OLSR)", RFC 3626, October 2003.

- [RFC5444] Clausen, T., Dearlove, C., Dean, J., and C. Adjih, "Generalized Mobile Ad Hoc Network (MANET) Packet/Message Format", RFC 5444, February 2009.
- [RFC6130] Clausen, T., Dearlove, C., and J. Dean, "Mobile Ad Hoc Network (MANET) Neighborhood Discovery Protocol (NHDP)", RFC 6130, April 2011.
- [RFC7181] Clausen, T., Dearlove, C., Jacquet, P., and U. Herberg, "The Optimized Link State Routing Protocol Version 2", RFC 7181, April 2014.

Appendix A. MPR Routing Property

In order for routers to find and use shortest routes in a network while using the minimum reduced topology supported by OLSRV2 (that a router only advertises its MPR selectors in TC messages), routing MPR selection must result in the property that there are shortest routes with all intermediate routers being routing MPRs.

This appendix uses the following terminology and assumptions:

- o The network is a graph of nodes connected by arcs, where nodes correspond to routers with willingness not equal to `WILL_NEVER` (except possibly at the ends of routes). An arc corresponds to the set of symmetric links connecting those routers; the OLSRV2 interfaces used by those links are not relevant.
- o Each arc has a metric in each direction, being the minimum of the corresponding link metrics in that direction, i.e., the corresponding neighbor metric. This metric must be positive.
- o A sequence of arcs joining two nodes is referred to as a path.
- o Node A is an MPR of node B if corresponding router A is a routing MPR of router B.

The required property (of using shortest routes with reduced topology) is equivalent to the following property: for any pair of distinct nodes X and Z, there is a shortest path from X to Z, $X - Y_1 - Y_2 - \dots - Y_m - Z$ such that Y_1 is an MPR of Y_2 , ..., Y_m is an MPR of Z. Call such a path a routable path, and call this property the routable path property.

The required definition for a node X selecting MPRs is that for each distinct node Z from which there is a two-arc path, there is a shorter, or equally short, path that is either $Z - Y - X$ where Y is an MPR of X or is the one-arc path $Z - X$. Note that the existence of locally known, shorter paths that have more than two arcs, which can be used to reduce the numbers of MPRs, is not considered here. (Such reductions are only when the remaining MPRs can be seen to retain all necessary shortest paths and therefore retain the required property.)

Although this appendix is concerned with paths with minimum total metric, not number of arcs (hop count), it proceeds by induction on the number of arcs in a path. Although it considers minimum metric routes with a bounded number of arcs, it then allows that number of arcs to increase so that overall minimum metric paths, regardless of the number of arcs, are considered.

Specifically, the routable path property is a corollary of the property that for all positive integers n and all distinct nodes X and Z , if there is any path from X to Z of n arcs or fewer, then there is a shortest path, from among those of n arcs or fewer, that is a routable path. This may be called the n -arc routable path property.

The n -arc routable path property is trivial for $n = 1$ and directly follows from the definition of the MPRs of Z for $n = 2$.

Proceeding by induction, assuming the n -arc routable path property is true for $n = k$, consider the case that $n = k+1$.

Suppose that $X - V_1 - V_2 - \dots - V_k - Z$ is a shortest $k+1$ arc path from X to Z . We construct a path that has no more than $k+1$ arcs, has the same or shorter length (hence has the same, shortest, length considering only paths of up to $k+1$ arcs, by assumption), and is a routable path.

First, consider whether V_k is an MPR of Z . If it is not, then consider the two-arc path $V_{k-1} - V_k - Z$. This can be replaced either by a one-arc path $V_{k-1} - Z$ or by a two-arc path $V_{k-1} - W_k - Z$, where W_k is an MPR of Z , such that the metric from V_{k-1} to Z by the replacement path is no longer. In the former case (replacement one-arc path), this now produces a path of length k , and the previous inductive step may be applied. In the latter case, we have replaced V_k by W_k , where W_k is an MPR of Z . Thus, we need only consider the case that V_k is an MPR of Z .

We now apply the previous inductive step to the path $X - V_1 - \dots - V_{k-1} - V_k$, replacing it by an equal length path $X - W_1 - \dots - W_{m-1} - V_k$, where $m \leq k$, where this path is a routable path. Then, because V_k is an MPR of Z , the path $X - W_1 - \dots - W_{m-1} - V_k - Z$ is a routable path and demonstrates the n -arc routable path property for $n = k+1$.

This thus shows that for any distinct nodes X and Z , there is a routable path using the MPR-reduced topology from X to Z , i.e., that OLSRv2 finds minimum length paths (minimum total metric routes).

Authors' Addresses

Christopher Dearlove
BAE Systems Advanced Technology Centre
West Hanningfield Road
Great Baddow, Chelmsford
United Kingdom

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

Thomas Heide Clausen
LIX, Ecole Polytechnique
91128 Palaiseau Cedex
France

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

Philippe Jacquet
Alcatel-Lucent Bell Labs

Phone: +33 6 7337 1880
EMail: philippe.jacquet@alcatel-lucent.com