

ROLL
Internet-Draft
Intended status: Standards Track
Expires: September 12, 2019

R. Koutsiamanis, Ed.
G. Papadopoulos
IMT Atlantique
E. Ingles Sanchez
University of Murcia
C. Ji
Alexander TEI of Thessaloniki
D. Dujovne
Universidad Diego Portales
N. Montavont
IMT Atlantique
March 11, 2019

Traffic-aware Objective Function
draft-koutsiamanis-roll-traffic-aware-of-00

Abstract

This document proposes a remaining throughput metric for parent and DODAG selection. This metric represents the amount of remaining traffic handling capacity that the node has. This document also proposes an Objective Function (OF) which uses the proposed metric for parent and DODAG selection to balance the amount of traffic between nodes and DODAGs.

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 <https://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 September 12, 2019.

Copyright Notice

Copyright (c) 2019 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 (<https://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.	DODAG construction in RPL	3
4.	Load distribution problem in RPL	4
4.1.	Parent selection problem	5
4.2.	DODAG selection problem	6
5.	TAOF description	8
6.	DIO Metric Container Type extension	10
6.1.	DAG Metric Container fields	12
7.	Enrollment	12
8.	Security Considerations	13
9.	IANA Considerations	13
10.	Informative references	13
	Authors' Addresses	14

[1.](#) Introduction

RPL [[RFC6550](#)] is an IPv6 Routing protocol for LLNs. It uses Objective Functions (OF) to construct the Destination Oriented Directed Acyclic Graph (DODAG) containing the nodes of the network. The existing OFs defined are OF Zero (OF0) [[RFC6552](#)] and Minimum Rank with Hysteresis OF (MRHOF) [[RFC6719](#)]. These OFs specify how nodes in a DODAG select their preferred parent using different metrics.

The metrics can be separated into two different types, link metrics (e.g. ETX) and node metrics (e.g. energy). Experimental results [[I-D.gasem-roll-rpl-load-balancing](#)] conclude that using the current OFs leads to an unbalanced network within which some nodes are overloaded. Here, a node is overloaded in the sense that it forwards many more packets than it otherwise would if the network were balanced. This problem has consequences for the lifetime of the network because overloaded nodes drain quicker than others, a problem which becomes even more significant when the overloaded nodes are near the DODAG root [[I-D.gasem-roll-rpl-load-balancing](#)].

Similarly, one DODAG might be overloaded in the same sense compared to another DODAG, and this will lead to the same consequences for the whole DODAG as for a specific node.

This problem is still an open issue. This draft proposes a new way of parent and DODAG selection as an attempt towards a solution. This draft proposes a new OF that considers the remaining throughput as a representation of the remaining traffic handling capacity each node possesses and which uses this information to balance the amount of traffic between nodes and DODAGs.

In brief, each node tracks its remaining throughput and appends this information as a DAG Metric Container option to DIO messages it sends. When the DIO message is received by child nodes or potential child nodes, the remaining throughput information is stored and used to influence the result when RPL parent or DODAG selection is performed.

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](#)].

3. DODAG construction in RPL

RPL uses OFs to construct a DODAG. OFs define the way the nodes select their preferred parent and DODAG and how they compute the new rank. A node's rank is always larger than its parent's rank because the calculation of rank is based on an increment to the parent's rank. This increment differs for each OF but all OFs include the MinHopRankIncrease, which is the minimum increase in rank between a node and a node's parent and a step. Different OFs use different metrics or constraints to select the preferred parent and DODAG and to define the step, depending on application requirements. Nodes obtain these values from DODAG Information Object (DIO) control messages sent by their neighbor nodes.

The construction of a DODAG starts when the root node sends DIO messages to its neighbors. After receiving the DIO, these neighbor nodes select the root as their preferred parent if they wish to join the DODAG. To announce that they joined the DODAG as its child node, they send a Destination Advertisement Object (DAO) to their preferred parent - the DODAG root. After joining the DODAG, these nodes send their own DIO messages with the new computed rank to their neighbors. This procedure repeats for every node which joins the DODAG.

4. Load distribution problem in RPL

According to the experiments conducted using existing OFs RPL faces a load distribution problem in large LLNs. With RPL using existing OFs, such as MRHOF, an unbalanced network is formed with some nodes overloaded and other nodes at rest. This problem is severe for network performance because overloaded nodes will use up their available energy faster than other nodes. This is exacerbated for nodes near the root (within 1 hop distance) or nodes which are the only parent candidate for other nodes. Additionally, when the overloaded node shuts down, a big part of the network will become disconnected and will have to be transferred to another parent or DODAG. There is a high probability that the children nodes will also select the same new node as their parent or the same DODAG, leading to another overloaded node/DODAG. Also, when a node has selected its parent, it will change only when the parent node is not reachable (due to battery depletion or packet losses).

The existing OFs usually use a single metric to compare parent candidates, for example, as described in [\[RFC6719\]](#) the default metric used in MRHOF is ETX [\[RFC6551\]](#), which represents the number of transmissions a node expects to make to a destination to successfully deliver one packet. The result from using a single metric is that nodes prefer to select the same node as their parent, which according to [\[I-D.qasem-roll-rpl-load-balancing\]](#) leads to an unbalanced network with overloaded nodes (node load is indicated by a node's child count). But the child count does not accurately indicate the load because among the child nodes some may have higher traffic load and others may have lower.

The network traffic can be quantified by tracking the packets a node generates/sends/receives and the amount of energy it consumes. Energy consumption is strongly correlated to the amount of network traffic handled by a node since the energy consumption for the operation of the radio is the primary energy consumer in typical nodes. However, directly measuring the remaining throughput is both more accurate and also works when nodes have atypical energy consumption profiles (e.g. increased node processing or high energy consumption sensors).

Calculating the remaining throughput then requires knowledge of the total throughput supported by a node and subtraction of the used throughput.

4.1. Parent selection problem

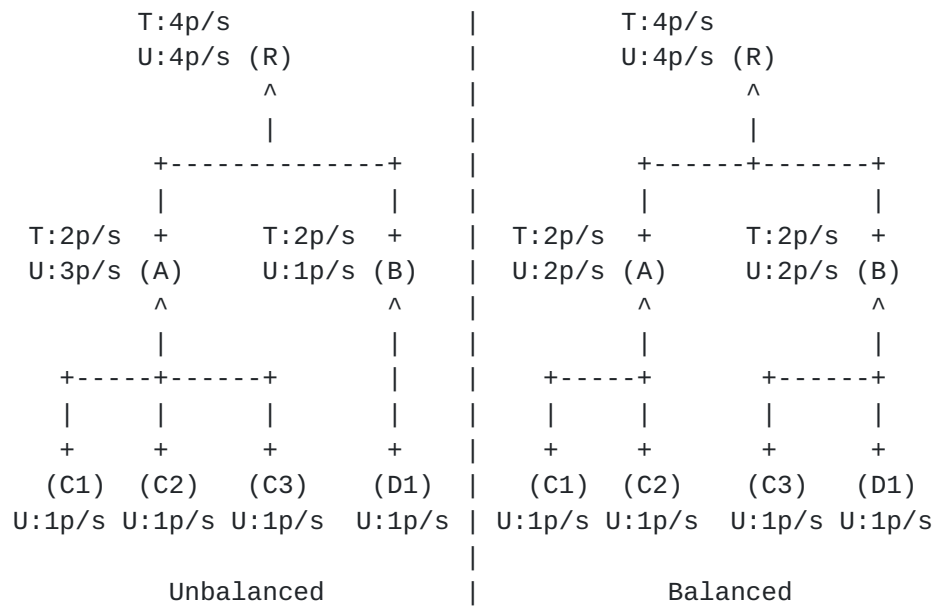


Figure 1: Use of Remaining Throughput with nodes with the same requirements

As a first simple example, an unbalanced network with nodes which all use the same throughput ("U:") is shown in Figure 1. Nodes A and B have the same total throughput ("T:"), but node A is overloaded due to trying to handle more than its ability while node B has a spare throughput of $2-1=1\text{p/s}$. Its transformation into a balanced network is shown on the right and it involves a node (C3) switching parents from A to B so that the capacity of its parent is no longer exceeded.

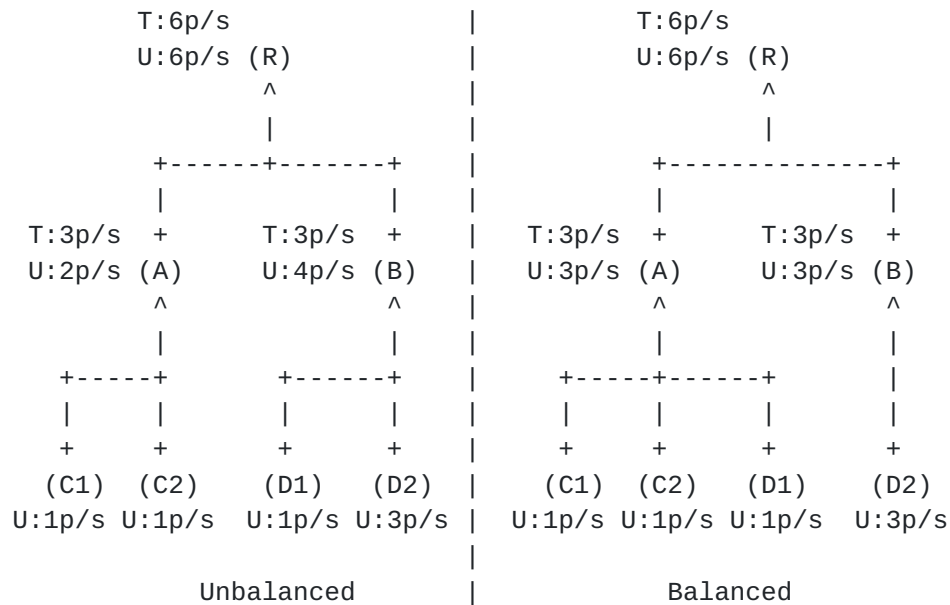


Figure 2: Use of Remaining Throughput with nodes with different requirements

As a second simple example, an unbalanced network with nodes which have different throughput ("U:") is shown in Figure 2. In this case, node B is overloaded and node D1 should move to parent A, which has a spare throughput of 1p/s. Its transformation into a balanced equivalent network is shown on the right.

4.2. DODAG selection problem

The purpose of the following example is to show the problem of DODAG selection, and not to focus on selecting the best parent.

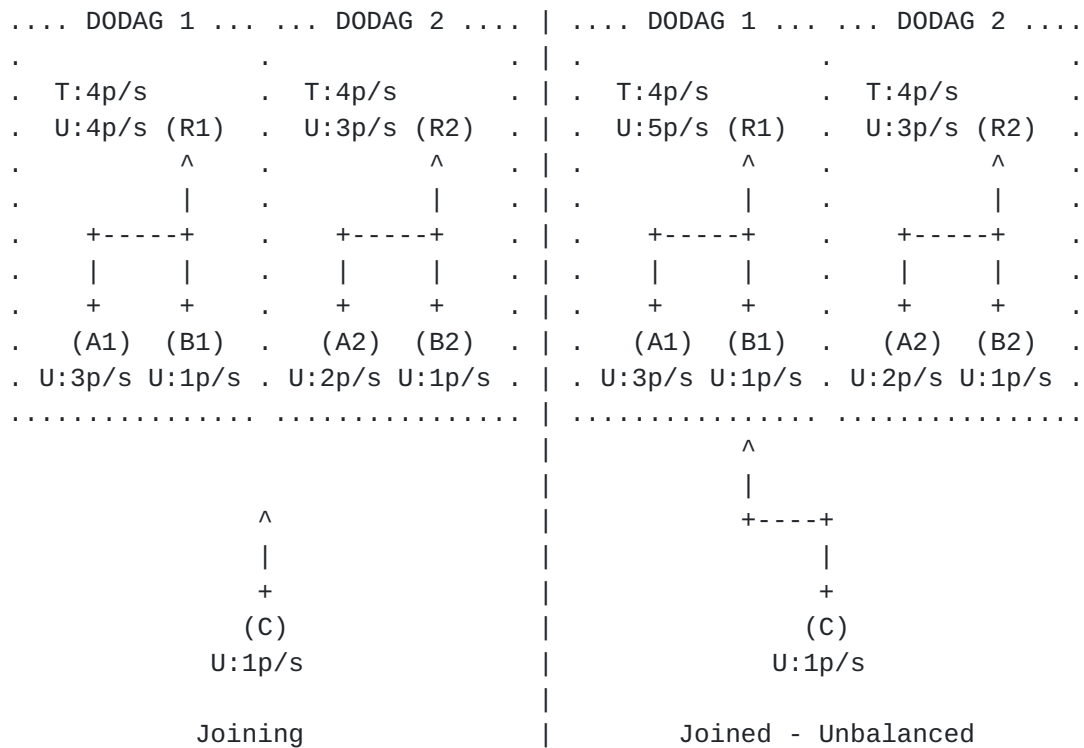


Figure 3: DODAG selection example leading to unbalanced traffic with RT metric

In the example in Figure 3, there are two DODAGs (DODAG 1 and DODAG 2) that belong to the same RPL Instance and a node (C) that must select a DODAG. Node C has to pick from the information provided by its two reachable neighbors: B1 and A2. On the left, node C is shown before selecting the preferred DODAG, while on the right it is shown after the DODAG selection.

Node C might choose B1 in DODAG 1 to be its preferred parent since the traffic information of the two DODAGs is not available. However, at the root node (R1), it can be observed that the total network traffic is higher in DODAG 1 and that after node C joins it, the traffic handling capacity of the root R1 has been exceeded.

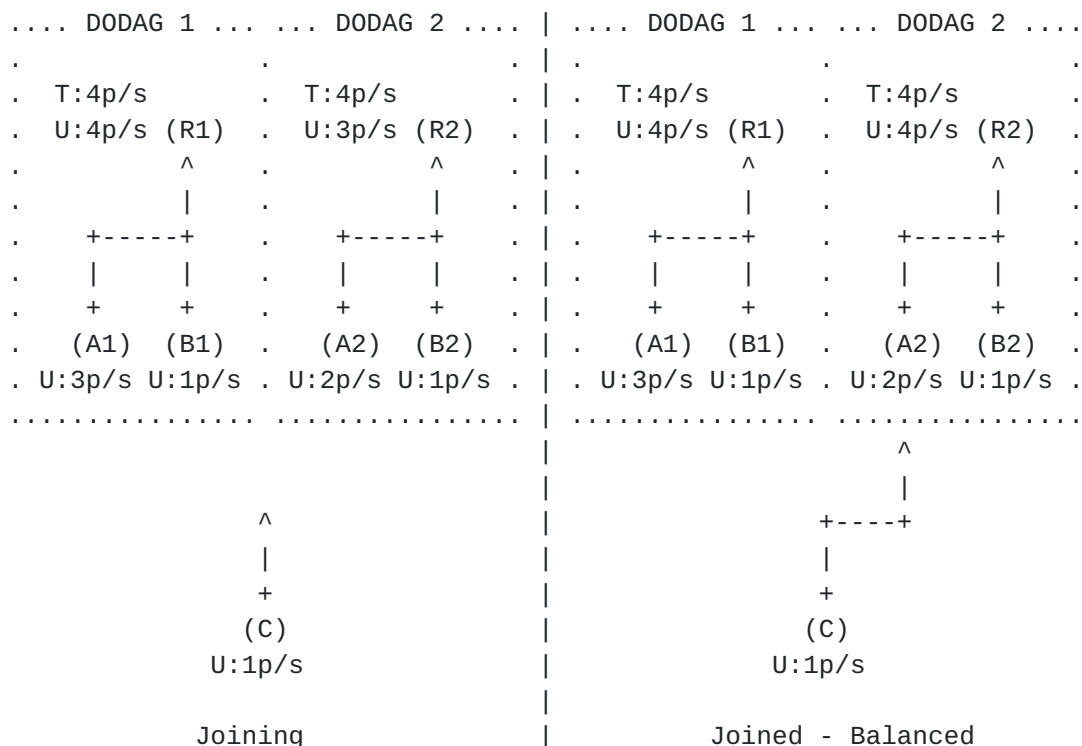


Figure 4: DODAG selection example leading to balanced traffic with RT metric

If the traffic handling capacity information is available, then node C could make a more efficient decision by using DODAG 2 and selecting node A2 as the preferred parent, as shown in Figure 4. Such a selection is based on the traffic of the entire DODAG and would not lead to exceeding the traffic handling capacity of the root R2 since this root had spare capacity.

5. TAOF description

In this specification, a metric is proposed to be used in the parent and DODAG selection mechanism, the Remaining Throughput (RT) which represents the number of packets each node can transmit (send or forward) during a certain time window. The window used, named THROUGHPUT_WINDOW, and the time unit used for this window THROUGHPUT_WINDOW_UNIT, are parameters common to the whole RPL instance. The window used SHOULD coincide with a sliding window of the same size used to calculate the packets transferred during this window to facilitate the calculation of the remaining possible packet transmissions. Therefore, whenever the RT value is reported it will refer to the previous THROUGHPUT_WINDOW window of time, expressed in THROUGHPUT_WINDOW_UNIT time units. This information is added in DIO messages and is broadcast to every neighbor. These parameters (THROUGHPUT_WINDOW, THROUGHPUT_WINDOW_UNIT) CAN be pre-configured on

all the nodes or CAN be provided along with the RT value. If they are transmitted along with the RT value, they MUST be set by the root node and retransmitted unchanged by all other nodes.

At first, each node MUST identify from their neighbor set which nodes are acceptable to be selected as a parent. For this purpose, the metric ETX is used as a filter to filter out parent candidates with low link quality with a preference for nodes with link quality below a given threshold. The ETX threshold SHOULD be different depending on application requirements. The suggested value for the relevant threshold MAX_PATH_COST from MRHOF [RFC6719] is 32768, which means the specific path has expected transmission counts greater than 256.

When the ETX value is used as a filter, nodes with bad link quality will not be included in the parent set. This ensures that undue retransmissions caused by bad links will be avoided. After all the filtering is done, if any, the node chooses the parent candidate or DODAG with the highest remaining throughput.

For the purpose of DODAG specifically, the A field in the Routing Metric/Constraint Flag field object [RFC6551] SHOULD be set to 2, indicating that the value reported is a minimum. More specifically, when a node is calculating the value of RT to broadcast in a DIO, the value reported SHOULD be the minimum of two values: its parent RT and the node's own calculated remaining throughput. Thus, the value broadcasted will be the available remaining throughput in the whole path from the node to the DODAG root.

This proposal is expected to increase the frequency of parent changes because the remaining throughput is more likely to be different between DIO messages, even for DIO messages from the same node. There are multiple ways to minimize the frequency of unnecessary parent changes:

- a. Use the remaining throughput in combination with another metric (e.g. child count, hop counts).
- b. Use a threshold when comparing the remaining throughput, similar to the approach in MRHOF [RFC6719]. Switch parents when the difference of remaining throughput between the original parent and the alternative parent is above a threshold. This threshold depends on different factors (e.g. network size, average traffic load) and SHOULD be defined differently for each use case.
- c. Take advantage of multiple parents in the parent set and slowly switch over traffic to less congested parents instead of switching all the traffic to a new preferred parent. The

difficulty lies in the method of partitioning the traffic to multiple parents.

6. DIO Metric Container Type extension




Figure 5: DAG metric container type format.

A DIO message carries fields as described in [RFC6550](#) [[RFC6550](#)] and the available options for the DAG metric container are described in [RFC6551](#) [[RFC6551](#)]. In this specification, a metric container option is proposed and the detailed format is shown in Figure 5. The information carried is the RT, represented as a 2 byte unsigned integer. Optional TLVs can define the THROUGHPUT_WINDOW time and the THROUGHPUT_WINDOW_UNIT time unit.

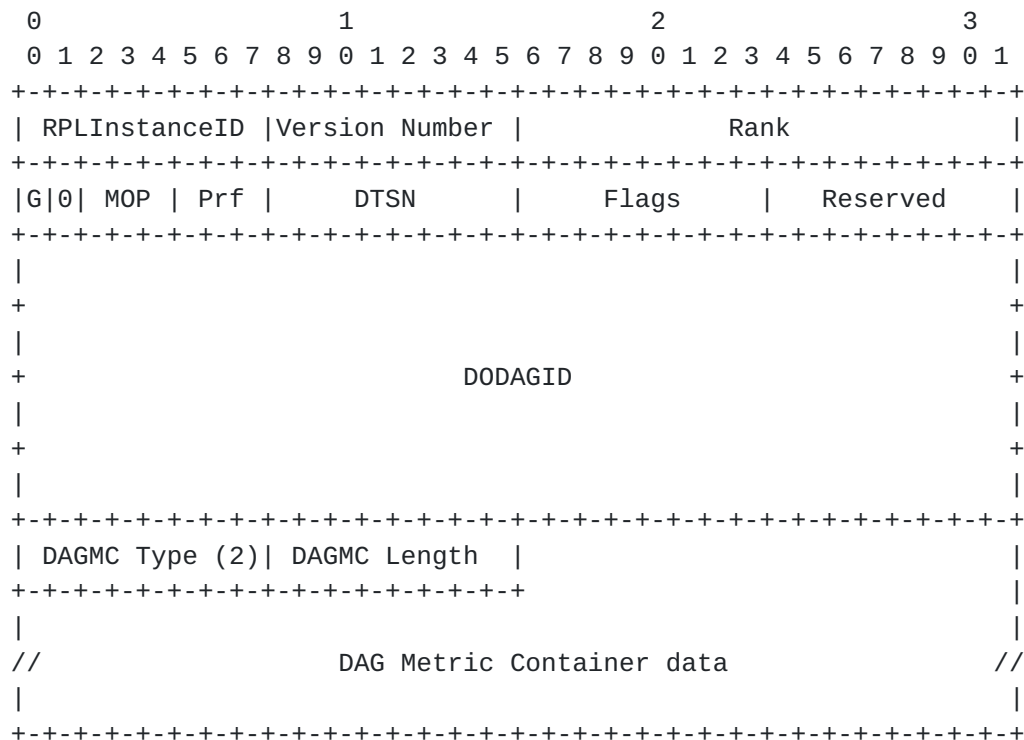


Figure 6: Example DIO Message with a DAG Metric Container option

The structure of the DIO Control Message when a DAG Metric Container option is included is shown in Figure 6. The DAG Metric Container

option type (DAGMC Type in Figure 6) has the value 0x02 as per the IANA registry for the RPL Control Message Options and is defined in [RFC6550]. The DAG Metric Container option length (DAGMC Length in Figure 6) expresses the DAG Metric Container length in bytes. DAG Metric Container data holds the actual data and is shown further expanded in Figure 7.

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Routing-MC-Type|Res Flags|P|C|O|R| A   |   Prec | Length (bytes)|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|
|           RT                       | Optional TLVs ...
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 7: DAG Metric Container (MC) data with Remaining Throughput (RT) object body

An example DAG Metric Container containing the proposed Metric Container object is shown in Figure 7. The explicit definition of the fields is:

Routing-MC-Type: TBD1. The type of the proposed DAGMC extension.
To be assigned by IANA.

Remaining Throughput (RT): The remaining throughput, represented as a 2-byte unsigned integer in units of packets per THROUGHPUT_WINDOW time in THROUGHPUT_WINDOW_UNIT time units.

Optional TLVs:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Type (TBD2)   | Length (2)   |   THROUGHPUT_WINDOW   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 8: THROUGHPUT_WINDOW TLV.

THROUGHPUT_WINDOW: Shown in Figure 8. Type is TBD2: The type of the proposed THROUGHPUT_WINDOW TLV. Length is 2 bytes consisting of only THROUGHPUT_WINDOW. THROUGHPUT_WINDOW expresses the size of the window used in units of THROUGHPUT_WINDOW_UNIT.




Figure 9: THROUGHPUT_WINDOW_UNIT TLV.

THROUGHPUT_WINDOW_UNIT: Shown in Figure 9. Type is TBD3: The type of the proposed THROUGHPUT_WINDOW_UNIT TLV. Length is 1 byte consisting of only THROUGHPUT_WINDOW_UNIT.

THROUGHPUT_WINDOW_UNIT expresses the time units used for THROUGHPUT_WINDOW. The resulting time unit is $2^{\text{THROUGHPUT_WINDOW_UNIT}}$ milliseconds (i.e. $\text{POWER}(2, \text{THROUGHPUT_WINDOW_UNIT}) * 0.001$ seconds). This results in a range for the time unit between 1 ms and approximately $5.7 * 10^{73}$ seconds.

6.1. DAG Metric Container fields

Given the intended usage, when using the RT, the DAGMC data it is contained in CAN be used as a constraint or as a metric in the DAG Metric Container. More specifically, using the RT places the following requirements on the DAG Metric Container header fields:

- o 'P' flag: Used as per [\[RFC6550\]](#).
- o 'C' flag: Used as per [\[RFC6550\]](#). When RT is used as a constraint, it indicates the minimum acceptable value of RT. Below this, the constraint MUST be considered violated.
- o 'O' flag: Used as per [\[RFC6550\]](#), to indicated optionality.
- o 'R' flag: Used as per [\[RFC6550\]](#).
- o 'A' Field: MUST be set to 2 "reports a minimum". This is the only aggregation function which makes sense for this metric.
- o 'Prec' Field: Used as per [\[RFC6550\]](#).

7. Enrollment

The RT metric SHOULD be used not only for ongoing parent selection but also especially within enrollment, i.e. the process a node follows to join a 6TiSCH network. In accordance to [\[I-D.ietf-6tisch-enrollment-enhanced-beacon\]](#), the value in RT SHOULD be used to affect the enrollment process so that a new node will be able to directly select a DODAG which will be able to cover its

traffic needs and to spread the traffic load between different DODAGs. More specifically, the pan priority field described in [\[I-D.ietf-6tisch-enrollment-enhanced-beacon\]](#) can be derived from the RT value. For this purpose, the RT value SHOULD be used with the maximum value aggregation mode (A field in the Routing Metric/Constraint Flag field object [\[RFC6551\]](#) set to 1), to report the maximum remaining throughput in the whole path to the DODAG root. The pan priority field is an unsigned 8-bit integer with lower values signifying higher priority while the RT value is a 16-bit unsigned integer with higher values signifying more remaining throughput. To convert the RT value to a pan priority the following formula should be used:

$$\text{pan priority} = 16 - \text{FLOOR}(\text{LOG}_2(\text{RT} + 1))$$

where LOG2 is the logarithm function with a base of 2. The use of the LOG function allows having higher accuracy in the low values of the remaining throughput, where small value differences are significant, and lower accuracy in the high values of the remaining throughput, where small differences are less significant. The addition of 1 to the RT allows converting RT=0.

8. Security Considerations

The structure of the DIO control message is extended, within the pre-defined DIO options. Therefore, the security mechanisms defined in RPL [\[RFC6550\]](#) apply to this proposed extension.

9. IANA Considerations

This proposal requests the allocation of a new value TBD1 for the metric type "RT" in the Routing-MC-Type field in the DAG MC from IANA. It also requests the allocation of two TLV type values (TBD2, TBD3) for use in the metric type "RT" in the Routing-MC-Type field in the DAG MC from IANA.

Additionally, an Objective Code Point (OCP) with value TBD4 for TAOF needs to be assigned in the Objective Code Point Registry as described in [Section 20.5 of \[RFC6550\]](#).

10. Informative references

[I-D.ietf-6tisch-enrollment-enhanced-beacon]
Dujovne, D. and M. Richardson, "IEEE802.15.4 Informational Element encapsulation of 6tisch Join and Enrollment Information", [draft-ietf-6tisch-enrollment-enhanced-beacon-01](#) (work in progress), January 2019.

[I-D.qasem-roll-rpl-load-balancing]

Qasem, M., Al-Dubai, A., Romdhani, I., Ghaleb, B., Hou, J., and R. Jadhav, "Load Balancing Objective Function in RPL", [draft-qasem-roll-rpl-load-balancing-02](#) (work in progress), October 2017.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", [RFC 6550](#), DOI 10.17487/RFC6550, March 2012, <<https://www.rfc-editor.org/info/rfc6550>>.

[RFC6551] Vasseur, JP., Ed., Kim, M., Ed., Pister, K., Dejean, N., and D. Barthel, "Routing Metrics Used for Path Calculation in Low-Power and Lossy Networks", [RFC 6551](#), DOI 10.17487/RFC6551, March 2012, <<https://www.rfc-editor.org/info/rfc6551>>.

[RFC6552] Thubert, P., Ed., "Objective Function Zero for the Routing Protocol for Low-Power and Lossy Networks (RPL)", [RFC 6552](#), DOI 10.17487/RFC6552, March 2012, <<https://www.rfc-editor.org/info/rfc6552>>.

[RFC6719] Gnawali, O. and P. Levis, "The Minimum Rank with Hysteresis Objective Function", [RFC 6719](#), DOI 10.17487/RFC6719, September 2012, <<https://www.rfc-editor.org/info/rfc6719>>.

Authors' Addresses

Remous-Aris Koutsiamanis (editor)
IMT Atlantique
Office B00 - 126A
2 Rue de la Chataigneraie
Cesson-Sevigne - Rennes 35510
FRANCE

Phone: +33 299 12 70 49
Email: aris@ariskou.com

Georgios Papadopoulos
IMT Atlantique
Office B00 - 114A
2 Rue de la Chataigneraie
Cesson-Sevigne - Rennes 35510
FRANCE

Phone: +33 299 12 70 04
Email: georgios.papadopoulos@imt-atlantique.fr

Eduardo Ingles Sanchez
University of Murcia
Campus de Espinardo S/N
Faculty of Computer Science
Murcia 30100
SPAIN

Phone: +34 868 88 96 73
Email: eduardo.ingles@um.es

Chenyang Ji
Alexander TEI of Thessaloniki
Department of Informatics
Thessaloniki 57400
GREECE

Email: jichenyang920@gmail.com

Diego Dujovne
Universidad Diego Portales
Escuela de Informatica y Telecomunicaciones, Av. Ejercito 441
Santiago, Region Metropolitana
Chile

Phone: +56 (2) 676-8121
Email: diego.dujovne@mail.udp.cl

Nicolas Montavont
IMT Atlantique
Office B00 - 106A
2 Rue de la Chataigneraie
Cesson-Sevigne - Rennes 35510
FRANCE

Phone: +33 299 12 70 23

Email: nicolas.montavont@imt-atlantique.fr