Networking Working Group	O.G. Gnawali
Internet-Draft	P.L. Levis
Intended status: Standards Track	Stanford University
Expires: September 01, 2011	February 28, 2011

The Minimum Rank Objective Function with Hysteresis
draft-ietf-roll-minrank-hysteresis-of-01

Abstract

Hysteresis delays the effect of changes in link metric on parent selection. Such delay makes the topology stable despite jitters in link metrics. The Routing Protocol for Low Power and Lossy Networks (RPL) allows the use of objective functions to construct routes that optimize or constrain a routing metric on the paths. This specification describes the Minimum Rank Objective Function with Hysteresis (MRHOF), an objective function that minimizes the node rank in terms of a given metric, while using hysteresis to prevent excessive rank churn. The use of MRHOF with RPL results in nodes selecting stable paths that minimize the given routing metric to the roots of a Directed Acyclic Graph (DAG).

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

An objective function allows RPL [I-D.ietf-roll-rpl] to select paths that are best in terms of a given routing metric or select paths that meet certain constraints in terms of the routing metric. RPL achieves this goal by selecting the parent among the alternate parents as dictated by that objective function. 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. A metric can be used by different objective functions to optimize or constrain the metric in different ways.

This specification describes MRHOF, an objective function for RPL. MRHOF uses hysteresis while selecting the path with the smallest metric value. The path with the minimum cost has different property depending on the metric used for path selection. 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 be used only with additive metric that must be minimized on the paths selected for routing. Although MRHOF can be used with a number of metrics, this draft is based on experiences with the ETX metric.

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 composed using the selected metric of the links along the path. Path cost can be used by RPL to compare different paths.

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 switching to the minimum cost path only if the path cost of the current path is larger than the path cost of the minimum cost path by 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. MRHOF cannot be used if the routing objective is to maximize the metric.

3.1. Computing the Path cost

Nodes compute the path cost for each candidate neighbor reachable on all the interfaces. 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. The selected metric for the link to a candidate neighbor.

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 all the interfaces. 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 SHOULD join one of the candidate neighbors as a leaf node.
If the selected 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.
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. A node receives a new metric advertisement from the candidate neighbor.

This computation MAY also be performed periodically. Deferring the path cost computation for too long after new metric advertisements or updates to the selected link metric results in nodes making parent selection decision based on stale link and path information.

3.2. Parent Selection

After computing the path cost for all the candidate neighbors reachable through all the interfaces 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 or stopping the use of worse paths than what is 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 that 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 candidates which are first hop 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.

The preferred parent is the only node in the parent set at a given time. Any candidate neighbor may become the preferred parent as indicated above.

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 preferred parent, it can use the following table to covert its path cost to the DAG root through its preferred 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.

Node rank is undefined for these node/link metrics: Node state and attributes, throughput, and link color.

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 whenever DIO messages are sent.

4. 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 to trigger new preferred parent selection.

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.

5. Acknowledgements

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

6. IANA Considerations

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

7. Security Considerations

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

8. References

8.1. Normative References

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

8.2. 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