

Networking Working Group	O.G. Gnawali
Internet-Draft	P.L. Levis
Intended status: Standards Track	Stanford University
Expires: November 05, 2011	May 04, 2011

The Minimum Rank Objective Function with Hysteresis
draft-ietf-roll-minrank-hysteresis-of-03

Abstract

The Routing Protocol for Low Power and Lossy Networks (RPL) uses objective functions to construct routes that optimize or constrain the routes it selects and uses. This specification describes the Minimum Rank Objective Function with Hysteresis (MRHOF), an objective function that selects routes that minimize a metric, while using hysteresis to reduce churn in response to small metric changes. MRHOF works with metrics that are additive along a route, and the metric it uses is determined by the metrics RPL Destination Information Object (DIO) messages advertise.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on November 05, 2011.

Copyright Notice

Copyright (c) 2011 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. [Terminology](#)
- *3. [The Minimum Rank Objective Function with Hysteresis](#)
 - *3.1. [Computing the Path cost](#)
 - *3.2. [Parent Selection](#)
 - *3.3. [Computing Rank](#)
 - *3.4. [Advertising the Path Cost](#)
 - *3.5. [Working Without Metric Containers](#)
- *4. [Using MRHOF for Metric Maximization](#)
- *5. [MRHOF Variables and Parameters](#)
- *6. [Acknowledgements](#)
- *7. [IANA Considerations](#)
- *8. [Security Considerations](#)
- *9. [References](#)
 - *9.1. [Normative References](#)
 - *9.2. [Informative References](#)
- *[Authors' Addresses](#)

1. Introduction

An objective function specifies how RPL [\[I-D.ietf-roll-rpl\]](#) selects paths. Objective functions can choose paths based on routing metrics or constraints. For example, if an RPL instance uses an objective function that minimizes hop-count, RPL will select paths with minimum hop count. The nodes running RPL might use a number of metrics to describe a link or a node [\[I-D.ietf-roll-routing-metrics\]](#) and make it available for route selection. These metrics are advertised in RPL Destination Information Object (DIO) messages using a Metric Container suboption. An objective function can use these metrics to choose routes. To decouple the details of an individual metric or objective function from forwarding and routing, RPL describes routes through a value called Rank. Rank, roughly speaking, corresponds to the distance

associated with a route. An objective function is responsible for computing a node's advertised Rank value based on the Rank of its potential parents, metrics, and other network properties. This specification describes MRHOF, an objective function for RPL. MRHOF uses hysteresis while selecting the path with the smallest metric value. The metric that MRHOF uses is determined by the metrics in the DIO Metric Container. For example, the use of MRHOF with the latency metric allows RPL to find stable minimum-latency paths from the nodes to a root in the DAG instance. The use of MRHOF with the ETX metric allows RPL to find the stable minimum-ETX paths from the nodes to a root in the DAG instance. MRHOF can only be used with an additive metric that must be minimized on the paths selected for routing.

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 [RFC 2119 \[RFC2119\]](#).

This terminology used in this document is consistent with the terminologies described in [\[I-D.ietf-roll-terminology\]](#), [\[I-D.ietf-roll-rpl\]](#), and [\[I-D.ietf-roll-routing-metrics\]](#).

This document introduces two terms:

Selected metric: The metric chosen by the network operator to use for path selection. This metric can be any additive metric listed in [\[I-D.ietf-roll-routing-metrics\]](#).

Path cost: Path cost quantifies a property of an end-to-end path. Path cost is obtained by summing up the selected metric of the links or nodes along the path. Path cost can be used by RPL to compare different paths.

Worst parent: The node in the parent set with the largest path cost.

3. The Minimum Rank Objective Function with Hysteresis

The Minimum Rank Objective Function with Hysteresis, MRHOF, is designed to find the paths with the smallest path cost while preventing excessive churn in the network. It does so by finding the minimum cost path and switching to that path only if it is shorter (in terms of path cost) than the current path by at least a given threshold. MRHOF may be used with any additive metric listed in [\[I-D.ietf-roll-routing-metrics\]](#) as long the routing objective is to minimize the given routing metric.

3.1. Computing the Path cost

Nodes compute the path cost for each candidate neighbor reachable on an interface. The Path cost represents the cost of the path, in terms of

the selected metric, from a node to the root of the DODAG through the neighbor.

Root nodes (Grounded or Floating) set the variable `cur_min_path_cost` to `MIN_PATH_COST`.

A non-root node computes the path cost for a path to the root through each candidate neighbor by adding these two components:

1. If the selected metric is a link metric, the selected metric for the link to a candidate neighbor. If the selected metric is a node metric, the selected metric for the node.
2. The value of the selected metric in the metric container in the DIO sent by that neighbor.

A node SHOULD compute the path cost for the path through each candidate neighbor reachable through an interface. If a node cannot compute the path cost for the path through a candidate neighbor, the node MUST NOT select the candidate neighbor as its preferred parent, with one exception. If the node does not have metrics to compute the path cost through any of the candidate neighbors, it MUST join one of the candidate neighbors as a leaf node.

If the selected metric is a link metric and the metric of the link to a neighbor is not available, the path cost for the path through that neighbor SHOULD be set to `MAX_PATH_COST`. This cost value will prevent this path from being considered for path selection.

If the selected metric is a node metric, and the metric is not available, the path cost through all the neighbors SHOULD be set to `MAX_PATH_COST`.

The path cost corresponding to a neighbor SHOULD be re-computed each time:

1. The selected metric of the link to the candidate neighbor is updated.
2. If the selected metric is a node metric and the metric is updated.
3. A node receives a new metric advertisement from the candidate neighbor.

This computation MAY also be performed periodically. Too much delay in updating the path cost after the metric is updated or a new metric advertisement is received can lead to stale Rank or parent set.

3.2. Parent Selection

After computing the path cost for all the candidate neighbors reachable through an interface for the current DODAG iteration, a node selects the preferred parent. This process is called parent selection. Parent Selection SHOULD be performed each time:

1. The path cost for an existing candidate neighbor, including the preferred parent, changes. This condition can be checked immediately after the path cost is computed.
2. A new candidate neighbor is inserted into the neighbor table.

The parent selection MAY be deferred until a later time. Deferring the parent selection can delay the use of better paths available in the network.

A node MUST select a candidate neighbor as its preferred parent if the path cost corresponding to that neighbor is smaller than the path cost corresponding to the rest of the neighbors, except as indicated below:

1. If the smallest path cost for paths through the candidate neighbors is smaller than `cur_min_path_cost` by less than `PARENT_SWITCH_THRESHOLD`, the node MAY continue to use the current preferred parent.
2. If there are multiple paths with the smallest path cost and the smallest path cost is smaller than `cur_min_path_cost` by at least `PARENT_SWITCH_THRESHOLD`, a node MAY use a different objective function to select the preferred parent among the candidate neighbors on the path with the minimum cost.
3. A node MAY declare itself as a Floating root, and hence no preferred parent, depending on the configuration.
4. If the selected metric for a link is greater than `MAX_LINK_METRIC`, the node SHOULD exclude that link from consideration for parent selection.
5. If `cur_min_path_cost` is greater than `MAX_PATH_COST`, the node MAY declare itself as a Floating root.
6. If the configuration disallows a node to be a Floating root and no neighbors are discovered, the node does not have a preferred parent, and MUST set `cur_min_path_cost` to `MAX_PATH_COST`.

Except in the cases above, the candidate neighbor on the path with the smallest path cost is the preferred parent. A node MAY include a total of `PARENT_SET_SIZE` candidate neighbors in the parent set. The cost of path through the nodes in the parent set is smaller than or equal to the cost of the paths through any of the nodes that are not in the parent set. If the cost of the path through the preferred parent and the worst parent is too large, a node MAY keep a smaller parent set.

3.3. Computing Rank

The DAG roots set their rank to `MIN_PATH_COST` for the selected metric.

Once a non-root node selects its parent set, it can use the following table to convert the path cost of the worst parent (written as Cost in the table) to its rank:

Node/link Metric	Rank
Node Energy	255 - Cost
Hop-Count	Cost
Latency	Cost/65536
Link Quality Level	Cost
ETX	Cost

Conversion of metric to rank.

Nodes MUST support at least one of the above metrics. Nodes SHOULD support the ETX metric.

Node rank is undefined for these node/link metrics: Node state and attributes, throughput, and link color. If the rank is undefined, the node MUST join one of the neighbors as a leaf node.

3.4. Advertising the Path Cost

Once the preferred parent is selected, the node sets its `cur_min_path_cost` variable to the path cost corresponding to the preferred parent. Thus, `cur_min_path_cost` is the cost of the minimum cost path from the node to the root. The value of the `cur_min_path_cost` is carried in the metric container corresponding to the selected metric when DIO messages are sent.

3.5. Working Without Metric Containers

In the absence of metric container, MRHOF uses ETX as its metric. It locally computes the ETX of links to its neighbors and adds this value to their advertised Rank to compute the associated Rank of routes. Once parent selection and rank computation is performed using the ETX metric, the node advertises a Rank equal to the ETX cost and SHOULD NOT include a metric container in its DIO messages.

4. Using MRHOF for Metric Maximization

MRHOF cannot be directly used for parent selection using metrics which require finding paths with maximum value of the selected metric, such as path reliability. It is possible to convert such a metric maximization problem to a metric minimization problem and use MRHOF provided:

- *There is a fixed and well-known maximum metric value corresponding to the best path. This is the path cost for the DAG root. Example, the best link reliability has a value of 1.

*Metrics are all positive. Example, link reliability is always positive.

For metrics meeting the above conditions, the problem of maximizing the metric value is equivalent to minimizing the negative of the metric value. MRHOF is not required to work with these metrics.

5. MRHOF Variables and Parameters

MRHOF uses the following variable:

*cur_min_path_cost: The cost of the path from a node through its preferred parent to the root computed at the last parent selection.

MRHOF uses the following parameters:

*MAX_LINK_METRIC: Maximum allowed value for the selected link metric for each link on the path.

*MAX_PATH_COST: Maximum allowed value for the path metric of a selected path.

*MIN_PATH_COST: The minimum allowed value for the path metric of the selected path.

*PARENT_SWITCH_THRESHOLD: The difference between metric of the path through the preferred parent and the minimum-metric path in order to trigger the selection of a new preferred parent.

*PARENT_SET_SIZE: The number of candidate parents, including the preferred parent, in the parent set.

The parameter values are assigned depending on the selected metric. The best values for these parameters should be experimentally determined. The working group has long experience routing with the ETX metric. Based on those experiences, these ETX parameters are known to work in many settings:

*MAX_LINK_METRIC: 10. Disallow links with greater than 10 expected transmission count on the selected path.

*MAX_PATH_COST: 100. Disallow paths with greater than 100 expected transmission count.

*MIN_PATH_COST: 0. At root, the expected transmission count is 0.

*PARENT_SWITCH_THRESHOLD: 1.5. Switch to a new path only if it is expected to require at least 1.5 fewer transmission than the current path.

*PARENT_SET_SIZE: 3. If the preferred parent is not available, two candidate parents are still available without triggering a new round of route discovery.

6. Acknowledgements

Thanks to Antonio Grilo, Nicolas Tsiftes, Matteo Paris, JP Vasseur, and Phoebus Chen for their comments.

7. IANA Considerations

This specification requires an allocated OCP. A value of 1 is requested.

8. Security Considerations

Security considerations to be developed in accordance to the output of the WG.

9. References

9.1. Normative References

[RFC2119]	Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels" , BCP 14, RFC 2119, March 1997.
-----------	---

9.2. Informative References

[I-D.ietf-roll-rpl]	Winter, T, Thubert, P and R Team, " RPL: IPv6 Routing Protocol for Low power and Lossy Networks ", Internet-Draft draft-ietf-roll-rpl-05, December 2009.
[I-D.ietf-roll-routing-metrics]	Vasseur, J and D Networks, " Routing Metrics used for Path Calculation in Low Power and Lossy Networks ", Internet-Draft draft-ietf-roll-routing-metrics-01, October 2009.
[I-D.ietf-roll-terminology]	Vasseur, J, " Terminology in Low power And Lossy Networks ", Internet-Draft draft-ietf-roll-terminology-01, May 2009.

Authors' Addresses

Omprakash Gnawali
Gnawali Stanford University S255 Clark Center, 318
Campus Drive Stanford, CA 94305 USA Phone: +1 650 725 6086 Email:
gnawali@cs.stanford.edu

Philip Levis
Levis Stanford University 358 Gates Hall, Stanford
University Stanford, CA 94305 USA Email: pal@cs.stanford.edu