

ROLL
Internet-Draft
Intended status: Standards Track
Expires: April 21, 2019

S. Anamalamudi
SRM University-AP
M. Zhang
Huawei Technologies
C. Perkins
Futurewei
S.V.R.Anand
Indian Institute of Science
B. Liu
Huawei Technologies
October 18, 2018

Asymmetric AODV-P2P-RPL in Low-Power and Lossy Networks (LLNs)
draft-ietf-roll-aodv-rpl-05

Abstract

Route discovery for symmetric and asymmetric Point-to-Point (P2P) traffic flows is a desirable feature in Low power and Lossy Networks (LLNs). For that purpose, this document specifies a reactive P2P route discovery mechanism for both hop-by-hop routing and source routing: Ad Hoc On-demand Distance Vector Routing (AODV) based RPL protocol. Paired Instances are used to construct directional paths, in case some of the links between source and target node are asymmetric.

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 April 21, 2019.

Copyright Notice

Copyright (c) 2018 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	3
2. Terminology	4
3. Overview of AODV-RPL	6
4. AODV-RPL DIO Options	7
4.1. AODV-RPL DIO RREQ Option	7
4.2. AODV-RPL DIO RREP Option	9
4.3. AODV-RPL DIO Target Option	10
5. Symmetric and Asymmetric Routes	11
6. AODV-RPL Operation	13
6.1. Route Request Generation	13
6.2. Receiving and Forwarding RREQ messages	14
6.2.1. General Processing	14
6.2.2. Additional Processing for Multiple Targets	15
6.3. Generating Route Reply (RREP) at TargNode	16
6.3.1. RREP-DIO for Symmetric route	16
6.3.2. RREP-DIO for Asymmetric Route	16
6.3.3. RPLInstanceID Pairing	16
6.4. Receiving and Forwarding Route Reply	17
7. Gratuitous RREP	18
8. Operation of Trickle Timer	19
9. IANA Considerations	19
9.1. New Mode of Operation: AODV-RPL	19
9.2. AODV-RPL Options: RREQ, RREP, and Target	19
10. Security Considerations	20
11. Future Work	20
12. Contributors	20
13. References	20
13.1. Normative References	21
13.2. Informative References	22
Appendix A. Example: ETX/RSSI Values to select S bit	22
Appendix B. Changelog	23

B.1. Changes to version 02	23
B.2. Changes to version 03	23
B.3. Changes to version 04	24
Authors' Addresses	24

1. Introduction

RPL[RFC6550] is a IPv6 distance vector routing protocol for Low-power and Lossy Networks (LLNs), and is designed to support multiple traffic flows through a root-based Destination-Oriented Directed Acyclic Graph (DODAG). Typically, a router does not have routing information for most other routers. Consequently, for traffic between routers within the DODAG (i.e., Point-to-Point (P2P) traffic) data packets either have to traverse the root in non-storing mode, or traverse a common ancestor in storing mode. Such P2P traffic is thereby likely to traverse longer routes and may suffer severe congestion near the DAG root [RFC6997], [RFC6998].

To discover better paths for P2P traffic flows in RPL, P2P-RPL [RFC6997] specifies a temporary DODAG where the source acts as a temporary root. The source initiates DIOs encapsulating the P2P Route Discovery option (P2P-RDO) with an address vector for both hop-by-hop mode (H=1) and source routing mode (H=0). Subsequently, each intermediate router adds its IP address and multicasts the P2P mode DIOs, until the message reaches the target node (TargNode), which then sends the "Discovery Reply" object. P2P-RPL is efficient for source routing, but much less efficient for hop-by-hop routing due to the extra address vector overhead. However, for symmetric links, when the P2P mode DIO message is being multicast from the source hop-by-hop, receiving nodes can infer a next hop towards the source. When TargNode subsequently replies to the source along the established forward route, receiving nodes determine the next hop towards TargNode. For hop-by-hop routes (H=1) over symmetric links, this would allow efficient use of routing tables for P2P-RDO messages instead of the "Address Vector".

RPL and P2P-RPL both specify the use of a single DODAG in networks of symmetric links, where the two directions of a link MUST both satisfy the constraints of the objective function. This disallows the use of asymmetric links which are qualified in one direction. But, application-specific routing requirements as defined in IETF ROLL Working Group [RFC5548], [RFC5673], [RFC5826] and [RFC5867] may be satisfied by routing paths using bidirectional asymmetric links. For this purpose, [I-D.thubert-roll-asymlink] described bidirectional asymmetric links for RPL [RFC6550] with Paired DODAGs, for which the DAG root (DODAGID) is common for two Instances. This can satisfy application-specific routing requirements for bidirectional asymmetric links in core RPL [RFC6550]. Using P2P-RPL twice with

Paired DODAGs, on the other hand, requires two roots: one for the source and another for the target node due to temporary DODAG formation. For networks composed of bidirectional asymmetric links (see Section 5), AODV-RPL specifies P2P route discovery, utilizing RPL with a new MoP. AODV-RPL makes use of two multicast messages to discover possibly asymmetric routes, which can achieve higher route diversity. AODV-RPL eliminates the need for address vector overhead in hop-by-hop mode. This significantly reduces the control packet size, which is important for Constrained LLN networks. Both discovered routes (upward and downward) meet the application specific metrics and constraints that are defined in the Objective Function for each Instance [RFC6552].

The route discovery process in AODV-RPL is modeled on the analogous procedure specified in AODV [RFC3561]. The on-demand nature of AODV route discovery is natural for the needs of peer-to-peer routing in RPL-based LLNs. AODV terminology has been adapted for use with AODV-RPL messages, namely RREQ for Route Request, and RREP for Route Reply. AODV-RPL currently omits some features compared to AODV -- in particular, flagging Route Errors, blacklisting unidirectional links, multihoming, and handling unnumbered interfaces.

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 [RFC2119]. This document uses the following terms:

AODV

Ad Hoc On-demand Distance Vector Routing[RFC3561].

AODV-RPL Instance

Either the RREQ-Instance or RREP-Instance

Asymmetric Route

The route from the OrigNode to the TargNode can traverse different nodes than the route from the TargNode to the OrigNode. An asymmetric route may result from the asymmetry of links, such that only one direction of the series of links fulfills the constraints in route discovery.

Bi-directional Asymmetric Link

A link that can be used in both directions but with different link characteristics.

DIO

DODAG Information Object

DODAG RREQ-Instance (or simply RREQ-Instance)

RPL Instance built using the DIO with RREQ option; used for control message transmission from OrigNode to TargNode, thus enabling data transmission from TargNode to OrigNode.

DODAG RREP-Instance (or simply RREP-Instance)

RPL Instance built using the DIO with RREP option; used for control message transmission from TargNode to OrigNode thus enabling data transmission from OrigNode to TargNode.

Downward Direction

The direction from the OrigNode to the TargNode.

Downward Route

A route in the downward direction.

hop-by-hop routing

Routing when each node stores routing information about the next hop.

on-demand routing

Routing in which a route is established only when needed.

OrigNode

The IPv6 router (Originating Node) initiating the AODV-RPL route discovery to obtain a route to TargNode.

Paired DODAGs

Two DODAGs for a single route discovery process between OrigNode and TargNode.

P2P

Point-to-Point -- in other words, not constrained a priori to traverse a common ancestor.

reactive routing

Same as "on-demand" routing.

RREQ-DIO message

An AODV-RPL MoP DIO message containing the RREQ option. The RPLInstanceID in RREQ-DIO is assigned locally by the OrigNode.

RREP-DIO message

An AODV-RPL MoP DIO message containing the RREP option. The RPLInstanceID in RREP-DIO is typically paired to the one in the associated RREQ-DIO message.

Source routing

A mechanism by which the source supplies the complete route towards the target node along with each data packet [RFC6550].

Symmetric route

The upstream and downstream routes traverse the same routers.

TargNode

The IPv6 router (Target Node) for which OrigNode requires a route and initiates Route Discovery within the LLN network.

Upward Direction

The direction from the TargNode to the OrigNode.

Upward Route

A route in the upward direction.

ART option

AODV-RPL Target option: a target option defined in this document.

3. Overview of AODV-RPL

With AODV-RPL, routes from OrigNode to TargNode within the LLN network established are "on-demand". In other words, the route discovery mechanism in AODV-RPL is invoked reactively when OrigNode has data for delivery to the TargNode but existing routes do not satisfy the application's requirements. The routes discovered by AODV-RPL are not constrained to traverse a common ancestor. Unlike RPL [RFC6550] and P2P-RPL [RFC6997], AODV-RPL can enable asymmetric communication paths in networks with bidirectional asymmetric links. For this purpose, AODV-RPL enables discovery of two routes: namely, one from OrigNode to TargNode, and another from TargNode to OrigNode. When possible, AODV-RPL also enables symmetric route discovery along Paired DODAGs (see Section 5).

In AODV-RPL, routes are discovered by first forming a temporary DAG rooted at the OrigNode. Paired DODAGs (Instances) are constructed according to the AODV-RPL Mode of Operation (MoP) during route formation between the OrigNode and TargNode. The RREQ-Instance is formed by route control messages from OrigNode to TargNode whereas the RREP-Instance is formed by route control messages from TargNode to OrigNode. Intermediate routers join the Paired DODAGs based on the rank as calculated from the DIO message. Henceforth in this document, the RREQ-DIO message means the AODV-RPL mode DIO message from OrigNode to TargNode, containing the RREQ option (see Section 4.1). Similarly, the RREP-DIO message means the AODV-RPL mode DIO message from TargNode to OrigNode, containing the RREP option (see Section 4.2). The route discovered in the RREQ-Instance is used for transmitting data from TargNode to OrigNode, and the

route discovered in RREP-Instance is used for transmitting data from OrigNode to TargNode.

4. AODV-RPL DIO Options

4.1. AODV-RPL DIO RREQ Option

OrigNode sets its IPv6 address in the DODAGID field of the RREQ-DIO message. A RREQ-DIO message MUST carry exactly one RREQ option.




Figure 1: DIO RREQ option format for AODV-RPL MoP

OrigNode supplies the following information in the RREQ option:

Type

The type assigned to the RREQ option (see Section 9.2).

Option Length

The length of the option in octets, excluding the Type and Length fields. Variable due to the presence of the address vector and the number of octets elided according to the Compr value.

S

Symmetric bit indicating a symmetric route from the OrigNode to the router transmitting this RREQ-DIO.

H

Set to one for a hop-by-hop route. Set to zero for a source route. This flag controls both the downstream route and upstream route.

X

Reserved.

Compr

4-bit unsigned integer. Number of prefix octets that are elided from the Address Vector. The octets elided are shared with the IPv6 address in the DODAGID. This field is only used in source routing mode (H=0). In hop-by-hop mode (H=1), this field MUST be set to zero and ignored upon reception.

L

2-bit unsigned integer determining the duration that a node is able to belong to the temporary DAG in RREQ-Instance, including the OrigNode and the TargNode. Once the time is reached, a node MUST leave the DAG and stop sending or receiving any more DIOs for the temporary DODAG. The definition for the "L" bit is similar to that found in [RFC6997], except that the values are adjusted to enable arbitrarily long route lifetime.

- * 0x00: No time limit imposed.
- * 0x01: 16 seconds
- * 0x02: 64 seconds
- * 0x03: 256 seconds

L is independent from the route lifetime, which is defined in the DODAG configuration option. The route entries in hop-by-hop routing and states of source routing can still be maintained even after the DAG expires.

MaxRank

This field indicates the upper limit on the integer portion of the rank (calculated using the DAGRank() macro defined in [RFC6550]). A value of 0 in this field indicates the limit is infinity.

Orig SeqNo

Sequence Number of OrigNode, defined similarly as in AODV [RFC3561].

Address Vector

A vector of IPv6 addresses representing the route that the RREQ-DIO has passed. It is only present when the 'H' bit is set to 0. The prefix of each address is elided according to the Compr field.

A node MUST NOT join a RREQ instance if its own rank would equal to or higher than MaxRank. Targnode can join the RREQ instance at a rank whose integer portion is equal to the MaxRank. A router MUST discard a received RREQ if the integer part of the advertised rank equals or exceeds the MaxRank limit. This definition of MaxRank is the same as that found in [RFC6997].

4.2. AODV-RPL DIO RREP Option

TargNode sets its IPv6 address in the DODAGID field of the RREP-DIO message. A RREP-DIO message MUST carry exactly one RREP option. TargNode supplies the following information in the RREP option:




Figure 2: DIO RREP option format for AODV-RPL MoP

Type

The type assigned to the RREP option (see Section 9.2)

Option Length

The length of the option in octets, excluding the Type and Length fields. Variable due to the presence of the address vector and the number of octets elided according to the Compr value.

G

Gratuitous route (see Section 7).

H

Requests either source routing (H=0) or hop-by-hop (H=1) for the downstream route. It MUST be set to be the same as the 'H' bit in RREQ option.

X

Reserved.

Compr

4-bit unsigned integer. Same definition as in RREQ option.

L

2-bit unsigned integer defined as in RREQ option.

MaxRank

Similarly to MaxRank in the RREQ message, this field indicates the upper limit on the integer portion of the rank. A value of 0 in this field indicates the limit is infinity.

Shift

6-bit unsigned integer. This field is used to recover the original InstanceID (see Section 6.3.3); 0 indicates that the original InstanceID is used.

Rsv

MUST be initialized to zero and ignored upon reception.

Address Vector

Only present when the 'H' bit is set to 0. For an asymmetric route, the Address Vector represents the IPv6 addresses of the route that the RREP-DIO has passed. For a symmetric route, it is the Address Vector when the RREQ-DIO arrives at the TargNode, unchanged during the transmission to the OrigNode.

4.3. AODV-RPL DIO Target Option

The AODV-RPL Target (ART) Option is defined based on the Target Option in core RPL [RFC6550]: the Destination Sequence Number of the TargNode is added.

A RREQ-DIO message MUST carry at least one ART Options. A RREP-DIO message MUST carry exactly one ART Option.

OrigNode can include multiple TargNode addresses via multiple AODV-RPL Target Options in the RREQ-DIO, for routes that share the same constraints. This reduces the cost to building only one DODAG. Furthermore, a single Target Option can be used for different TargNode addresses if they share the same prefix; in that case the use of the destination sequence number is not defined in this document.




Figure 3: Target option format for AODV-RPL MoP

Type

The type assigned to the ART Option

Dest SeqNo

In RREQ-DIO, if nonzero, it is the last known Sequence Number for TargNode for which a route is desired. In RREP-DIO, it is the destination sequence number associated to the route.

5. Symmetric and Asymmetric Routes

In Figure 4 and Figure 5, BR is the Border Router, O is the OrigNode, R is an intermediate router, and T is the TargNode. If the RREQ-DIO arrives over an interface that is known to be symmetric, and the 'S' bit is set to 1, then it remains as 1, as illustrated in Figure 4. If an intermediate router sends out RREQ-DIO with the 'S' bit set to 1, then all the one-hop links on the route from the OrigNode O to this router meet the requirements of route discovery, and the route can be used symmetrically.




Figure 4: AODV-RPL with Symmetric Paired Instances

Upon receiving a RREQ-DIO with the 'S' bit set to 1, a node determines whether this one-hop link can be used symmetrically, i.e., both the two directions meet the requirements of data transmission. If the RREQ-DIO arrives over an interface that is not known to be symmetric, or is known to be asymmetric, the 'S' bit is set to 0. If the 'S' bit arrives already set to be '0', it is set to be '0' on retransmission (Figure 5). Therefore, for asymmetric route, there is at least one hop which doesn't fulfill the constraints in the two directions. Based on the 'S' bit received in RREQ-DIO, the TargNode T determines whether or not the route is symmetric before transmitting the RREP-DIO message upstream towards the OrigNode O.

The criteria used to determine whether or not each link is symmetric is beyond the scope of the document, and may be implementation-specific. For instance, intermediate routers MAY use local information (e.g., bit rate, bandwidth, number of cells used in 6tisch), a priori knowledge (e.g. link quality according to previous communication) or use averaging techniques as appropriate to the application.

Appendix A describes an example method using the ETX and RSSI to estimate whether the link is symmetric in terms of link quality is given in using an averaging technique.




Figure 5: AODV-RPL with Asymmetric Paired Instances

6. AODV-RPL Operation

6.1. Route Request Generation

The route discovery process is initiated when an application at the OrigNode has data to be transmitted to the TargNode, but does not have a route for the target that fulfills the requirements of the data transmission. In this case, the OrigNode builds a local RPLInstance and a DODAG rooted at itself. Then it transmits a DIO message containing exactly one RREQ option (see Section 4.1) via link-local multicast. The DIO MUST contain at least one ART Option (see Section 4.3). The 'S' bit in RREQ-DIO sent out by the OrigNode is set to 1.

Each node maintains a sequence number, which rolls over like a lollipop counter [Perlman83], detailed operation can refer to the section 7.2 of [RFC6550]. When the OrigNode initiates a route discovery process, it MUST increase its own sequence number to avoid conflicts with previous established routes. The increased number is carried in the OrigSeqNo field of the RREQ option.

The address in the ART Option can be a unicast IPv6 address or a prefix. The OrigNode can initiate the route discovery process for multiple targets simultaneously by including multiple ART Options,

and within a RREQ-DIO the requirements for the routes to different TargNodes MUST be the same.

OrigNode can maintain different RPLInstances to discover routes with different requirements to the same targets. Using the InstanceID pairing mechanism (see Section 6.3.3), route replies (RREP-DIOs) for different RPLInstances can be distinguished.

The transmission of RREQ-DIO obeys the Trickle timer. If the duration specified by the "L" bit has elapsed, the OrigNode MUST leave the DODAG and stop sending RREQ-DIOs in the related RPLInstance.

6.2. Receiving and Forwarding RREQ messages

6.2.1. General Processing

Upon receiving a RREQ-DIO, a router which does not belong to the RREQ-instance goes through the following steps:

Step 1:

If the 'S' bit in the received RREQ-DIO is set to 1, the router MUST check the two directions of the link by which the RREQ-DIO is received. In case that the downward (i.e. towards the TargNode) direction of the link can't fulfill the requirements, the link can't be used symmetrically, thus the 'S' bit of the RREQ-DIO to be sent out MUST be set as 0. If the 'S' bit in the received RREQ-DIO is set to 0, the router only checks into the upward direction (towards the OrigNode) of the link.

If the upward direction of the link can fulfill the requirements indicated in the constraint option, and the router's rank would not exceed the MaxRank limit, the router joins the DODAG of the RREQ-Instance. The router that transmitted the received RREQ-DIO is selected as the preferred parent. Later, other RREQ-DIO messages might be received. How to maintain the parent set, select the preferred parent, and update the router's rank obeys the core RPL and the OFs defined in ROLL WG. In case that the constraint or the MaxRank limit is not fulfilled, the router MUST discard the received RREQ-DIO and MUST NOT join the DODAG.

Step 2:

Then the router checks if one of its addresses is included in one of the ART Options. If so, this router is one of the TargNodes. Otherwise, it is an intermediate router.

Step 3:

If the 'H' bit is set to 1, then the router (TargNode or intermediate) MUST build the upward route entry accordingly. The route entry MUST include at least the following items: Source Address, InstanceID, Destination Address, Next Hop, Lifetime, and Sequence Number. The Destination Address and the InstanceID can be respectively learned from the DODAGID and the RPLInstanceID of the RREQ-DIO, and the Source Address is copied from the ART Option. The next hop is the preferred parent. And the lifetime is set according to DODAG configuration and can be extended when the route is actually used. The sequence number represents the freshness of the route entry, and it is copied from the Orig SeqNo field of the RREQ option. A route entry with same source and destination address, same InstanceID, but stale sequence number, SHOULD be deleted.

If the 'H' bit is set to 0, an intermediate router MUST include the address of the interface receiving the RREQ-DIO into the address vector.

Step 4:

An intermediate router transmits a RREQ-DIO via link-local multicast. TargNode prepares a RREP-DIO.

6.2.2. Additional Processing for Multiple Targets

If the OrigNode tries to reach multiple TargNodes in a single RREQ-instance, one of the TargNodes can be an intermediate router to the others, therefore it SHOULD continue sending RREQ-DIO to reach other targets. In this case, before rebroadcasting the RREQ-DIO, a TargNode MUST delete the Target Option encapsulating its own address, so that downstream routers with higher ranks do not try to create a route to this TargetNode.

An intermediate router could receive several RREQ-DIOs from routers with lower ranks in the same RREQ-instance but have different lists of Target Options. When rebroadcasting the RREQ-DIO, the intersection of these lists SHOULD be included. For example, suppose two RREQ-DIOs are received with the same RPLInstance and OrigNode. Suppose further that the first RREQ has (T1, T2) as the targets, and the second one has (T2, T4) as targets. Then only T2 needs to be included in the generated RREQ-DIO. If the intersection is empty, it means that all the targets have been reached, and the router SHOULD NOT send out any RREQ-DIO. Any RREQ-DIO message with different ART Options coming from a router with higher rank is ignored.

6.3. Generating Route Reply (RREP) at TargNode

6.3.1. RREP-DIO for Symmetric route

If a RREQ-DIO arrives at TargNode with the 'S' bit set to 1, there is a symmetric route along which both directions can fulfill the requirements. Other RREQ-DIOs might later provide asymmetric upward routes (i.e. S=0). Selection between a qualified symmetric route and an asymmetric route that might have better performance is implementation-specific and out of scope. If the implementation uses the symmetric route, the TargNode MAY delay transmitting the RREP-DIO for duration RREP_WAIT_TIME to await a better symmetric route.

For a symmetric route, the RREP-DIO message is unicast to the next hop according to the accumulated address vector (H=0) or the route entry (H=1). Thus the DODAG in RREP-Instance does not need to be built. The RPLInstanceID in the RREP-Instance is paired as defined in Section 6.3.3. In case the 'H' bit is set to 0, the address vector received in the RREQ-DIO MUST be included in the RREP-DIO. The sequence number of the TargNode is updated to the maximum of its current sequence number and the Dest SeqNo in the ART option of the RREQ-DIO, using a mechanism similar to that used in [RFC3561]. This updated sequence number is then copied to the Dest SeqNo field of the ART option. The address of the OrigNode MUST be encapsulated in the ART Option and included in this RREP-DIO message.

6.3.2. RREP-DIO for Asymmetric Route

When a RREQ-DIO arrives at a TargNode with the 'S' bit set to 0, the TargNode MUST build a DODAG in the RREP-Instance rooted at itself in order to discover the downstream route from the OrigNode to the TargNode. The RREP-DIO message MUST be re-transmitted via link-local multicast until the OrigNode is reached or MaxRank is exceeded.

The settings of the fields in RREP option and ART option are the same as for the symmetric route, except for the 'S' bit.

6.3.3. RPLInstanceID Pairing

Since the RPLInstanceID is assigned locally (i.e., there is no coordination between routers in the assignment of RPLInstanceID), the tuple (OrigNode, TargNode, RPLInstanceID) is needed to uniquely identify a discovered route. The upper layer applications may have different requirements and they can initiate the route discoveries simultaneously. Thus between the same pair of OrigNode and TargNode, there can be multiple AODV-RPL instances. To avoid any mismatch, the RREQ-Instance and the RREP-Instance in the same route discovery MUST be paired somehow, e.g. using the RPLInstanceID.

When preparing the RREP-DIO, a TargNode could find the RPLInstanceID to be used for the RREP-Instance is already occupied by another RPL Instance from an earlier route discovery operation which is still active. In other words, it might happen that two distinct OrigNodes need routes to the same TargNode, and they happen to use the same RPLInstanceID for RREQ-Instance. In this case, the occupied RPLInstanceID MUST NOT be used again. Then the second RPLInstanceID MUST be shifted into another integer so that the two RREP-instances can be distinguished. In RREP option, the Shift field indicates the shift to be applied to original RPLInstanceID. When the new InstanceID after shifting exceeds 63, it rolls over starting at 0. For example, the original InstanceID is 60, and shifted by 6, the new InstanceID will be 2. Related operations can be found in Section 6.4.

6.4. Receiving and Forwarding Route Reply

Upon receiving a RREP-DIO, a router which does not belong to the RREQ-instance goes through the following steps:

Step 1:

If the 'S' bit is set to 1, the router proceeds to step 2.

If the 'S' bit of the RREP-DIO is set to 0, the router MUST check the downward direction of the link (towards the TargNode) over which the RREP-DIO is received. If the downward direction of the link can fulfill the requirements indicated in the constraint option, and the router's rank would not exceed the MaxRank limit, the router joins the DODAG of the RREP-Instance. The router that transmitted the received RREP-DIO is selected as the preferred parent. Afterwards, other RREP-DIO messages can be received. How to maintain the parent set, select the preferred parent, and update the router's rank obeys the core RPL and the OFs defined in ROLL WG.

If the constraints are not fulfilled, the router MUST NOT join the DODAG; the router MUST discard the RREQ-DIO, and does not execute the remaining steps in this section.

Step 2:

The router next checks if one of its addresses is included in the ART Option. If so, this router is the OrigNode of the route discovery. Otherwise, it is an intermediate router.

Step 3:

If the 'H' bit is set to 1, then the router (OrigNode or intermediate) MUST build a downward route entry. The route entry SHOULD include at least the following items: OrigNode Address, InstanceID, TargNode Address as destination, Next Hop, Lifetime and Sequence Number. For a symmetric route, the next hop in the route entry is the router from which the RREP-DIO is received. For an asymmetric route, the next hop is the preferred parent in the DODAG of RREQ-Instance. The InstanceID in the route entry MUST be the original RPLInstanceID (after subtracting the Shift field value). The source address is learned from the ART Option, and the destination address is learned from the DODAGID. The lifetime is set according to DODAG configuration and can be extended when the route is actually used. The sequence number represents the freshness of the route entry, and is copied from the Dest SeqNo field of the ART option of the RREP-DIO. A route entry with same source and destination address, same InstanceID, but stale sequence number, SHOULD be deleted.

If the 'H' bit is set to 0, for an asymmetric route, an intermediate router MUST include the address of the interface receiving the RREP-DIO into the address vector; for a symmetric route, there is nothing to do in this step.

Step 4:

If the receiver is the OrigNode, it can start transmitting the application data to TargNode along the path as provided in RREP-Instance, and processing for the RREP-DIO is complete. Otherwise, in case of an asymmetric route, the intermediate router transmits the RREP-DIO via link-local multicast. In case of a symmetric route, the RREP-DIO message is unicast to the next hop according to the address vector in the RREP-DIO (H=0) or the local route entry (H=1). The RPLInstanceID in the transmitted RREP-DIO is the same as the value in the received RREP-DIO. The local knowledge for the TargNode's sequence number SHOULD be updated.

7. Gratuitous RREP

In some cases, an Intermediate router that receives a RREQ-DIO message MAY transmit a "Gratuitous" RREP-DIO message back to OrigNode instead of continuing to multicast the RREQ-DIO towards TargNode. The intermediate router effectively builds the RREP-Instance on behalf of the actual TargNode. The 'G' bit of the RREP option is provided to distinguish the Gratuitous RREP-DIO (G=1) sent by the Intermediate node from the RREP-DIO sent by TargNode (G=0).

The gratuitous RREP-DIO can be sent out when an intermediate router R receives a RREQ-DIO for a TargNode T, and R happens to have a more

recent (larger destination sequence number) pair of downward and upward routes to T which also fulfill the requirements.

In case of source routing, the intermediate router R MUST unicast the received RREQ-DIO to TargNode T including the address vector between the OrigNode O and the router R. Thus T can have a complete upward route address vector from itself to O. Then R MUST send out the gratuitous RREP-DIO including the address vector from R to T.

In case of hop-by-hop routing, R MUST unicast the received RREQ-DIO hop-by-hop to T. The routers along the route SHOULD build new route entries with the related RPLInstanceID and DODAGID in the downward direction. Then T MUST unicast the RREP-DIO hop-by-hop to R, and the routers along the route SHOULD build new route entries in the upward direction. Upon receiving the unicast RREP-DIO, R sends the gratuitous RREP-DIO to the OrigNode as defined in Section 6.3.

8. Operation of Trickle Timer

The trickle timer operation to control RREQ-Instance/RREP-Instance multicast is similar to that in P2P-RPL [RFC6997].

9. IANA Considerations

9.1. New Mode of Operation: AODV-RPL

IANA is required to assign a new Mode of Operation, named "AODV-RPL" for Point-to-Point(P2P) hop-by-hop routing under the RPL registry. The value of TBD1 is assigned from the "Mode of Operation" space [RFC6550].

Value	Description	Reference
TBD1 (5)	AODV-RPL	This document

Figure 6: Mode of Operation

9.2. AODV-RPL Options: RREQ, RREP, and Target

Three entries are required for new AODV-RPL options "RREQ", "RREP" and "ART" with values of TBD2 (0x0A), TBD3 (0x0B) and TBD4 (0x0C) from the "RPL Control Message Options" space [RFC6550].

Value	Meaning	Reference
TBD2 (0x0A)	RREQ Option	This document
TBD3 (0x0B)	RREP Option	This document
TBD3 (0x0C)	ART Option	This document

Figure 7: AODV-RPL Options

10. Security Considerations

This document does not introduce additional security issues compared to base RPL. For general RPL security considerations, see [RFC6550].

11. Future Work

There has been some discussion about how to determine the initial state of a link after an AODV-RPL-based network has begun operation. The current draft operates as if the links are symmetric until additional metric information is collected. The means for making link metric information is considered out of scope for AODV-RPL. In the future, RREQ and RREP messages could be equipped with new fields for use in verifying link metrics. In particular, it is possible to identify unidirectional links; an RREQ received across a unidirectional link has to be dropped, since the destination node cannot make use of the received DODAG to route packets back to the source node that originated the route discovery operation. This is roughly the same as considering a unidirectional link to present an infinite cost metric that automatically disqualifies it for use in the reverse direction.

12. Contributors

Abdur Rashid Sangi
 Huaiyin Institute of Technology
 No.89 North Beijing Road, Qinghe District
 Huaian 223001
 P.R. China
 Email: sangi_bahrian@yahoo.com

13. References

13.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3561] Perkins, C., Belding-Royer, E., and S. Das, "Ad hoc On-Demand Distance Vector (AODV) Routing", RFC 3561, DOI 10.17487/RFC3561, July 2003, <<https://www.rfc-editor.org/info/rfc3561>>.
- [RFC5548] Dohler, M., Ed., Watteyne, T., Ed., Winter, T., Ed., and D. Barthel, Ed., "Routing Requirements for Urban Low-Power and Lossy Networks", RFC 5548, DOI 10.17487/RFC5548, May 2009, <<https://www.rfc-editor.org/info/rfc5548>>.
- [RFC5673] Pister, K., Ed., Thubert, P., Ed., Dwars, S., and T. Phinney, "Industrial Routing Requirements in Low-Power and Lossy Networks", RFC 5673, DOI 10.17487/RFC5673, October 2009, <<https://www.rfc-editor.org/info/rfc5673>>.
- [RFC5826] Brandt, A., Buron, J., and G. Porcu, "Home Automation Routing Requirements in Low-Power and Lossy Networks", RFC 5826, DOI 10.17487/RFC5826, April 2010, <<https://www.rfc-editor.org/info/rfc5826>>.
- [RFC5867] Martocci, J., Ed., De Mil, P., Riou, N., and W. Vermeulen, "Building Automation Routing Requirements in Low-Power and Lossy Networks", RFC 5867, DOI 10.17487/RFC5867, June 2010, <<https://www.rfc-editor.org/info/rfc5867>>.
- [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>>.
- [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>>.

[RFC6998] Goyal, M., Ed., Baccelli, E., Brandt, A., and J. Martocci, "A Mechanism to Measure the Routing Metrics along a Point-to-Point Route in a Low-Power and Lossy Network", RFC 6998, DOI 10.17487/RFC6998, August 2013, <<https://www.rfc-editor.org/info/rfc6998>>.

13.2. Informative References

[I-D.thubert-roll-asymlink] Thubert, P., "RPL adaptation for asymmetrical links", draft-thubert-roll-asymlink-02 (work in progress), December 2011.

[Perlman83] Perlman, R., "Fault-Tolerant Broadcast of Routing Information", December 1983.

[RFC6997] Goyal, M., Ed., Baccelli, E., Philipp, M., Brandt, A., and J. Martocci, "Reactive Discovery of Point-to-Point Routes in Low-Power and Lossy Networks", RFC 6997, DOI 10.17487/RFC6997, August 2013, <<https://www.rfc-editor.org/info/rfc6997>>.

Appendix A. Example: ETX/RSSI Values to select S bit

We have tested the combination of "RSSI(downstream)" and "ETX (upstream)" to determine whether the link is symmetric or asymmetric at the intermediate nodes. The example of how the ETX and RSSI values are used in conjunction is explained below:

Source----->NodeA----->NodeB----->Destination

Figure 8: Communication link from Source to Destination

RSSI at NodeA for NodeB	Expected ETX at NodeA for NodeB->NodeA
> -60	150
-70 to -60	192
-80 to -70	226
-90 to -80	662
-100 to -90	993

Table 1: Selection of 'S' bit based on Expected ETX value

We tested the operations in this specification by making the following experiment, using the above parameters. In our experiment,

a communication link is considered as symmetric if the ETX value of NodeA->NodeB and NodeB->NodeA (See Figure.8) are, say, within 1:3 ratio. This ratio should be taken as a notional metric for deciding link symmetric/asymmetric nature, and precise definition of the ratio is beyond the scope of the draft. In general, NodeA can only know the ETX value in the direction of NodeA -> NodeB but it has no direct way of knowing the value of ETX from NodeB->NodeA. Using physical testbed experiments and realistic wireless channel propagation models, one can determine a relationship between RSSI and ETX representable as an expression or a mapping table. Such a relationship in turn can be used to estimate ETX value at nodeA for link NodeB->NodeA from the received RSSI from NodeB. Whenever nodeA determines that the link towards the nodeB is bi-directional asymmetric then the "S" bit is set to "S=0". Later on, the link from NodeA to Destination is asymmetric with "S" bit remains to "0".

Appendix B. Changelog

B.1. Changes to version 02

- o Include the support for source routing.
- o Import some features from [RFC6997], e.g., choice between hop-by-hop and source routing, the "L" bit which determines the duration of residence in the DAG, MaxRank, etc.
- o Define new target option for AODV-RPL, including the Destination Sequence Number in it. Move the TargNode address in RREQ option and the OrigNode address in RREP option into ADOV-RPL Target Option.
- o Support route discovery for multiple targets in one RREQ-DIO.
- o New InstanceID pairing mechanism.

B.2. Changes to version 03

- o Updated RREP option format. Remove the 'T' bit in RREP option.
- o Using the same RPLInstanceID for RREQ and RREP, no need to update [RFC6550].
- o Explanation of Shift field in RREP.
- o Multiple target options handling during transmission.

B.3. Changes to version 04

- o Add description for sequence number operations.
- o Extend the residence duration L in the section 4.1.
- o Change AODV-RPL Target option to ART option.

Authors' Addresses

Satish Anamalamudi
SRM University-AP
Amaravati Campus
Amaravati, Andhra Pradesh 522 502
India

Email: satishnaidu80@gmail.com

Mingui Zhang
Huawei Technologies
No. 156 Beiqing Rd. Haidian District
Beijing 100095
China

Email: zhangmingui@huawei.com

Charles E. Perkins
Futurewei
2330 Central Expressway
Santa Clara 95050
Unites States

Email: charliep@computer.org

S.V.R Anand
Indian Institute of Science
Bangalore 560012
India

Email: anand@ece.iisc.ernet.in

Bing Liu
Huawei Technologies
No. 156 Beiqing Rd. Haidian District
Beijing 100095
China

Email: remy.liubing@huawei.com