

Internet Engineering Task Force (IETF)
Request for Comments: 7779
Category: Experimental
ISSN: 2070-1721

H. Rogge
Fraunhofer FKIE
E. Baccelli
INRIA
April 2016

Directional Airtime Metric Based on Packet Sequence Numbers for
Optimized Link State Routing Version 2 (OLSRv2)

Abstract

This document specifies a Directional Airtime (DAT) link metric for usage in Optimized Link State Routing version 2 (OLSRv2).

Status of This Memo

This document is not an Internet Standards Track specification; it is published for examination, experimental implementation, and evaluation.

This document defines an Experimental Protocol for the Internet community. 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/rfc7779>.

Copyright Notice

Copyright (c) 2016 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.

Table of Contents

- 1. Introduction 2
- 2. Terminology 3
- 3. Applicability Statement 4
- 4. Directional Airtime Metric Rationale 5
- 5. Metric Functioning and Overview 6
- 6. Protocol Constants 7
- 7. Protocol Parameters 8
 - 7.1. Recommended Values 8
- 8. Data Structures 8
 - 8.1. Initial Values 9
- 9. Packets and Messages 10
 - 9.1. Definitions 10
 - 9.2. Requirements for Using DAT Metric in OLSRv2 Implementations 10
 - 9.3. Link-Loss Data Gathering 11
 - 9.4. HELLO Message Processing 12
- 10. Timer Event Handling 12
 - 10.1. Packet Timeout Processing 12
 - 10.2. Metric Update 13
- 11. Security Considerations 14
- 12. References 14
 - 12.1. Normative References 14
 - 12.2. Informative References 15
- Appendix A. Future Work 17
- Appendix B. OLSR.org Metric History 17
- Appendix C. Link-Speed Stabilization 18
- Appendix D. Packet-Loss Hysteresis 19
- Appendix E. Example DAT Values 19
- Acknowledgements 20
- Authors' Addresses 21

1. Introduction

One of the major shortcomings of Optimized Link State Routing (OLSR) [RFC3626] is the lack of a granular link-cost metric between OLSR routers. Operational experience with OLSR networks gathered since its publication has revealed that wireless networks links can have highly variable and heterogeneous properties. This makes a hop-count metric insufficient for effective OLSR routing.

Based on this experience, OLSRv2 [RFC7181] integrates the concept of link metrics directly into the core specification of the routing protocol. The OLSRv2 routing metric is an external process, and it can be any kind of dimensionless additive cost function that reports to the OLSRv2 protocol.

Since 2004, the OLSR.org [OLSR.org] implementation of OLSR has included an Estimated Transmission Count (ETX) metric [MOBICOM04] as a proprietary extension. While this metric is not perfect, it proved to be sufficient for a long time for Community Mesh Networks (see Appendix B). But the increasing maximum data rate of IEEE 802.11 made the ETX metric less efficient than in the past, which is one reason to move to a different metric.

This document describes a Directional Airtime routing metric for OLSRv2, a successor of the OLSR.org ETX-derived routing metric for OLSR. It takes both the loss rate and the link speed into account to provide a more accurate picture of the links within the network.

This specification allows OLSRv2 deployments with a metric defined by the IETF Mobile Ad Hoc Networks (MANET) working group. It enables easier interoperability testing between implementations and targets to deliver a useful baseline to compare with, for experiments with this metric as well as other metrics. Appendix A contains a few possible steps to improve the Directional Airtime metric. Future experiments should also determine whether the DAT metric can be useful for other IETF protocols, both inside and outside of the MANET working group. This could lead to either moving this document to the Standards Track or replacing it with an improved document.

2. Terminology

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

The terminology introduced in [RFC5444], [RFC7181], and [RFC6130], including the terms "packet", "message" and "TLV", are to be interpreted as described therein.

Additionally, this document uses the following terminology and notational conventions:

DAT - Directional Airtime (metric). The link metric specified in this document, which is a directional variant of ETT. It does not take reverse path loss into account.

QUEUE - A first in, first out queue of integers.

QUEUE[TAIL] - The most recent element in the queue.

add(QUEUE, value) - Adds a new element to the TAIL of the queue.

remove(QUEUE) - Removes the HEAD element of the queue.

sum(QUEUE) - An operation that returns the sum of all elements in a QUEUE.

diff_seqno(new, old) - An operation that returns the positive distance between two elements of the circular sequence number space defined in Section 5.1 of [RFC5444]. Its value is either (new - old) if this result is positive, or else its value is (new - old + 65536).

MAX(a, b) - The maximum of a and b.

MIN(a, b) - The minimum of a and b.

UNDEFINED - A value not in the normal value range of a variable.

airtime - The time a transmitted packet blocks the link layer, e.g., a wireless link.

ETX - Expected Transmission Count. A link metric proportional to the number of transmissions to successfully send an IP packet over a link.

ETT - Estimated Travel Time. A link metric proportional to the amount of airtime needed to successfully transmit an IP packet over a link, not considering Layer 2 overhead created by preamble, backoff time, and queuing.

3. Applicability Statement

The Directional Airtime metric was designed and tested (see [COMNET15]) in wireless IEEE 802.11 OLSRv2 networks [RFC7181]. These networks employ link-layer retransmission to increase the delivery probability. A dynamic rate selection algorithm selects the unicast data rate independently for each neighbor.

As specified in OLSRv2, the metric calculates only the incoming link cost. It neither calculates the outgoing metric, nor decides the link status (heard, symmetric, lost).

The metric works both for nodes that can send/receive [RFC5444] packet sequence numbers and those that do not have this capability. In the absence of such sequence numbers, the metric calculates the packet loss based on HELLO message [RFC6130] timeouts.

The metric must learn about the unicast data rate towards each one-hop neighbor from an external process, either by configuration or by an external measurement process. This measurement could be done via gathering cross-layer data from the operating system, via an external

daemon like Dynamic Link Exchange Protocol [DLEP], or via indirect Layer 3 measurements like packet-pair (see [MOBICOM04]).

The metric uses [RFC5444] multicast control traffic to determine the link packet loss. The administrator should take care that link-layer multicast transmission do not have a higher reception probability than the slowest unicast transmission without retransmission. For example, with 802.11g, it might be necessary to increase the data-rate of the multicast transmissions, e.g., set the multicast data-rate to 6 Mbit/s.

The metric can only handle a certain range of packet loss and unicast data-rate. The maximum packet loss that can be encoded into the metric is a loss of 7 of 8 packets (87.5%), without link-layer retransmissions. The unicast data-rate that can be encoded by this metric can be between 1 kbit/s and 2 Gbit/s. This metric has been designed for data-rates of 1 Mbit/s and hundreds of Mbit/s.

4. Directional Airtime Metric Rationale

The Directional Airtime metric has been inspired by the publications on the ETX [MOBICOM03] and ETT [MOBICOM04] metric, but differs from both of these in several ways.

Instead of measuring the combined loss probability of a bidirectional transmission of a packet over a link in both directions, the Directional Airtime metric measures the incoming loss rate and integrates the incoming link speed into the metric cost. There are multiple reasons for this decision:

- o OLSRv2 [RFC7181] defines the link metric as directional costs between routers.
- o Not all link-layer implementations use acknowledgement mechanisms. Most link-layer implementations that do use them use less airtime and a more robust modulation for the acknowledgement than the data transmission, which makes it more likely for the data transmission to be disrupted compared to the acknowledgement.
- o Incoming packet loss and link speed can be measured locally, while symmetric link loss would need an additional signaling TLV in the HELLO [RFC6130] and would delay metric calculation by up to one HELLO interval.

The Directional Airtime metric does not integrate the packet size into the link cost. Doing so is not feasible in most link-state routing protocol implementations. The routing decision of most operation systems does not take packet size into account.

Multiplying all link costs of a topology with the size of a data-plane packet would never change the Dijkstra result in any way.

The queue-based packet-loss estimator specified in this document has been tested extensively in the OLSR.org ETX implementation; see Appendix B. The output is the average of the packet loss over a configured time period.

The metric normally measures the loss of a link by tracking the incoming [RFC5444] packet sequence numbers. Without these packet sequence numbers, the metric does calculate the loss of the link based on the received and lost [RFC6130] HELLO messages. It uses the incoming HELLO interval time (or if not present, the validity time) to decide when a HELLO is lost.

When a neighbor router resets, its packet sequence number might jump to a random value. The metric tries to detect jumps in the packet sequence number and removes them from the data set because the previously gathered link-loss data should still be valid (see Section 9.3). The link-loss data is only removed from memory when a link times out completely and its Link Set Tuple is removed from the database.

5. Metric Functioning and Overview

The Directional Airtime metric is calculated for each Link Set entry, as defined in [RFC6130], Section 7.1.

The metric processes two kinds of data into the metric value, namely packet-loss rate and link speed. The link speed is taken from an external process not defined in this document. The current packet-loss rate is defined in this document by keeping track of packet reception and packet-loss events. It could also be calculated by an external process with a compatible output.

Multiple incoming packet-loss/reception events must be combined into a loss rate to get a smooth metric. Experiments with exponential weighted moving average (EWMA) lead to a highly fluctuating or a slow converging metric (or both). To get a smoother and more controllable metric result, this metric uses two fixed-length queues to measure and average the incoming packet events, one queue for received packets and one for the estimated number of packets sent by the other side of the link.

Because the rate of incoming packets is not uniform over time, the queue contains a number of counters, each representing a fixed time interval. Incoming packet-loss and packet-reception events are accumulated in the current queue element until a timer adds a new

empty counter to both queues and removes the oldest counter from both.

In addition to the packet loss stored in the queue, this metric uses a timer to detect a total link loss. For every [RFC6130] HELLO interval in which the metric received no packet from a neighbor, it scales the number of received packets in the queue based on the total time interval the queue represents compared to the total time of the lost HELLO intervals.

The average packet-loss ratio is calculated as the sum of the 'total packets' counters divided by the sum of the 'packets received' counters. This value is then divided through the current link speed and then scaled into the range of metrics allowed for OLSRv2.

The metric value is then used as `L_in_metric` of the Link Set (as defined in Section 8.1. of [RFC7181]).

While this document does not add new [RFC5444] elements to HELLO [RFC6130] or TC messages [RFC7181], it works best when both the `INTERVAL_TIME` message TLV is present in the HELLO messages and when each [RFC5444] packet contains an interface-specific sequence number. It also adds a number of new data entries to be stored for each [RFC6130] link.

6. Protocol Constants

This specification defines the following constants, which define the range of metric values that can be encoded by the DAT metric (see Table 1). They cannot be changed without making the metric outputs incomparable and should only be changed for a MANET with a very slow or a very fast link layer. See Appendix E for example metric values.

`DAT_MAXIMUM_LOSS` - Fraction of the loss rate used in this routing metric. Loss rate will be between $0/\text{DAT_MAXIMUM_LOSS}$ and $(\text{DAT_MAXIMUM_LOSS}-1)/\text{DAT_MAXIMUM_LOSS}$.

`DAT_MINIMUM_BITRATE` - Minimal bitrate in Bit/s used by this routing metric.

Name	Value
<code>DAT_MAXIMUM_LOSS</code>	8
<code>DAT_MINIMUM_BITRATE</code>	1000

Table 1: DAT Protocol Constants

7. Protocol Parameters

This specification defines the following parameters for this routing metric. These parameters are:

DAT_MEMORY_LENGTH - Queue length for averaging packet loss. All received and lost packets within the queue length are used to calculate the cost of the link.

DAT_REFRESH_INTERVAL - Interval in seconds between two metric recalculations as described in Section 10.2. This value SHOULD be smaller than a typical HELLO interval. The interval can be a fraction of a second.

DAT_HELLO_TIMEOUT_FACTOR - Multiplier relative to the HELLO_INTERVAL (see Section 5.3.1 of [RFC6130]) after which the DAT metric considers a HELLO as lost.

DAT_SEQNO_RESTART_DETECTION - Threshold in the number of missing packets (based on received packet sequence numbers) at which point the router considers the neighbor has restarted. This parameter is only used for loss estimation based on packet sequence numbers. This number MUST be larger than DAT_MAXIMUM_LOSS.

7.1. Recommended Values

The proposed values of the protocol parameters are for Community Mesh Networks, which mostly use routers that are not mobile. Using this metric for mobile networks might require shorter DAT_REFRESH_INTERVAL and/or DAT_MEMORY_LENGTH.

DAT_MEMORY_LENGTH := 64

DAT_REFRESH_INTERVAL := 1

DAT_HELLO_TIMEOUT_FACTOR := 1.2

DAT_SEQNO_RESTART_DETECTION := 256

8. Data Structures

This specification extends the Link Set of the Interface Information Base, as defined in Section 7.1 of [RFC6130], by the adding the following elements to each Link Tuple:

L_DAT_received - A QUEUE with DAT_MEMORY_LENGTH integer elements. Each entry contains the number of successfully received packets within an interval of DAT_REFRESH_INTERVAL.

L_DAT_total - A QUEUE with DAT_MEMORY_LENGTH integer elements. Each entry contains the estimated number of packets transmitted by the neighbor, based on the received packet sequence numbers within an interval of DAT_REFRESH_INTERVAL.

L_DAT_packet_time - The time when the next [RFC5444] packet should have arrived.

L_DAT_hello_interval - The interval between two HELLO messages of the links neighbor as signaled by the INTERVAL_TIME TLV [RFC5497] of NHDP messages [RFC6130].

L_DAT_lost_packet_intervals - The estimated number of HELLO intervals from this neighbor from which the metric has not received a single packet.

L_DAT_rx_bitrate - The current bitrate of incoming unicast traffic for this neighbor.

L_DAT_last_pkt_seqno - The last received packet sequence number received from this link.

Methods to obtain the value of L_DAT_rx_bitrate are out of the scope of this specification. Such methods may include static configuration via a configuration file or dynamic measurement through mechanisms described in a separate specification (e.g., [DLEP]). Any Link Tuple with L_status = HEARD or L_status = SYMMETRIC MUST have a specified value of L_DAT_rx_bitrate if it is to be used by this routing metric.

The incoming bitrate value should be stabilized by a hysteresis filter to improve the stability of this metric. See Appendix D for an example.

This specification updates the L_in_metric field of the Link Set of the Interface Information Base, as defined in Section 8.1. of [RFC7181]).

8.1. Initial Values

When generating a new tuple in the Link Set, as defined in item 3 of Section 12.5 of [RFC6130], the values of the elements specified in Section 8 are set as follows:

- o L_DAT_received := 0, ..., 0. The queue always has DAT_MEMORY_LENGTH elements.
- o L_DAT_total := 0, ..., 0. The queue always has DAT_MEMORY_LENGTH elements.

- o L_DAT_packet_time := EXPIRED (no earlier [RFC5444] packet received).
- o L_DAT_hello_interval := UNDEFINED (no earlier NHDP HELLO received).
- o L_DAT_lost_packet_intervals := 0 (no HELLO interval without packets).
- o L_DAT_last_pkt_seqno := UNDEFINED (no earlier [RFC5444] packet with sequence number received).

9. Packets and Messages

This section describes the necessary changes of [RFC7181] implementations with DAT metric for the processing and modification of the incoming and outgoing [RFC5444] data.

9.1. Definitions

For the purpose of this section, note the following definitions:

- o "pkt_seqno" is defined as the [RFC5444] packet sequence number of the received packet.
- o "interval_time" is the time encoded in the INTERVAL_TIME message TLV of a received HELLO message [RFC6130].
- o "validity_time" is the time encoded in the VALIDITY_TIME message TLV of a received HELLO message [RFC6130].

9.2. Requirements for Using DAT Metric in OLSRv2 Implementations

An implementation of OLSRv2 using the metric specified by this document SHOULD include the following parts into its [RFC5444] output:

- o An INTERVAL_TIME message TLV in each HELLO message, as defined in [RFC6130], Section 4.3.2.
- o An interface-specific packet sequence number as defined in [RFC5444], Section 5.1 that is incremented by 1 for each outgoing [RFC5444] packet on the interface.

An implementation of OLSRv2 using the metric specified by this document that inserts packet sequence numbers in some, but not all, outgoing [RFC5444] packets will make this metric ignore all packets without the sequence number. Putting the INTERVAL_TIME TLV into

some, but not all, HELLO messages will make the timeout-based loss detection slower. This will only matter in the absence of packet sequence numbers.

9.3. Link-Loss Data Gathering

For each incoming [RFC5444] packet, additional processing SHOULD be carried out after the packet messages have been processed as specified in [RFC6130] and [RFC7181] as described in this section.

[RFC5444] packets without packet sequence numbers MUST NOT be processed in the way described in this section.

The router updates the Link Set Tuple corresponding to the originator of the packet:

1. If L_DAT_last_pkt_seqno = UNDEFINED, then:
 - * L_DAT_received[TAIL] := 1.
 - * L_DAT_total[TAIL] := 1.
2. Otherwise:
 - * L_DAT_received[TAIL] := L_DAT_received[TAIL] + 1.
 - * diff := diff_seqno(pkt_seqno, L_DAT_last_pkt_seqno).
 - * If diff > DAT_SEQNO_RESTART_DETECTION, then:
 - diff := 1.
 - * L_DAT_total[TAIL] := L_DAT_total[TAIL] + diff.
3. L_DAT_last_pkt_seqno := pkt_seqno.
4. If L_DAT_hello_interval != UNDEFINED, then:
 - * L_DAT_packet_time := current time + (L_DAT_hello_interval * DAT_HELLO_TIMEOUT_FACTOR).
5. L_DAT_lost_packet_intervals := 0.

9.4. HELLO Message Processing

For each incoming HELLO Message, after it has been processed as defined in Section 12 of [RFC6130], the Link Set Tuple corresponding to the incoming HELLO message MUST be updated.

1. If the HELLO message contains an INTERVAL_TIME message TLV, then:

L_DAT_hello_interval := interval_time.

2. Otherwise:

L_DAT_hello_interval := validity_time.

3. If L_DAT_last_pkt_seqno = UNDEFINED, then:

* L_DAT_received[TAIL] := L_DAT_received[TAIL] + 1.

* L_DAT_total[TAIL] := L_DAT_total[TAIL] + 1.

* L_DAT_packet_time := current time + (L_DAT_hello_interval * DAT_HELLO_TIMEOUT_FACTOR).

10. Timer Event Handling

In addition to changes in the [RFC5444] processing/generation code, the DAT metric also uses two timer events.

10.1. Packet Timeout Processing

When L_DAT_packet_time has timed out, the following step MUST be done:

1. If L_DAT_last_pkt_seqno = UNDEFINED, then:

L_DAT_total[TAIL] := L_DAT_total[TAIL] + 1.

2. Otherwise:

L_DAT_lost_packet_intervals := L_DAT_lost_packet_intervals + 1.

3. L_DAT_packet_time := L_DAT_packet_time + L_DAT_hello_interval.

10.2. Metric Update

Once every DAT_REFRESH_INTERVAL, all L_in_metric values in all Link Set entries MUST be recalculated:

1. `sum_received := sum(L_DAT_received)`.
2. `sum_total := sum(L_DAT_total)`.
3. If `L_DAT_hello_interval != UNDEFINED` and `L_DAT_lost_packet_intervals > 0`, then:
 - * `lost_time_proportion := L_DAT_hello_interval * L_DAT_lost_packet_intervals / DAT_MEMORY_LENGTH`.
 - * `sum_received := sum_received * MAX(0, 1 - lost_time_proportion)`;
4. If `sum_received < 1`, then:
 - `L_in_metric := MAXIMUM_METRIC`, as defined in [RFC7181], Section 5.6.1.
5. Otherwise:
 - * `loss := MIN(sum_total / sum_received, DAT_MAXIMUM_LOSS)`.
 - * `bitrate := MAX(L_DAT_rx_bitrate, DAT_MINIMUM_BITRATE)`.
 - * `L_in_metric := (2^24 / DAT_MAXIMUM_LOSS) * loss / (bitrate / DAT_MINIMUM_BITRATE)`.
6. `remove(L_DAT_total)`
7. `add(L_DAT_total, 0)`
8. `remove(L_DAT_received)`
9. `add(L_DAT_received, 0)`

The calculated L_in_metric value should be stabilized by a hysteresis function. See Appendix D for an example.

11. Security Considerations

Artificial manipulation of metrics values can drastically alter network performance. In particular, advertising a higher `L_in_metric` value may decrease the amount of incoming traffic, while advertising lower `L_in_metric` may increase the amount of incoming traffic.

For example, by artificially attracting mesh routes and then dropping the incoming traffic, an attacker may achieve a Denial of Service (DoS) against other mesh nodes. Similarly, an attacker may achieve Man-in-the-Middle (MITM) attacks or traffic analysis by concentrating traffic being routed over a node the attacker controls (and end-to-end encryption is not used or somehow broken). Protection mechanisms against such MITM or DoS attacks are nevertheless out of scope of this document.

Security threats also include potential attacks on the integrity of the control traffic passively monitored by DAT to measure link quality. For example, an attacker might inject packets pretending to be somebody else and using incorrect sequence numbers. This attack can be prevented by the true originator of the [RFC5444] packets by adding an ICV Packet TLV and TIMESTAMP Packet TLV [RFC7182] to each packet. This allows the receiver to drop all incoming packets that have a forged packet source, both packets generated by the attacker, or replayed packets. However, the security mechanism described in [RFC7183] does not protect the sequence number used by the DAT metric because it only signs the [RFC5444] messages, not the [RFC5444] packet header (which contains the [RFC5444] packet sequence number).

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC5444] Clausen, T., Dearlove, C., Dean, J., and C. Adjih, "Generalized Mobile Ad Hoc Network (MANET) Packet/Message Format", RFC 5444, DOI 10.17487/RFC5444, February 2009, <<http://www.rfc-editor.org/info/rfc5444>>.
- [RFC5497] Clausen, T. and C. Dearlove, "Representing Multi-Value Time in Mobile Ad Hoc Networks (MANETs)", RFC 5497, DOI 10.17487/RFC5497, March 2009, <<http://www.rfc-editor.org/info/rfc5497>>.

- [RFC6130] Clausen, T., Dearlove, C., and J. Dean, "Mobile Ad Hoc Network (MANET) Neighborhood Discovery Protocol (NHDP)", RFC 6130, DOI 10.17487/RFC6130, April 2011, <<http://www.rfc-editor.org/info/rfc6130>>.
- [RFC7181] Clausen, T., Dearlove, C., Jacquet, P., and U. Herberg, "The Optimized Link State Routing Protocol Version 2", RFC 7181, DOI 10.17487/RFC7181, April 2014, <<http://www.rfc-editor.org/info/rfc7181>>.

12.2. Informative References

- [RFC3626] Clausen, T., Ed. and P. Jacquet, Ed., "Optimized Link State Routing Protocol (OLSR)", RFC 3626, DOI 10.17487/RFC3626, October 2003, <<http://www.rfc-editor.org/info/rfc3626>>.
- [RFC7182] Herberg, U., Clausen, T., and C. Dearlove, "Integrity Check Value and Timestamp TLV Definitions for Mobile Ad Hoc Networks (MANETs)", RFC 7182, DOI 10.17487/RFC7182, April 2014, <<http://www.rfc-editor.org/info/rfc7182>>.
- [RFC7183] Herberg, U., Dearlove, C., and T. Clausen, "Integrity Protection for the Neighborhood Discovery Protocol (NHDP) and Optimized Link State Routing Protocol Version 2 (OLSRv2)", RFC 7183, DOI 10.17487/RFC7183, April 2014, <<http://www.rfc-editor.org/info/rfc7183>>.
- [COMNET15] Barz, C., Fuchs, C., Kirchhoff, J., Niewiejska, J., and H. Rogge, "OLSRv2 for Community Networks: Using Directional Airtime Metric with external radios", Elsevier Computer Networks 2015, DOI 10.1016/j.comnet.2015.09.022, September 2015, <<http://dx.doi.org/10.1016/j.comnet.2015.09.022>>.
- [CONFINE] "Community Networks Testbed for the Future Internet (CONFINE)", <<http://www.confine-project.eu>>.
- [DLEP] Ratliff, S., Berry, B., Jury, S., Satterwhite, D., and R. Taylor, "Dynamic Link Exchange Protocol (DLEP)", Work in Progress, draft-ietf-manet-dlep-22, April 2016.
- [BATMAN] Neumann, A., Aichele, C., Lindner, M., and S. Wunderlich, "Better Approach To Mobile Ad-hoc Networking (B.A.T.M.A.N.)", Work in Progress, draft-wunderlich-openmesh-manet-routing-00, April 2008.

[MOBICOM03]

De Couto, D., Aguayo, D., Bicket, J., and R. Morris, "A High-Throughput Path Metric for Multi-Hop Wireless Routing", Proceedings of the MOBICOM Conference, DOI 10.1145/938985.939000, 2003.

[MOBICOM04]

Draves, R., Padhye, J., and B. Zill, "Routing in Multi-Radio, Multi-Hop Wireless Mesh Networks", Proceedings of the MOBICOM Conference, DOI 10.1145/1023720.1023732, 2004.

[OLSR.org] "OLSR.org Wiki", <<http://www.olsr.org/>>.

[FREIFUNK] "Freifunk Wireless Community Networks", <<http://www.freifunk.net>>.

[FUNKFEUER]

"Austria Wireless Community Network", <<http://www.funkfeuer.at>>.

Appendix A. Future Work

As the DAT metric proved to work reasonably well for non- or slow-moving ad hoc networks [COMNET15], it should be considered a solid first step on a way to better MANET metrics. There are multiple parts of the DAT metric that need to be reviewed again in the context of real world deployments and can be subject to later improvements.

The easiest part of the DAT metric to change and test would be the timings parameters. A 1-minute interval for packet-loss statistics might be a good compromise for some MANETs, but could easily be too large or too small for others. More data is needed to verify or improve the current parameter selection.

The DAT metric considers only the multicast [RFC5444] packet loss for estimating the link, but it would be good to integrate the unicast data loss into the loss estimation. This information could be provided directly from the link layer. This could increase the accuracy of the loss rate estimation in scenarios where the assumptions regarding the ratio of multicast vs. unicast loss do not hold.

The packet-loss averaging algorithm could also be improved. While the DAT metric provides a stable sliding time interval to average the incoming packet loss and does not give the recent input too much influence, first experiments suggest that the algorithm tends to be less agile in detecting major changes of link quality. This makes it less suited for mobile networks. A more agile algorithm is needed for detecting major changes while filtering out random fluctuations regarding frame loss. However, the current "queue of counters" algorithm suggested for DAT outperforms the binary queue algorithm and the exponential aging algorithms used for the ETX metric in the OLSR [RFC3626] codebase of OLSR.org.

Appendix B. OLSR.org Metric History

The Funkfeuer [FUNKFEUER] and Freifunk networks [FREIFUNK] are based on OLSR [RFC3626] or B.A.T.M.A.N. [BATMAN] wireless community networks with hundreds of routers in permanent operation. The Vienna Funkfeuer network in Austria, for instance, consists of 400 routers covering the whole city of Vienna and beyond, spanning roughly 40 km in diameter. It has been supplying its users with Internet access since 2003. A particularity of the Vienna Funkfeuer network is that it manages to provide Internet access through a city-wide, large-scale Wi-Fi MANET, with just a single Internet uplink.

Operational experience of the OLSR project [OLSR.org] with these networks has revealed that the use of hop-count as a routing metric leads to unsatisfactory network performance. Experiments with the ETX metric [MOBICOM03] were therefore undertaken in parallel in the Berlin Freifunk network as well as in the Vienna Funkfeuer network in 2004, and found satisfactory, i.e., sufficiently easy to implement and providing sufficiently good performance. This metric has now been in operational use in these networks for several years.

The ETX metric of a link is the estimated number of transmissions required to successfully send a packet (each packet equal to or smaller than MTU) over that link, until a link-layer acknowledgement is received. The ETX metric is additive, i.e., the ETX metric of a path is the sum of the ETX metrics for each link on this path.

While the ETX metric delivers a reasonable performance, it does not handle networks with heterogeneous links that have different bitrates well. When using the ETX metric, since every wireless link is characterized only by its packet-loss ratio, long-ranged links with low bitrate (with low loss ratios) are preferred over short-ranged links with high bitrate (with higher but reasonable loss ratios). Such conditions, when they occur, can degrade the performance of a network considerably, by not taking advantage of higher capacity links.

Because of this, the OLSR.org project has implemented the Directional Airtime metric for OLSRv2, which has been inspired by the Estimated Travel Time (ETT) metric [MOBICOM04]. This metric uses a unidirectional packet loss, but also takes the bitrate into account to create a more accurate description of the relative costs or capabilities of OLSRv2 links.

Appendix C. Link-Speed Stabilization

The DAT metric specifies how to generate a reasonably stable packet-loss rate value based on incoming packet reception/loss events, but the source of the link speed used in this document is considered an external process.

In the presence of a Layer 2 technology with variable link speed, it is likely that the raw link speed will be fluctuating too fast to be useful for the DAT metric.

The amount of stabilization necessary for the link speed depends on the implementation of the MAC layer, especially the rate-control algorithm.

Experiments with the Linux 802.11 Wi-Fi stack have shown that a simple Median filter over a series of raw link-speed measurements can smooth the calculated value without introducing intermediate link-speed values one would obtain by using averaging or an exponential weighted moving average.

Appendix D. Packet-Loss Hysteresis

While the DAT metric uses a sliding window to compute a reasonably stable frame loss, the implementation might choose to integrate an additional hysteresis to prevent undesirable oscillations between two values (i.e., metric flapping).

In Section 10.2, DAT calculates a fractional loss rate. The fraction of "loss := sum_total / sum_received" may result in minor fluctuations in the advertised L_in_metric due to minimal changes in sum_total or sum_received, which can cause undesirable protocol churn.

A hysteresis function applied to the fraction could reduce the amount of changes in the loss rate and help to further stabilize the metric output.

Appendix E. Example DAT Values

The DAT metric value can be expressed in terms of link speed (bit/s) or used airtime (s). When using the default protocol constants (see Section 6), DAT encodes link speeds between 119 bit/s and 2 Gbit/s.

Table 2 contains a few examples for metric values and their meaning as a link speed:

Metric	bit/s
MINIMUM_METRIC (1)	2 Gbit/s
MAXIMUM_METRIC (16776960)	119 bit/s
2000	1 Mbit/s

Table 2: DAT Link Cost Examples

A path metric value could also be expressed as a link speed, but this would be less intuitive. An easier way to transform a path metric value into a textual representation is to divide it by the hop count of the path and express the path cost as the average link speed together with the hop count (see Table 3).

Metric	hops	average bit/s
4	2	1 Gbit/s
4000000	6	3 kbit/s

Table 3: DAT Link Cost Examples

Acknowledgements

The authors would like to acknowledge the network administrators from Freifunk Berlin [FREIFUNK] and Funkfeuer Vienna [FUNKFEUER] for endless hours of testing and suggestions to improve the quality of the original ETX metric for the OLSR.org routing daemon.

This effort/activity is supported by the European Community Framework Program 7 within the Future Internet Research and Experimentation Initiative (FIRE), Community Networks Testbed for the Future Internet ([CONFINE]), contract FP7-288535.

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): Teco Boot (Infinity Networks), Juliusz Chroboczek (PPS, University of Paris 7), Thomas Clausen, Christopher Dearlove (BAE Systems Advanced Technology Centre), Ulrich Herberg (Fujitsu Laboratories of America), Markus Kittenberger (Funkfeuer Vienna), Joseph Macker (Naval Research Laboratory), Fabian Nack (Freie Universitaet Berlin), and Stan Ratliff (Cisco Systems).

Authors' Addresses

Henning Rogge
Fraunhofer FKIE

Email: henning.rogge@fkie.fraunhofer.de
URI: <http://www.fkie.fraunhofer.de>

Emmanuel Baccelli
INRIA

Email: Emmanuel.Baccelli@inria.fr
URI: <http://www.emmanuelbaccelli.org/>