

Network Working Group
Internet-Draft
Intended status: Informational
Expires: August 2, 2012

T. Clausen
A. Colin de Verdiere
J. Yi
LIX, Ecole Polytechnique
U. Herberg
Fujitsu Laboratories of America
January 30, 2012

Experiences with RPL: IPv6 Routing Protocol for Low power and Lossy
Networks

draft-clausen-lln-rpl-experiences-00

Abstract

With RPL - the "IPv6 Routing Protocol for Low-power Lossy Networks" - having been published as a Proposed Standard after a ~2-year development cycle, this document presents an evaluation of the resulting protocol of its applicability and of its limits. The documents presents a selection of observations of the protocol characteristics, exposes experiences acquired when producing various prototype implementations of RPL, and presents results obtained from testing this protocol - by way of network simulations, in network testbeds and in deployments. The document aims at providing a better understanding of possible weaknesses and limits of RPL, notably the possible directions that further protocol developments should explore, in order to address these.

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 August 2, 2012.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the

Internet-Draft

Experiences with RPL

January 2012

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	3
2.	Terminology	3
3.	RPL Overview	4
3.1.	RPL Message Emission Timing - Trickle Timers	6
4.	Requirement Of DODAG Root	6
4.1.	Observations	7
5.	RPL Data Traffic Flows	8
5.1.	Observations	9
6.	Fragmentation Of RPL Control Messages And Data Packet	9
6.1.	Observations	10
7.	The DAO Mechanism: Downward and Point-to-Point Routes	10
7.1.	Observations	11
8.	Address Aggregation and Summarization	12
8.1.	Observations	13
9.	Links Assumed Bi-Directional	14
9.1.	Observations	14
10.	Neighbor Unreachability Detection For Unidirectional Links	14
10.1.	Observations	15
11.	RPL Implementability and Complexity	16
11.1.	Observations	16
12.	Underspecification	17
12.1.	Observations	17
13.	Protocol Convergence	18
13.1.	Observations	19
14.	Loops	19
14.1.	Observations	19
15.	Security Considerations	20
16.	IANA Considerations	21

17. Acknowledgements	21
18. Informative References	21
Authors' Addresses	23

[1.](#) Introduction

RPL - the "Routing Protocol for Low Power and Lossy Networks" [[I-D.ietf-roll-rpl](#)] - is a proposal for an IPv6 routing protocol for Low-power Lossy Networks (LLNs), by the ROLL Working Group in the Internet Engineering Task Force (IETF). This routing protocol is intended to be the IPv6 routing protocol for LLNs and sensor networks, applicable in all kinds of deployments and applications of LLNs.

The objective of RPL and ROLL is to target networks which "comprise up to thousands of nodes" [[roll-charter](#)], where the majority of the nodes have very constrained resources, where the network to a large degree is "managed" by a (single or few) central "supernodes" [[I-D.ietf-roll-rpl-terminology](#)], and where handling mobility is not an explicit design criteria ([[RFC5867](#)], [[RFC5826](#)], [[RFC5673](#)], [[RFC5548](#)]). Supported traffic patterns include multipoint-to-point, point-to-multipoint and point-to-point traffic. The emphasis among these traffic patterns is to optimize for multipoint-to-point traffic, to reasonably support point-to-multipoint traffic and to provide basic features for point-to-point traffic, in that order.

As of early 2011, RPL has been approved by the IESG, for publication as a "Proposed Standard" RFC (Request for Comments). The implication of a protocol being labeled "Proposed Standard" is that it is considered generally stable: well-understood and community reviewed, no known design issues pending, and with some community support [[RFC2026](#)].

"Proposed Standard" is, however, only the first step on the Standards Track, and it is thus opportune to consider the protocol in order to understand which aspects of it necessitate further investigations, and in order to identify possibly weak points which may restrict the deployment scope of the protocol. This document has as objective to provide observations of RPL, in the spirit of better understanding its characteristics and limits.

2. Terminology

This document uses the terminology and notation defined in [\[I-D.ietf-roll-rpl\]](#).

Additionally, it uses the following terminology:

Clausen, et al.

Expires August 2, 2012

[Page 3]

Internet-Draft

Experiences with RPL

January 2012

RPL Router - A device, running the RPL protocol, as specified by [\[I-D.ietf-roll-rpl\]](#).

3. RPL Overview

The basic construct in RPL is a "Destination Oriented Directed Acyclic Graph" (DODAG), depicted in Figure 1. In a converged LLN, each RPL Router has identified a stable set of parents, each of which is a potential next-hop on a path towards the "root" of the DODAG, as well as a preferred parent. Each RPL Router, which is part of a DODAG (i.e. has selected parents) will emit DODAG Information Object (DIO) messages, using link-local multicast, indicating its respective rank in the DODAG (i.e. distance to the DODAG Root according to some metric(s), in the simplest form hop-count). Upon having received a (number of such) DIO messages, an RPL Router will calculate its own rank such that it is greater than the rank of each of its parents, select a preferred parent and then itself start emitting DIO messages.




Figure 1: RPL DODAG

The DODAG formation thus starts at the DODAG Root (initially, the only RPL Router which is part of a DODAG), and spreads gradually to cover the whole LLN as DIOs are received, parents and preferred parents are selected and further RPL Routers participate in the DODAG. The DODAG Root also includes, in DIO messages, a DODAG Configuration Object, describing common configuration attributes for all RPL Routers in that network – including their mode of operation, timer characteristics etc. RPL Routers in a DODAG include a verbatim copy of the last received DODAG Configuration Object in their DIO messages, permitting also such configuration parameters propagating through the network.

As a Distance Vector protocol, RPL restricts the ability for an RPL Router to change rank. An RPL Router can freely assume a smaller rank than previously advertised (i.e. logically move closer to the DODAG Root) if it discovers a parent advertising a lower rank, and

must then disregard all previous parents of ranks higher than the router's new rank. The ability for an RPL Router to assume a greater rank (i.e. logically move farther from the DODAG Root) than previously advertised is restricted in order to avoid count-to-infinity problems. The DODAG Root can trigger "global recalculation" of the DODAG by increasing a sequence number, DODAG version, in DIO messages.

The DODAG so constructed is used for installing routes: the "preferred parent" of an RPL Router can serve as a default route towards the DODAG Root, or the DODAG Root can embed in its DIO messages the destination prefixes, included by DIOs generated by RPL Routers through the LLN, to which connectivity is provided by the DODAG Root. Thus, RPL by way of DIO generation provides "upward routes" or "multipoint-to-point routes" from the sensors inside the LLN and towards the DODAG Root.

"Downward routes" are enabled by having sensors issue Destination Advertisement Object (DAO) messages, propagating as unicast via parents towards the DODAG Root. These describe which prefixes belong to, and can be reached via, which RPL Router. In a network, all RPL Routers must operate in either of storing-mode or non-storing-mode,

specified by way of a "Mode of Operation" (MOP) flag in the DODAG Configuration Object from the DODAG Root. Depending on the MOP, DAO messages are forwarded differently towards the DODAG Root:

- o In "non-storing-mode", an RPL Router originates DAO messages, advertising one or more of its parents, and unicast it to the DODAG Root. Once the DODAG Root has received DAOs from an RPL Router, and from all RPL Routers on the path between it and the DODAG Root, it can use source routing for reaching advertised destinations inside the LLN.
- o In "storing-mode", each RPL Router on the path between the originator of a DAO and the DODAG Root records a route to the prefixes advertised in the DAO, as well as the next-hop towards these (the RPL Router, from which the DAO was received), then forwards the DAO to its preferred parent.

"Point-to-point routes", for communication between devices inside the LLN and where neither of the communicating devices are the DODAG Root, are as default supported by having the source sensor transmit via its default route to the DODAG Root (i.e., using the upward routes) which will then, depending on the "Mode of Operation" for the DODAG, either add a source-route to the received data for reaching the destination sensor (downward routes in non-storing-mode) or simply use hop-by-hop routing (downward routes in storing-mode). In the case of storing-mode, if the source and the destination for a

point-to-point communication share a common ancestor other than the DODAG Root, a downward route may be available in an RPL Router (and, thus, used) before the data packet reaching the DODAG Root.

3.1. RPL Message Emission Timing - Trickle Timers

RPL message generation is timer-based, with the DODAG Root being able to configure back-off of message emission intervals using Trickle [[RFC6206](#)]. Trickle, as used in RPL, stipulates that an RPL Router transmits a DIO "every so often" - except if receiving a number of DIOs from neighbor RPL routers, enabling the RPL Router to determine if its DIO transmission is redundant.

When an RPL Router transmits a DIO, there are two possible outcomes: either every neighbor RPL Router that hears the message finds that

the information contained is consistent with its own state (i.e., the received DODAG version number corresponds with that which the RPL Router has recorded and no better rank is advertised than that which is recorded in the parent set) - or, a recipient RPL Router detects that either the sender of the DIO or itself has out-of-date information. If the sender has out-of-date information, then the recipient RPL Router schedules transmission of a DIO to update this information. If the recipient RPL Router has out-of-date information, then it updates based on the information received in the DIO.

With Trickle, an RPL Router will schedule emission of a DIO at some time, t , in the future. When receiving a DIO containing information consistent with its own information, the RPL Router will record that "redundant information has been received" by incrementing a redundancy counter, c . At the time t , if c is below some "redundancy threshold", then it transmits its DIO. Otherwise, transmission of a DIO at this time is suppressed, c is reset and a new t is selected to twice as long time in the future - bounded by a pre-configured maximum value for t . If, on the other hand, the RPL Router has received an out-of-date DIO from one of its neighbors, t is reset to a pre-configured minimum value and c is set to zero. In both cases, at the expiration of t , the RPL Router will verify if c is below the "redundancy threshold" and if so transmit - otherwise, increase t and stay quiet.

[4.](#) Requirement Of DODAG Root

As indicated in [Section 3](#), the DODAG Root has both a special responsibility and is subject to special requirements. The DODAG Root is responsible for determining and maintaining the configuration parameters for the DODAG, and for initiating DIO emissions.

The DODAG Root is also responsible (in both storing and non-storing mode) for being able to, when downward routes are supported, maintain sufficient topological information to be able to construct routes to all destinations in the network. This entails that the DODAG Root is required to have sufficient memory and sufficient computational resources to be able to record a network graph containing all paths from itself and to all destinations and calculate paths - this, regardless of if the network is operating in storing or non-storing

mode.

The DODAG Root is also required to have sufficient energy available so as to be able to ensure the relay functions required (this, especially for non-storing mode, where all traffic transits through the DODAG Root).

4.1. Observations

In a given deployment, select RPL Routers can be provisioned with the required energy, memory and computational resources so as to serve as DODAG Roots, and be administratively configured as such - with the remainder of the RPL routers in the network being of typically lesser capacity. [[rpl-eval-UCB](#)] indicates that, in storing mode, a device with 10KB of memory scales up to about 30 RPL Routers - in a larger network (in storing or non-storing mode, both) the DODAG Root would require at least that much, most likely much more, memory. In non-storing mode, only the DODAG Root will require that much, or likely much more, memory with the other RPL Routers being able to make do with considerably less.

RPL Routers provisioned with resources to act as DODAG Roots, and administratively configured to act as such, represent a single point of failure in the network. As the memory requirements for DODAG Root and other RPL Routers are substantially different, unless all RPL Routers are provisioned with resources (memory, energy, ...) to act as DODAG Roots, effectively if the designated DODAG Root fails, the network fails and RPL is unable to operate. Even if electing another RPL Router as temporary DODAG root (e.g., for forming a "Floating" DODAG) for providing internal connectivity between RPL Routers, this router may not - likely, will not - have the necessary resources to satisfy the role as DODAG Root.

Thus, although in principle RPL provides, by way of "Floating DODAGs", protocol mechanisms for establishing a DODAG for providing internal connectivity even in case of failure of the administratively provisioned DODAG Root - especially in non-storing mode - it is unlikely that any RPL Routers not explicitly provisioned as DODAG Roots will have sufficient resources to undertake this task.

Another possible LLN scenario is that only internal point-to-point

connectivity is sought, and no RPL Router has a more "central" role than any other - a self-organizing LLN. Requiring special provisioning of a specific "super-node" as DODAG Root is both unnecessary and undesirable.

5. RPL Data Traffic Flows

RPL makes a-priori assumptions of data traffic patterns: sensor-to-root data traffic (multipoint-to-point) is predominant, root-to-sensor data traffic (point-to-multipoint) is rare and sensor-to-sensor (point-to-point) data traffic is extremely rare. While not specifically called out thus in [\[I-D.ietf-roll-rpl\]](#), the resulting protocol design reflects these assumptions in that the mechanism constructing multipoint-to-point paths is efficient in terms of control traffic generated and state required, point-to-multipoint path construction much less so - and point-to-point paths subject to potentially significant route stretch (routes going through the DODAG Root in non-storing mode) and over-the-wire overhead from using source routing (from the DODAG Root to the destination) (cf [Section 7](#)) - or, in case of storing mode, considerable memory requirements in all LLN routers inside the network (cf [Section 7](#)).

An RPL Router selects from among its parents a "preferred parent", to serve as a default route towards the DODAG Root (and to prefixes advertised by the DODAG Root). Thus, RPL provides "upward routes" or "multipoint-to-point routes" from the RPL Routers below the DODAG Root and towards the DODAG Root.

An RPL Router which wishes to act as a destination for data traffic ("downward routes" or "point-to-multipoint") issues DAOs upwards in the DODAG towards the DODAG Root, describing which prefixes belong to, and can be reached via, that RPL Router.

Point-to-Point routes between RPL Routers below the DODAG Root are supported by having the source RPL Router transmit, via its default route, data traffic towards the DODAG Root. In non-storing mode, the data traffic will reach the DODAG Root, which will reflect the data traffic downward towards the destination RPL Router, adding a strict source routing header indicating the precise path for the data traffic to reach the intended destination RPL Router. In storing mode, the source and the destination may possibly (although, may also not) have a common ancestor other than the DODAG Root, which may provide a downward route to the destination before data traffic reaching the DODAG Root.

[5.1.](#) Observations

The data traffic characteristics assumed by RPL do not represent a universal distribution of traffic patterns in LLNs:

- o There are scenarios where sensor-to-sensor traffic is a more common occurrence, documented e.g. in [[RFC5867](#)].
- o There are scenarios, where all traffic is bi-directional, e.g. in case sensor devices in the LLN are, in majority, "actively read": a request is issued by the root to a specific sensor, and the sensor value is expected returned.

For the former, all sensor-to-sensor paths include the root, possibly causing congestion on the communications medium near the root, and draining energy from the intermediate RPL Routers on an unnecessarily long path. If sensor-to-sensor traffic is common, RPL Routers near the root will be particularly solicited as relays, especially in non-storing mode. For the latter, all RPL Routers are required to generate DAOs, which generates a considerable control traffic overhead [[bidir](#)].

[6.](#) Fragmentation Of RPL Control Messages And Data Packet

Fragmentation of IP packets appears when the size of the IP datagram is larger than the Maximum Transmission Unit (MTU) supported by the link layer. When an IP packet is fragmented, all fragments of that IP packet must be successfully received by a router, in order that the IP packet is successfully received - otherwise, the whole IP packet is lost. Moreover, the additional link-layer frame overhead for each of the fragments increases the capacity required from the medium, and may consume more energy for transmitting a higher number of frames on the network interface.

RPL is an IPv6 routing protocol, designed to operate on constrained link layers, such as IEEE 802.15.4 [[ieee802154](#)], with a maximum MTU of 127 bytes - a deviation from the otherwise specified minimum MTU of 1280 bytes for IPv6 [[RFC2460](#)]. Reducing the need of fragmentation of packets on such a link layer, compression adaptation layers exist [[RFC4944](#)], [[RFC6282](#)], reducing the overhead of the IPv6 header from at least 40 octets to a minimum of 2 octets. With a physical layer packet size of 127 octets, a maximum frame overhead of 25 octets and 21 octets for link layer security [[RFC4944](#)], 81 octets remain for L2 payload. Further subtracting 2 octets for the compressed IPv6 header leaves 79 octets for L3 data payload.

The second L in LLN indicating Lossy [[roll-charter](#)], higher loss

rates than typically seen in IP networks are expected, rendering fragmentation important to avoid.

DIO messages consist of a mandatory base object, facilitating DODAG formation, and additional options for e.g. autoconfiguration and network management. The base object contains two unused octets, reserved for future use, resulting in two bytes of unnecessary zeros, sent with each DIO message. The Prefix Information option, used for automatic configuration of address, carries even four unused octets in order to be compatible with IPv6 neighbor discovery.

[6.1.](#) Observations

While 79 octets may seem to be sufficient to carry RPL control messages, consider the following: RPL control messages are carried in ICMPv6, and the mandatory ICMPv6 header consumes 4 octets. The DIO base another 24 octets. If link metrics are used, that consumes at least another 8 octets – and this is using a hop count metric; other metrics may require more. The DODAG Configuration Object consumes up to a further 16 octets, for a total of 52 octets. Adding a Prefix Information Object for address configuration consumes another 32 octets, for a total of 84 octets – thus exceeding the 79 octets available for L3 data payload and causing fragmentation of such a DIO. As a point of reference, the ContikiRPL [[rpl-contiki](#)] implementation includes both the DODAG Configuration option and the Prefix Information option in all DIO messages. Any other options, e.g. Route Information options indicating prefixes reachable through the DODAG Root, increase the overhead and thus the probability of fragmentation.

RPL may further increase the probability of fragmentation of also data traffic: for non-storing-mode, RPL employs source-routing for all downward traffic. [[I-D.ietf-6man-header](#)] specifies the RPL Source Routing header, which imposes a fixed overhead of 8 octets per IP packet leaving 71 octets remaining from the MTU – from which must be deducted a variable number of octets, depending on the length of the route. With fewer octets available for data payload, RPL thus increases the probability for fragmentation of also data packets. This, in particular, for longer paths, e.g. in point-to-point data

traffic between sensors inside the LLN, where data traffic transit through the DODAG Root and are then source-routed to the destination.

[7.](#) The DAO Mechanism: Downward and Point-to-Point Routes

RPL specifies two distinct and incompatible "modes of operation" for downward traffic: storing mode, where each RPL Router is assumed to maintain routes to all destinations in its sub-DODAG, i.e. RPL

Clausen, et al.

Expires August 2, 2012

[Page 10]

Internet-Draft

Experiences with RPL

January 2012

Routers that are "deeper down" in the DODAG, and non-storing mode, where only the DODAG Root stores routes to destinations inside the LLN, and where the DODAG Root employs strict source routing in order to route data traffic to the destination RPL Router.

[7.1.](#) Observations

In addition to possible fragmentation, as occurs when using potentially long source routing headers over a medium with a small MTU - similar to what is discussed in [Section 6](#) - the maximum length of the source routing header [[I-D.ietf-6man-header](#)] is limited to 136 octets, including an 8 octet long header. As each IPv6 address has a length of 16 octets, not more than 8 hops from the source to the destination are possible for "raw IPv6". Using address compression (e.g. as specified in [[RFC4944](#)]), the maximum path length may not exceed 64 hops. This excludes deployment of RPL for scenarios with long "chain-like" topologies, such as traffic lights along a street.

In storing mode, each RPL Router has to store routes for destinations in its sub-DODAG. This implies that, for RPL Routers near the DODAG Root, the required storage is only bounded by the number of paths to all other destinations in the network. As RPL targets constrained devices with little memory, but also has as ambition to be operating networks consisting of thousands of routers [[roll-charter](#)], the storing capacity on these RPL Routers may not be sufficient - or, at least, the storage requirements in RPL Routers "near the DODAG Root" and "far from the DODAG Root" is not homogenous, thus some sort of administrative deployment, and continued administrative maintenance as the network evolves, of devices is needed. Indeed, [[rpl-eval-UCB](#)] argues that practical experiences suggest that RPL in storing mode, with RPL routers having 10kB of RAM, should be limited to networks of less than ~30 RPL Routers. Aggregation / summarization of addresses

may be advanced as a possible argument that this issue is of little significance - [Section 8](#) discusses why such an argument does not apply.

In short, the mechanisms in RPL, when using storing mode, force the choice between requiring all RPL Routers to have sufficient memory to store route entries for all destinations (storing-mode) - or, suffer increased risk of fragmentation, and thus loss of data packets, while consuming network capacity by way of source routing through the DODAG Root.

In RPL, the "mode of operation" stipulates that either downward routes are not supported (MOP=0), or that they are supported by way of either storing or non-storing mode. In case downward routes are supported, RPL does not provide any mechanism for discriminating between which routes should or should not be maintained. In

particular, in order to calculate paths to a given destination, all intermediaries between the DODAG Root and that destination must themselves be reachable - effectively rendering downward routes in RPL an "all-or-none" situation. In case a destination is unreachable, all the DODAG Root may do is require all destinations to re-issue their DAOs, by way of issuing a DIO with an increased DODAG version number, possibly provoking a broadcast-storm-like situation. This, in particular, as [\[I-D.ietf-roll-rpl\]](#) does not specify DAO message transmission constraints, nor any mechanism for adapting DAO emission to the network capacity.

A final point on the DAO mechanism: RPL supports point-to-point traffic only by way of relaying through the DODAG. In networks where point-to-point traffic is no rare occasion, this causes unduly long paths (with possibly increased energy consumption, increased probability of packet losses) as well as possibly congestion around the DODAG Root.

[8.](#) Address Aggregation and Summarization

As indicated in [Section 7](#), in storing mode, an RPL Router is expected to be able to store routing entries for all destinations in its "sub-DODAG", i.e., routing entries for all destinations in the network where the path to the DODAG Root includes that RPL Router.

In the Internet, no single router stores explicit routing entries for all destinations. Rather, IP addresses are assigned hierarchically, such that an IP address does not only uniquely identify a network interface, but also its topological location in the network, as illustrated in Figure 2. All addresses with the same prefix are reachable by way of the same router - which can, therefore, advertise only that prefix. Other routers need only record a single routing entry for that prefix, knowing that as the IP packet reaches the router advertising that prefix, more precise routing information is available.




Figure 2: Address Hierarchies

8.1. Observations

In RPL, each RPL Router acquires a number of parents, as described in [Section 3](#), from among which it selects one as its preferred parent and, thus, next-hop on the path to the DODAG Root. RPL Routers maintain a parent set containing possibly more than a single parent so as to be able to rapidly select an alternative preferred parent, should the previously selected such become unavailable. Thus expected behavior is for an RPL Router to be able to change its point of attachment towards the DODAG Root. If IP addresses are assigned in a strictly hierarchical fashion, and if scalability of the routing state maintained in storing mode is based on this hierarchy, then this entails that each time an RPL Router changes its preferred parent, it must also change its own IP address - as well as cause RPL routers in its "sub-DODAG" to do the same. RPL does not specify signaling for reconfiguring addresses in a sub-DODAG.

A slightly less strict hierarchy can be envisioned, where an RPL Router can change its preferred parent without necessarily changing addresses of itself and of its sub-DODAG, provided that its former and new preferred parents both have the same preferred parent, and have addresses hierarchically assigned from that - from the "preferred grandparent". With reference to Figure 1, this could be changing its preferred parent from d to c, provided that both d and c have b as preferred parent. Doing so would impose a restriction on the parent-set selection, admitting only parents which have themselves the same parent - thus, no longer having a DODAG but a simple tree, losing redundancy in the network connectivity. RPL does

not specify rules for admitting only parents with identical grand-parents into the parent set - although such is not prohibited either, if the loss of redundancy from constructing a tree is acceptable.

The DODAG Root incrementing the DODAG version number is the mechanism by which RPL enables global reconfiguration of the network, reconstructing the DODAG with (intended) more optimal paths. In case of addressing hierarchies being enforced, so as to enable aggregation, this will either restrict the ability for an optimal DODAG construction, or will trigger global address autoconfiguration so as to ensure addressing hierarchies.

Finally, with IP addresses serving a dual role of an identifier of both an end-point for communication and a topological location in the network, changing the IP address of a device, so as to reflect a change in network topology, also entails interrupting ongoing communication to or through that device. Additional mechanisms (e.g. a DNS-like system) mapping "communications identifiers" and "IP addresses" are required.

9. Links Assumed Bi-Directional

Parents (and the preferred parent) are selected based on receipt of DIOs, without verification of the ability for an RPL Router to successfully communicate with the parent - i.e. without any bidirectionality check of links. However, the basic use of links is for "upward" routes, i.e. for the RPL Router to use a parent (the preferred parent) as relay towards the DODAG Root - in the opposite direction of the one in which the DIO was received.

9.1. Observations

Unidirectional links are no rare occurrence, such as is known from wireless multi-hop networks. If an RPL Router receives a DIO on such a unidirectional link, and selects the originator of the DIO as parent, that would be a bad choice: unicast traffic in the upward direction would be lost. If the RPL Router had verified the bidirectionality of links, it might have selected a better parent, to which it has a bidirectional link.

10. Neighbor Unreachability Detection For Unidirectional Links

[I-D.ietf-roll-rpl] suggests using Neighbor Unreachability Detection (NUD) [[RFC4861](#)] to detect and recover from the situation of unidirectional links between an RPL Router and its (preferred) parent(s). When, e.g., an RPL Router a tries (and fails) to actually

use RPL Router b for forwarding traffic, NUD is supposed engaged to detect and prompt corrective action, e.g. by way of selecting an alternative preferred parent.

NUD is based upon observing if a data packet is making forward progress towards the destination, either by way of indicators from upper-layer protocols (such as TCP and, though not called out in [\[RFC4861\]](#), also from lower-layer protocols such as Link Layer ACKs) or - failing that - by unicast probing by way of transmitting a unicast Neighbor Solicitation message and expecting that a solicited Neighbor Advertisement message be returned.

[10.1](#). Observations

An RPL Router may receive, transiently, a DIO from an RPL Router, closer (in terms of rank) to the DODAG Root than any other RPL Router from which a DIO has been received. Some, especially wireless, link layers may exhibit different transmission characteristics between multicast and unicast transmissions (such is the case for some implementations of IEEE 802.11b, where multicast/broadcast transmissions are sent at much lower bit-rates than are unicast. IEEE 802.11b is, of course, not suggested as a viable interface for LLNs, but serves to illustrate that such asymmetric designs exist), leading to a (multicast) DIO being received from farther away than a unicast transmission can reach. DIOs are sent (downward) using link-local multicast, whereas the traffic flowing in the opposite direction (upward) is unicast. Thus, a received (multicast) DIO may not be indicative of useful unicast connectivity - yet, RPL might cause this RPL Router to select this attractive RPL Router as its preferred parent. This may happen both at initialization or at any time during the LLN lifetime, as RPL allows attachment to a "better parent" at any time.

A DODAG so constructed may appear stable and converged until such time that unicast traffic is to be sent and, thus, NUD invoked. Detecting only at that point that unicast connectivity is not maintained, and causing local (and possibly global) repairs exactly at that time, may lead to traffic not being deliverable. As indicated in [Section 8](#), if scalability is dependent on addresses being assigned hierarchically, changing point-of-attachment may entail more than switching preferred parent.

Also, absent all RPL Routers consistently advertising their reachability through DAO messages, a protocol requiring bidirectional flows between the communicating devices, such as TCP, will be unable to operate.

Finally, upon having been notified by NUD that the "next hop" is

unreachable, an RPL Router must discard the preferred parent and select another - hoping that this time, the preferred parent is actually reachable. Also, if NUD indicates "no forward progress" based on an upper-layer protocol, there is no guarantee that the problem stems exclusively from the preferred parent being unreachable. Indeed, it may be a problem farther ahead, possibly outside the LLN, thus changing preferred parent will not alleviate the situation.

Incidentally, this stems from a fundamental difference between "fixed links in the Internet" and "wireless links": whereas the former, as a rule, are reliable, predictable and with losses being rare exceptions, the latter are characterized by frequent losses and general unpredictability.

11. RPL Implementability and Complexity

RPL is designed to operate on "RPL Routers [...] with constraints on processing power, memory, and energy (battery power)" [[I-D.ietf-roll-rpl](#)]. However, the 163 pages long specification of RPL, plus additional specifications for routing headers [[I-D.ietf-6man-header](#)], Trickle timer [[RFC6206](#)], routing metrics [[I-D.ietf-roll-routing-metrics](#)] and objective function [[I-D.ietf-roll-rpl-of0](#)], describes complex mechanisms (e.g. the upwards and downward data traffic, a security solution, manageability of RPL Routers, auxiliary functions for autoconfiguration of RPL Routers, etc.), and provides no less than 9 message types, and 10 different message options.

To give one example, the ContikiRPL implementation (<http://www.sics.se/contiki>), which provides only storing-mode and no security features, consumes about 50 KByte of memory. Sensor hardware, such as MSP430 sensor platforms, does not contain much more memory than that, i.e. there may not be much space left to deploy any application on the RPL Router.

11.1. Observations

Since RPL is intended as the routing protocol for LLNs, which covers all the diverse applications requirements listed in [[RFC5867](#)], [[RFC5673](#)], [[RFC5826](#)], [[RFC5548](#)], it is possible that (i) due to limited memory capacity of the RPL Routers, and (ii) due to expensive development cost of the routing protocol implementation, many RPL implementations will only support a partial set of features from the specification, leading to non-interoperable implementations.

stretching and source routing for point-to-point traffic, several additional Internet-Drafts are being discussed for adoption in the ROLL Working Group - adding complexity to an already complex specification which, it is worth recalling, was intended to be of a protocol for low-capacity devices.

[12.](#) Underspecification

While [[I-D.ietf-roll-rpl](#)] is verbose in many parts, as described in [Section 11](#), some mechanisms are underspecified.

While for DIOs, the Trickle timer specifies a relatively efficient and easy-to-understand timing for message transmission, the timing of DAO transmission is not explicit. As each DAO may have a limited lifetime, one "best guess" for implementers would be to send DAO periodically, just before the life-time of the previous DAO expires. Since DAOs may be lost, another "best guess" would be to send several DAOs shortly one after the other in order to increase probability that at least one DAO is successfully received.

The same underspecification applies for DAO-ACK messages: optionally, on reception of a DAO, an RPL Router may acknowledge successful reception by returning a DAO-ACK. Timing of DAO-ACK messages is unspecified by RPL.

[12.1.](#) Observations

By not specifying details about message transmission intervals and required actions when receiving DAO and DAO-ACKs, implementations may exhibit a bad performance if not carefully implemented. Some examples are:

1. If DAO messages are not sent in due time before the previous DAO expires (or if the DAO is lost during transmission), the routing entry will expire before it is renewed, leading to a possible data traffic loss.
2. RPL does not specify to use jitter [[RFC5148](#)] (i.e. small random delay for message transmissions). If DAOs are sent periodically,

adjacent RPL Routers may transmit DAO messages at the same time, leading to link layer collisions.

3. In non-storing mode, the "piece-wise calculation" of routes to a destination from which a DAO has been received, relies on previous reception of DAOs from intermediate RPL Routers along the path. If not all of these DAOs from intermediate RPL Routers have been received, route calculation is not possible, and DAO-

ACKs or data traffic cannot be sent to that destination.

Other examples of underspecification include the local repair mechanism, which may lead to loops and thus data traffic loss, if not carefully implemented: an RPL Router discovering that all its parents are unreachable, may - according to the RPL specification - "detach" from the DODAG, i.e. increase its own rank to infinity. It may then "poison" its sub-DODAG by advertising its infinite rank in its DIOs. If, however, the RPL Router receives a DIO before it transmits the "poisoned" DIO, it may attach to its own sub-DODAG, creating a loop. If, instead, it had waited some time before processing DIOs again, chances are it would have succeeded in poisoning its sub-DODAG and thus avoided the loop.

13. Protocol Convergence

Trickle [[RFC6206](#)] is used by RPL to schedule transmission of DIO messages, with the objective to minimize the amount of transmitted DIOs while ensuring a low convergence time of the network. The theoretical behavior of Trickle is well understood, and the convergence properties are well studied. Simulations of the mechanism, such as documented [[trickle-multicast](#)], confirm these theoretical studies.

In real-world environments, however, varying link qualities may cause the algorithm to converge less well: frequent message losses entail resets of the Trickle timer and more frequent and unpredicted message emissions.

This has been observed, e.g., in an experimental testbed: 69 RPL Routers (MSP430-based wireless sensor routers with IEEE 802.15.4, using [[rpl-contiki](#)] IPv6 stack and RPL without downward routes; the

parameters of the Trickle timer were set to the implementation defaults (minimum DIO interval: 4 s, DIO interval doublings: 8, redundancy constant: 10)) were positioned in a fixed grid topology. This resulted in DODAGs being constructed with an average of 2.45 children per RPL Router and an average rank of 3.58.

In this small test network, the number of DIO messages emitted - expectedly spiked within the first ~10 seconds. Alas, rather than taper off to become zero (as the simulation studies would suggest), the DIO emission rate remained constant at about 70 DIOs per second. Details on this experiment can be found in [[rpl-eval](#)].

[13.1.](#) Observations

The varying link quality in real-world environments results in frequent changes of the best parent, which triggers a reset of the Trickle timer and thus the emission of DIOs. Therefore Trickle does not converge as well for links that are fluctuating in quality as in theory.

The resulting higher control overhead due to frequent DIO emission, leads to higher bandwidth and energy consumption as well as possibly to an increased number of collisions of frames, as observed in [[trickle-multicast](#)].

[14.](#) Loops

[I-D.ietf-roll-rpl] states that it "guarantees neither loop free path selection nor tight delay convergence times, but can detect and repair a loop as soon as it is used. RPL uses this loop detection to ensure that packets make forward progress [...] and trigger repairs when necessary". This implies that a loop may only then be detected and fixed when data traffic is sent through the network.

In order to trigger a local repair, RPL relies on the "direction" information (with values "up" or "down"), contained in an IPv6 hop-by-hop option header from received a data packet. If an "upward"

data packet is received by an RPL Router, but the previous hop of the packet is listed with a lower rank in the neighbor set, the RPL Router concludes that there must be a routing loop and it may therefore trigger a local repair. For downward traffic in non-storing mode, the DODAG Root can detect loops if the same router identifier (i.e. IP address) appears at least twice in the path towards a destination.

14.1. Observations

The reason for RPL to repair loops only when detected by a data traffic transmission is to reduce control traffic overhead. However, there are two problems in repairing loops only when so triggered: (i) the triggered local repair mechanism delays forward progress of data packets, increasing end-to-end delays, and (ii) the data packet has to be buffered during repair.

(i) may seem as the lesser of the two problems, since in a number of applications, such as data acquisition in smart metering applications, an increased delay may be acceptable. However, for applications such as alarm signals or in home automation (e.g. a light switch), increased delay may be undesirable.

As for (ii), RPL is supposed to run on LLN routers with "constraints on [...] memory" [[I-D.ietf-roll-rpl](#)]; buffering incoming packets during the route repair may not be possible for all incoming data packets, leading to dropped packets. Depending on the transport protocol, these data packets must be retransmitted by the source or are definitely lost.

If carefully implemented with respect to avoiding loops before they occur, the impact of the loop detection in RPL may be minimized. However, it can be observed that with current implementations of RPL, such as the ContikiRPL implementation, loops do occur - and that frequently. During the same experiments described in [Section 13](#), a snapshot of the DODAG was made every ten seconds. In 74.14% of the 4114 snapshots, at least one loop was observed. Further investigation revealed that in all these cases the DODAG was partitioned, and the loop occurred in the sub-DODAG that no longer had a connection to the DODAG Root. When the link to the only parent of an RPL Router breaks, the RPL Router may increase its rank and - when receiving a DIO from an RPL Router in its sub-DODAG - attach

itself to its own sub-DODAG, thereby creating a loop - as detailed in [Section 12.1](#).

While it can be argued that the observed loops are harmless since they occur in a DODAG partition that has no connection to the DODAG Root, they show that the state of the network is inconsistent. Even worse, when the broken link re-appears, it is possible that in certain situations, the loop is only repaired when data traffic is sent, possibly leading to data loss (as described above). This can occur if the link to the previous parent is reestablished, but the rank of that previous parent has increased in the meantime.

Another problem with the loop repair mechanism arises in non-storing mode when using only downward traffic: while the DODAG Root can easily detect loops (as described above), it has no direct means to trigger a local repair where the loop occurs. Indeed, it can only trigger a global repair by increasing the DODAG version number, leading to a Trickle timer reset and increased control traffic overhead in the network caused by DIO messages, and therefore a possible energy drain of the RPL Routers and congestion of the channel.

[15.](#) Security Considerations

This document does currently not specify any security considerations. This document also does not provide any evaluation of the security mechanisms of RPL.

[16.](#) IANA Considerations

This document has no actions for IANA.

[17.](#) Acknowledgements

The authors would like to thank Matthias Philipp (INRIA) for his contributions to conducting many of the experiments, revealing or confirming the issues described in this document.

18. Informative References

- [I-D.ietf-6man-header]
Hui, J., Vasseur, J., Culler, D., and V. Manral, "An IPv6 Routing Header for Source Routes with RPL", work in progress [draft-ietf-6man-rpl-routing-header-07.txt](#), December 2011.
- [I-D.ietf-roll-routing-metrics]
Vasseur, J., Pister, K., Dejan, N., and D. Barthel, "Routing Metrics used for Path Calculation in Low Power and Lossy Networks", work in progress [draft-ietf-roll-routing-metrics-19](#), March 2011.
- [I-D.ietf-roll-rpl]
Winther, T., Thubert, P., Hui, J., Vasseur, J., Brandt, A., Kelsey, R., Levis, P., Piester, K., and R. Struik, "RPL: IPv6 Routing Protocol for Low power and Lossy Networks", work in progress [draft-ietf-roll-rpl-19.txt](#), March 2011.
- [I-D.ietf-roll-rpl-of0]
Thubert, P., "RPL Objective Function Zero", work in progress [draft-ietf-roll-of0-20](#), September 2011.
- [I-D.ietf-roll-rpl-terminology]
Vasseur, JP., "Terminology in Low power And Lossy Networks", work in progress [draft-ietf-roll-terminology-06](#), September 2011.
- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), October 1996.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), Decemer 1998.

- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#), September 2007.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler,

"Transmission of IPv6 Packets over IEEE 802.15.4 Networks", [RFC 4944](#), September 2007.

- [RFC5148] Clausen, T., Dearlove, C., and B. Adamson, "Jitter Considerations in Mobile Ad Hoc Networks (MANETs)", [RFC 5148](#), February 2008.
- [RFC5548] Dohler, M., Watteyne, T., Winter, T., and D. Barthel, "Routing Requirements for Urban Low-Power and Lossy Networks", [RFC 5548](#), May 2009.
- [RFC5673] Pister, K., Thubert, P., Dwars, S., and T. Phinney, "Industrial Routing Requirements in Low-Power and Lossy Networks", [RFC 5673](#), October 2009.
- [RFC5826] Brandt, A., Buron, J., and G. Porcu, "Home Automation Routing Requirements in Low-Power and Lossy Networks", [RFC 5826](#), April 2010.
- [RFC5867] Martocci, J., Mi, P., Riou, N., and W. Vermeylen, "Building Automation Routing Requirements in Low Power and Lossy Networks", [RFC 5867](#), June 2010.
- [RFC6206] Levis, P., Clausen, T., Gnawali, O., and J. Ko, "The Trickle Algorithm", [RFC 6206](#), March 2011.
- [RFC6282] Hui, J. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", [RFC 6282](#), September 2011.
- [bidir] Clausen, T. and U. Herberg, "A Comparative Performance Study of the Routing Protocols LOAD and RPL with Bi-Directional Traffic in Low-power and Lossy Networks (LLN)", Proceedings of the Eighth ACM International Symposium on Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks (PE-WASUN), 2011.
- [ieee802154] Computer Society, IEEE., "IEEE Std. 802.15.4-2003", October 2003.
- [roll-charter] "ROLL Charter",

web <http://datatracker.ietf.org/wg/roll/charter/>.

[rpl-contiki]

Tsiftes, N., Eriksson, J., and A. Dunkels, "Low-Power Wireless IPv6 Routing with ContikiRPL", Proceedings of the 9th ACM/IEEE International Conference on Information Processing in Sensor Networks (ISPN), 2011.

[rpl-eval]

Clausen, T., Herberg, U., and M. Philipp, "A Critical Evaluation of the IPv6 Routing Protocol for Low Power and Lossy Networks (RPL)", Proceedings of the 5th IEEE International Conference on Wireless & Mobile Computing, Networking & Communication (WiMob), 2011.

[rpl-eval-UCB]

Ko, J., Dawson-Haggerty, S., Culler, D., and A. Terzis, "Evaluating the Performance of RPL and 6LoWPAN in TinyOS", Proceedings of the Workshop on Extending the Internet to Low power and Lossy Networks (IP+SN), 2011.

[trickle-multicast]

Clausen, T. and U. Herberg, "Study of Multipoint-to-Point and Broadcast Traffic Performance in the 'IPv6 Routing Protocol for Low Power and Lossy Networks' (RPL)", Journal of Ambient Intelligence and Humanized Computing, 2011.

Authors' Addresses

Thomas Clausen
LIX, Ecole Polytechnique
91128 Palaiseau Cedex,
France

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

Internet-Draft

Experiences with RPL

January 2012

Axel Colin de Verdier
LIX, Ecole Polytechnique
91128 Palaiseau Cedex,
France

Phone: +33 6 1264 7119
Email: axel@axelcdv.com
URI: <http://www.axelcdv.com/>

Jiazi Yi
LIX, Ecole Polytechnique
91128 Palaiseau Cedex,
France

Phone: +33 1 6933 4031
Email: jiazi@jiaziyi.com
URI: <http://www.jiaziyi.com/>

Ulrich Herberg
Fujitsu Laboratories of America
1240 E Arques Ave
Sunnyvale, CA 94085
USA

Email: ulrich@herberg.name
URI: <http://www.herberg.name/>

