

Internet Engineering Task Force (IETF)
Request for Comments: 7733
Category: Standards Track
ISSN: 2070-1721

A. Brandt
Sigma Designs
E. Baccelli
INRIA
R. Cragie
ARM Ltd.
P. van der Stok
Consultant
February 2016

Applicability Statement: The Use of the Routing Protocol
for Low-Power and Lossy Networks (RPL) Protocol Suite
in Home Automation and Building Control

Abstract

The purpose of this document is to provide guidance in the selection and use of protocols from the Routing Protocol for Low-Power and Lossy Networks (RPL) protocol suite to implement the features required for control in building and home environments.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7733>.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	4
1.1.	Relationship to Other Documents	5
1.2.	Terminology	6
1.3.	Required Reading	6
1.4.	Requirements That Are Out of Scope	6
2.	Deployment Scenario	6
2.1.	Network Topologies	7
2.2.	Traffic Characteristics	8
2.2.1.	General	9
2.2.2.	Source-Sink (SS) Communication Paradigm	10
2.2.3.	Publish-Subscribe (PS, or Pub/Sub) Communication Paradigm	10
2.2.4.	Peer-to-Peer (P2P) Communication Paradigm	10
2.2.5.	Peer-to-Multipeer (P2MP) Communication Paradigm	11
2.2.6.	Additional Considerations: Duocast and N-Cast	11
2.2.7.	RPL Applicability per Communication Paradigm	11
2.3.	Layer 2 Applicability	13
3.	Using RPL to Meet Functional Requirements	13
4.	RPL Profile	14
4.1.	RPL Features	14
4.1.1.	RPL Instances	15
4.1.2.	Storing vs. Non-Storing Mode	15
4.1.3.	DAO Policy	15
4.1.4.	Path Metrics	15
4.1.5.	Objective Function	16
4.1.6.	DODAG Repair	16
4.1.7.	Multicast	16
4.1.8.	Security	17
4.1.9.	P2P Communications	21
4.1.10.	IPv6 Address Configuration	21
4.2.	Layer 2 Features	21
4.2.1.	Specifics about Layer 2	21
4.2.2.	Services Provided at Layer 2	21
4.2.3.	IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Options Assumed	21
4.2.4.	Mesh Link Establishment (MLE) and Other Things	21
4.3.	Recommended Configuration Defaults and Ranges	21
4.3.1.	Trickle Parameters	22
4.3.2.	Other Parameters	22
5.	MPL Profile	23
5.1.	Recommended Configuration Defaults and Ranges	23
5.1.1.	Real-Time Optimizations	23
5.1.2.	Trickle Parameters	23
5.1.3.	Other Parameters	24
6.	Manageability Considerations	25

7. Security Considerations	25
7.1. Security Considerations during Initial Deployment	26
7.2. Security Considerations during Incremental Deployment	27
7.3. Security Considerations for P2P Implementations	27
7.4. MPL Routing	27
7.5. RPL Security Features	27
8. Other Related Protocols	28
9. References	28
9.1. Normative References	28
9.2. Informative References	32
Appendix A. RPL Shortcomings in Home and Building Deployments	35
A.1. Risk of Undesirable Long P2P Routes	35
A.1.1. Traffic Concentration at the Root	35
A.1.2. Excessive Battery Consumption in Source Nodes	35
A.2. Risk of Delayed Route Repair	35
A.2.1. Broken Service	36
Appendix B. Communication Failures	36
Acknowledgements	38
Authors' Addresses	38

1. Introduction

The primary purpose of this document is to give guidance in the use of the Routing Protocol for Low-Power and Lossy Networks (RPL) protocol suite in two application domains:

- o Home automation
- o Building automation

The guidance is based on the features required by the requirements documents "Home Automation Routing Requirements in Low-Power and Lossy Networks" [RFC5826] and "Building Automation Routing Requirements in Low-Power and Lossy Networks" [RFC5867], respectively. The Advanced Metering Infrastructure is also considered where appropriate. The applicability domains distinguish themselves in the way they are operated, their performance requirements, and the most likely network structures. An abstract set of distinct communication paradigms is then used to frame the applicability domains.

Home automation and building automation application domains share a substantial number of properties:

- o In both domains, the network can be disconnected from the ISP and must still continue to provide control to the occupants of the home or building. Routing needs to be possible independent of the existence of a border router.
- o Both domains are subject to unreliable links but require instant and very reliable reactions. This has an impact on routing because of timeliness and multipath routing.

The differences between the two application domains mostly appear in commissioning, maintenance, and the user interface, which do not typically affect routing. Therefore, the focus of this applicability document is on reliability, timeliness, and local routing.

It should be noted that adherence to the guidance in this document does not necessarily guarantee fully interoperable solutions in home automation networks and building control networks and that additional rigorous and managed programs will be needed to ensure interoperability.

1.1. Relationship to Other Documents

The Routing Over Low power and Lossy networks (ROLL) working group has specified a set of routing protocols for Low-Power and Lossy Networks (LLNs) [RFC6550]. This applicability text describes a subset of those protocols and the conditions under which the subset is appropriate, and it provides recommendations and requirements for the accompanying parameter value ranges.

In addition, [RFC6997] was written specifically as an extension to core RPL [RFC6550] and provides a solution for reactive discovery of point-to-point routes in LLNs. The present applicability document provides recommendations and requirements for the accompanying parameter value ranges.

[RFC7416] describes a common set of security threats. The applicability statements provided in Section 4.1.8.2.2 of this document complement [RFC7416] by describing preferred security settings and solutions within the applicability statement conditions. This applicability statement recommends lighter-weight security solutions appropriate for home and building environments and indicates why these solutions are appropriate.

1.2. 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].

Additionally, this document uses terminology from [RFC6997], [RFC7731], [RFC7102], [IEEE802.15.4], and [RFC6550].

1.3. Required Reading

Applicable requirements are described in [RFC5826] and [RFC5867]. A survey of the application field is described in [BC-Survey].

1.4. Requirements That Are Out of Scope

The considered network diameter is limited to a maximum diameter of 10 hops and a typical diameter of five hops; this captures the most common cases in home automation and building control networks.

This document does not consider the applicability of RPL-related specifications for urban and industrial applications [RFC5548] [RFC5673], which may exhibit significantly larger network diameters.

2. Deployment Scenario

The use of communications networks in buildings is essential to satisfy energy-saving regulations. Environmental conditions of buildings can be adapted to suit the comfort of the individuals present inside. Consequently, when no one is present, energy consumption can be reduced. Cost is the main driving factor behind deployment of wireless networking in buildings, especially in the case of retrofitting, where wireless connectivity saves costs incurred due to cabling and building modifications.

A typical home automation network is comprised of less than 100 nodes. Large building deployments may span 10,000 nodes, but to ensure uninterrupted service of light and air conditioning systems in individual zones of the building, nodes are typically organized in subnetworks. Each subnetwork in a building automation deployment typically contains tens to hundreds of nodes and, for critical operations, may operate independently from the other subnetworks.

The main purpose of the home or building automation network is to provide control over light and heating/cooling resources. User intervention via wall controllers is combined with movement, light and temperature sensors to enable automatic adjustment of window blinds, reduction of room temperature, etc. In general, the sensors

and actuators in a home or building typically have fixed physical locations and will remain in the same home or building automation network.

People expect an immediate and reliable response to their presence or actions. For example, a light not switching on after entry into a room may lead to confusion and a profound dissatisfaction with the lighting product.

Monitoring of functional correctness is at least as important as timely responses. Devices typically communicate their status regularly and send alarm messages to notify users or implementers that a malfunction of controlled equipment or a controlled network has occurred.

In building control, the infrastructure of the building management network can be shared with security/access, Internet Protocol (IP) telephony, and fire/alarm networks. This approach has a positive impact on the operation and cost of the network; however, care should be taken to ensure that the availability of the building management network does not become compromised beyond the ability of critical functions to perform adequately.

In homes, the entertainment network for audio/video streaming and gaming has different requirements, where the most important requirement is the need for high bandwidth not typically needed for home or building control. It is therefore expected that the entertainment network in the home will mostly be separate from the control network, as this will also lessen the impact on the availability of the control network.

2.1. Network Topologies

In general, the home automation network or building control network consists of wired and wireless subnetworks. In large buildings in particular, the wireless subnetworks can be connected to an IP backbone network where all infrastructure services (e.g., Domain Name System (DNS), automation servers) are located.

The wireless subnetwork can be configured according to any of the following topologies:

- o A stand-alone network of 10-100 nodes without a border router. This typically occurs in the home with a stand-alone control network, in low-cost buildings, and during installation of high-end control systems in buildings.
- o A connected network with one border router. This configuration will happen in homes where home appliances are controlled from outside the home, possibly via a smart phone, and in many building control scenarios.
- o A connected network with multiple border routers. This will typically happen in installations of large buildings.

Many of the nodes are battery powered and may be sleeping nodes that wake up according to clock signals or external events.

In a building control network, for a large installation with multiple border routers, subnetworks often overlap both geographically and from a wireless coverage perspective. Due to two purposes of the network -- (i) direct control and (ii) monitoring -- there may exist two types of routing topologies in a given subnetwork: (i) a tree-shaped collection of routes spanning from a central building controller via the border router, on to destination nodes in the subnetwork, and (ii) a flat, undirected collection of intra-network routes between functionally related nodes in the subnetwork.

The majority of nodes in home and building automation networks are typically Class 0 devices [RFC7228], such as individual wall switches. Only a few nodes (such as multi-purpose remote controls) are more expensive Class 1 devices, which can afford more memory capacity.

2.2. Traffic Characteristics

Traffic may enter the network originating from a central controller, or it may originate from an intra-network node. The majority of traffic is of a lightweight point-to-point control style, e.g., Put-Ack or Get-Response. There are, however, exceptions. Bulk data transfer is used for firmware updates and logging, where firmware updates enter the network and logs leave the network. Group communication is used for service discovery or to control groups of nodes, such as light fixtures.

Often, there is a direct physical relationship between a controlling sensor and the controlled equipment. For example, the temperature sensor and room controller are located in the same room, sharing the same climate conditions. Consequently, the bulk of senders and receivers are separated by a distance that allows one-hop direct path communication. A graph of the communication will show several fully connected subsets of nodes. However, due to interference, multipath fading, reflection, and other transmission mechanisms, the one-hop direct path may be temporarily disconnected. For reliability purposes, it is therefore essential that alternative n-hop communication routes exist for quick error recovery. (See Appendix B for motivation.)

Looking over time periods of a day, the networks are very lightly loaded. However, bursts of traffic can be generated by, for example, incessant pushing of the button of a remote control, the occurrence of a defect, and other unforeseen events. Under those conditions, the timeliness must nevertheless be maintained. Therefore, measures are necessary to remove any unnecessary traffic. Short routes are preferred. Long multi-hop routes via the border router should be avoided whenever possible.

Group communication is essential for lighting control. For example, once the presence of a person is detected in a given room, lighting control applies to that room only, and no other lights should be dimmed or switched on/off. In many cases, this means that a multicast message with a one-hop and two-hop radius would suffice to control the required lights. The same argument holds for Heating, Ventilating, and Air Conditioning (HVAC) and other climate-control devices. To reduce network load, it is advisable that messages to the lights in a room are not distributed any further in the mesh than necessary, based on intended receivers.

[Office-Light] provides an example of an office space, and [OccuSwitch] describes the current use of wireless lighting control products.

2.2.1. General

Although air conditioning and other environmental-control applications may accept response delays of tens of seconds or longer, alarm and light control applications may be regarded as soft real-time systems. A slight delay is acceptable, but the perceived quality of service degrades significantly if response times exceed 250 ms. If the light does not turn on at short notice, a user may activate the controls again, thus causing a sequence of commands such as Light{on,off,on,off,...} or Volume{up,up,up,up,up,...}. In

addition, the repetitive sending of commands creates an unnecessary loading of the network, which in turn increases the poor responsiveness of the network.

2.2.2. Source-Sink (SS) Communication Paradigm

This paradigm translates to many sources sending messages to the same sink, sometimes reachable via the border router. As such, Source-Sink (SS) traffic can be present in home and building networks. The traffic may be generated by environmental sensors (often present in a wireless subnetwork) that push periodic readings to a central server. The readings may be used for pure logging or, more often, processed to adjust light, heating, and ventilation. Alarm sensors may also generate SS-style traffic. The central server in a home automation network will be connected mostly to a wired network segment of the home network, although it is likely that cloud services will also be used. The central server in a building automation network may be connected to a backbone or placed outside the building.

With regard to message latency, most SS transmissions can tolerate worst-case delays measured in tens of seconds. Fire detectors, however, represent an exception; for example, special provisions with respect to the location of the fire detectors and smoke dampers need to be put in place to meet stringent delay requirements that are measured in seconds.

2.2.3. Publish-Subscribe (PS, or Pub/Sub) Communication Paradigm

This paradigm translates to a number of devices expressing their interest in a service provided by a server device. For example, a server device can be a sensor delivering temperature readings on the basis of delivery criteria, like changes in acquisition value or age of the latest acquisition. In building automation networks, this paradigm may be closely related to the SS paradigm, given that servers, which are connected to the backbone or outside the building, can subscribe to data collectors that are present at strategic places in the building automation network. The use of PS will probably differ significantly from installation to installation.

2.2.4. Peer-to-Peer (P2P) Communication Paradigm

This paradigm translates to a device transferring data to another device often connected to the same subnetwork. Peer-to-Peer (P2P) traffic is a common traffic type in home automation networks. Most building automation networks rely on P2P traffic as described in the next paragraph. Other building automation networks rely on P2P control traffic between controls and a local controller box for

advanced group control. A local controller box can be further connected to service control boxes, thus generating more SS or PS traffic.

P2P traffic is typically generated by remote controls and wall controllers that push Control Messages directly to light or heat sources. P2P traffic has a stringent requirement for low latency, since P2P traffic often carries application messages that are invoked by humans. As mentioned in Section 2.2.1, application messages should be delivered within a few hundred milliseconds, even when connections fail momentarily.

2.2.5. Peer-to-Multipeer (P2MP) Communication Paradigm

This paradigm translates to a device sending a message as many times as there are destination devices. Peer-to-Multipeer (P2MP) traffic is common in home and building automation networks. Often, a thermostat in a living room responds to temperature changes by sending temperature acquisitions to several fans and valves consecutively. This paradigm is also closely related to the PS paradigm in the case where a single server device has multiple subscribers.

2.2.6. Additional Considerations: Duocast and N-Cast

This paradigm translates to a device sending a message to many destinations in one network transfer invocation. Multicast is well suited for lighting where a presence sensor sends a presence message to a set of lighting devices. Multicast increases the probability that the message is delivered within strict time constraints. The recommended multicast algorithm (e.g., [RFC7731]) provides a mechanism for delivering messages to all intended destinations.

2.2.7. RPL Applicability per Communication Paradigm

In the case of the SS paradigm applied to a wireless subnetwork to a server reachable via a border router, the use of RPL [RFC6550] in non-storing mode is appropriate. Given the low resources of the devices, source routing will be used from the border router to the destination in the wireless subnetwork for messages generated outside the mesh network. No specific timing constraints are associated with the SS-type messages, so network repair does not violate the operational constraints. When no SS traffic takes place, it is good practice to load only RPL code that enables the P2P mode of operation [RFC6997] to reduce the code size and satisfy memory requirements.

To assure responsiveness, P2P-RPL [RFC6997] is required for all P2P and P2MP traffic taking place between nodes within a wireless subnetwork (excluding the border router). Source and destination devices are typically physically close, based on room layout. Consequently, most P2P and P2MP traffic is one-hop or two-hop traffic. Appendix A identifies shortcomings of using RPL for this type of communication; these shortcomings are counteracted through the use of P2P-RPL. Appendix B explains why reliability measures such as multipath routing are necessary even when one-hop communication dominates.

Examples of additional advantages of P2P-RPL for home and building automation networks are as follows:

- o Individual wall switches are typically inexpensive Class 0 devices [RFC7228] with extremely low memory capacities. Multi-purpose remote controls for use in a home environment typically have more memory, but such devices are asleep when there is no user activity. P2P-RPL reactive discovery allows a node to wake up and find new routes within a few seconds, while memory-constrained nodes only have to keep routes to relevant targets.
- o The reactive discovery features of P2P-RPL ensure that commands are normally delivered within the 250 ms time window. When connectivity needs to be restored, discovery is typically completed within seconds. In most cases, an alternative route (a route that was discovered earlier) will work and route rediscovery is not necessary.
- o Broadcast storms typically associated with route discovery for the Ad hoc On-Demand Distance Vector (AODV) [RFC3561] are less disruptive for P2P-RPL. P2P-RPL has a "Stop" bit, which is set by the target of a route discovery to notify all other nodes that no more Destination-Oriented Directed Acyclic Graph (DODAG) Information Object (DIO) messages should be forwarded for this temporary DAG. Something that looks like a broadcast storm may happen when no target is responding; however, in this case, the Trickle suppression mechanism kicks in, limiting the number of DIO forwards in dense networks.

Due to the limited memory of the majority of devices, P2P-RPL SHOULD be deployed with source routing in non-storing mode, as explained in Section 4.1.2.

Multicast with the Multicast Protocol for Low-Power and Lossy Networks (MPL) [RFC7731] is preferably deployed for N-cast over the wireless network. Configuration constraints that are necessary to meet reliability and timeliness with MPL are discussed in Section 4.1.7.

2.3. Layer 2 Applicability

This document applies to [IEEE802.15.4] and [G.9959], which are adapted to IPv6 by the adaptation layers [RFC4944] and [RFC7428]. Other Layer 2 technologies, accompanied by an "IP-over-Foo" specification, are also relevant, provided there is no frame size issue and there are link-layer acknowledgements.

The above-mentioned adaptation layers leverage on the compression capabilities of [RFC6554] and [RFC6282]. Header compression allows small IP packets to fit into a single Layer 2 frame, even when source routing is used. A network diameter limited to five hops helps to achieve this, even while using source routing.

Dropped packets are often experienced in the targeted environments. Internet Control Message Protocol (ICMP), User Datagram Protocol (UDP), and even Transmission Control Protocol (TCP) flows may benefit from link-layer unicast acknowledgements and retransmissions. Link-layer unicast acknowledgements SHOULD be enabled when [IEEE802.15.4] or [G.9959] is used with RPL and P2P-RPL.

3. Using RPL to Meet Functional Requirements

Several features required by [RFC5826] and [RFC5867] challenge the P2P paths provided by RPL. Appendix A reviews these challenges. In some cases, a node may need to spontaneously initiate the discovery of a path towards a desired destination that is neither the root of a DAG nor a destination originating Destination Advertisement Object (DAO) signaling. Furthermore, P2P paths provided by RPL are not satisfactory in all cases because they involve too many intermediate nodes before reaching the destination.

P2P-RPL [RFC6997] SHOULD be used in home automation and building control networks, as traffic of a point-to-point style is substantial and route repair needs to be completed within seconds. P2P-RPL provides a reactive mechanism for quick, efficient, and root-independent route discovery/repair. The use of P2P-RPL furthermore allows data traffic to avoid having to go through a central region around the root of the tree and drastically reduces path length [SOFT11] [INTEROP12]. These characteristics are desirable in home and building automation networks because they substantially decrease unnecessary network congestion around the root of the tree.

When more reliability is required, P2P-RPL enables the establishment of multiple independent paths. For one-hop destinations, this means that one one-hop communication and a second two-hop communication take place via a neighboring node. Such a pair of redundant communication paths can be achieved by using MPL, where the source is an MPL Forwarder while a second MPL Forwarder is one hop away from both the source and the destination node. When the source multicasts the message, it may be received by both the destination and the second MPL Forwarder. The second MPL Forwarder forwards the message to the destination, thus providing two routes from sender to destination.

To provide more reliability with multiple paths, P2P-RPL can maintain two independent P2P source routes per destination, at the source. Good practice is to use the paths alternately to assess their existence. When one P2P path has failed (possibly only temporarily), as described in Appendix B, the alternative P2P path can be used without discarding the failed path. The failed P2P path, unless proven to work again, can be safely discarded after a timeout (typically 15 minutes). A new route discovery is done when the number of P2P paths is exhausted due to persistent link failures.

4. RPL Profile

P2P-RPL SHOULD be used in home automation and building control networks. Its reactive discovery allows for low application response times, even when on-the-fly route repair is needed. Non-storing mode SHOULD be used to reduce memory consumption in repeaters with constrained memory when source routing is used.

4.1. RPL Features

An important constraint on the application of RPL is the presence of sleeping nodes.

For example, in a stand-alone network, the master node (or coordinator) providing the logical Layer 2 identifier and unique node identifiers to connected nodes may be a remote control that returns to sleep once new nodes have been added. Due to the absence of the border router, there may be no global routable prefixes at all. Likewise, there may be no authoritative always-on root node, since there is no border router to host this function.

In a network with a border router and many sleeping nodes, there may be battery-powered sensors and wall controllers configured to contact other nodes in response to events and then return to sleep. Such nodes may never detect the announcement of new prefixes via multicast.

In each of the above-mentioned constrained deployments, a link-layer node (e.g., coordinator or master) SHOULD assume the role of an authoritative root node, transmitting unicast Router Advertisement (RA) messages with a Unique Local Address (ULA) prefix information option to nodes during the joining process to prepare the nodes for a later operational phase, where a border router is added.

A border router SHOULD be designed to be aware of sleeping nodes in order to support the distribution of updated global prefixes to such sleeping nodes.

4.1.1. RPL Instances

When operating P2P-RPL on a stand-alone basis, there is no authoritative root node maintaining a permanent RPL DODAG. A node MUST be able to join at least one RPL Instance, as a new, temporary instance is created during each P2P-RPL route discovery operation. A node MAY be designed to join multiple RPL Instances.

4.1.2. Storing vs. Non-Storing Mode

Non-storing mode MUST be used to cope with the extremely constrained memory of a majority of nodes in the network (such as individual light switches).

4.1.3. DAO Policy

Nodes send DAO messages to establish downward paths from the root to themselves. In order to minimize the power consumption overhead associated with path discovery, DAO messages are not acknowledged in networks composed of battery-operated field devices. The DAO messages build up a source route because the nodes MUST be in non-storing mode.

If devices in LLNs participate in multiple RPL Instances and DODAGs, both the RPLInstance ID and the DODAGID SHOULD be included in the DAO.

4.1.4. Path Metrics

Expected Transmission Count (ETX) is the RECOMMENDED metric. [RFC6551] provides other options.

Packets from asymmetric and/or unstable links SHOULD be deleted at Layer 2.

4.1.5. Objective Function

Objective Function Zero (OF0) [RFC6552] MUST be the Objective Function. Other Objective Functions MAY be used when dictated by circumstances.

4.1.6. DODAG Repair

Since P2P-RPL only creates DODAGs on a temporary basis during route repair or route discovery, there is no need to repair DODAGs.

For SS traffic, local repair is sufficient. The accompanying process is known as "poisoning" and is described in Section 8.2.2.5 of [RFC6550]. Given that the majority of nodes in the building do not physically move around, creating new DODAGs should not happen frequently.

4.1.7. Multicast

Commercial lighting deployments may have a need for multicast to distribute commands to a group of lights in a timely fashion. Several mechanisms exist for achieving such functionality; [RFC7731] is the RECOMMENDED protocol for home and building deployments. This section relies heavily on the conclusions of [RT-MPL].

At reception of a packet, the MPL Forwarder starts a series of consecutive Trickle timer intervals, where the first interval has a minimum size of I_{min} . Each consecutive interval is twice as long as the former, with a maximum value of I_{max} . There is a maximum number of intervals given by $max_expiration$. For each interval of length I , a time t is randomly chosen in the period $[I/2, I]$. For a given packet, p , MPL counts the number of times it receives p during the period $[0, t]$ in a counter c . At time t , MPL rebroadcasts p when $c < k$, where k is a predefined constant with a value $k > 0$.

The density of forwarders and the frequency of message generation are important aspects to obtain timeliness during control operations. A high frequency of message generation can be expected when a remote-control button is incessantly pressed or when alarm situations arise.

Guaranteeing timeliness is intimately related to the density of the MPL routers. In ideal circumstances, the message is propagated as a single wave through the network, such that the maximum delay is related to the number of hops times the smallest repetition interval of MPL. Each forwarder that receives the message passes the message on to the next hop by repeating the message. When several copies of a message reach the forwarder, it is specified that the copy need not

be repeated. Repetition of the message can be inhibited by a small value of k . To assure timeliness, the chosen value of k should be high enough to make sure that messages are repeated at the first arrival of the message in the forwarder. However, a network that is too dense leads to a saturation of the medium that can only be prevented by selecting a low value of k . Consequently, timeliness is assured by choosing a relatively high value of k but assuring at the same time a low enough density of forwarders to reduce the risk of medium saturation. Depending on the reliability of the network links, it is advisable to configure the density of the network such that at least two forwarders per hop repeat messages to the same set of destinations.

There are no rules about selecting forwarders for MPL. In buildings with central management tools, the forwarders can be selected, but at the time of this writing it is not possible to automatically configure the forwarder topology in the home.

4.1.8. Security

RPL MAY use unsecured RPL messages to reduce message size. If there is a single node that uses unsecured RPL messages, link-layer security MUST be used on all nodes. Therefore, all RPL messages MUST be secured using:

- o RPL message security, or
- o Link-layer security, or
- o Both RPL message security and link-layer security

A symmetric key is used to secure a RPL message using either RPL message security or link-layer security. The symmetric key MUST be distributed or established in a secure fashion. There may be more than one symmetric key in use by any node at any one time. The same symmetric key MUST NOT be used for both RPL message security and link-layer security between two peer nodes.

4.1.8.1. Symmetric Key Distribution

The scope of symmetric key distribution MUST be no greater than the network itself, i.e., a group key. This document describes what needs to be implemented to meet this requirement. The scope of symmetric key distribution MAY be smaller than the network -- for example:

- o A pairwise symmetric key between two peers.
- o A group key shared between a subset of nodes in the network.

4.1.8.2. Symmetric Key Distribution Mechanism

The authentication mechanism as described in Section 6.9 of [ZigBeeIP] SHALL be used to securely distribute a network-wide symmetric key.

The purpose of the authentication procedure is to provide mutual authentication resulting in:

- o Preventing untrusted nodes without appropriate credentials from joining the trusted network.
- o Preventing trusted nodes with appropriate credentials from joining an untrusted network.

There is an Authentication Server, which is responsible for authenticating the nodes on the network. If the authentication is successful, the Authentication Server sends the network security material to the joining node through the Protocol for Carrying Authentication for Network Access (PANA) [RFC5191] [RFC6345]. The joining node becomes a full participating node in the network and is able to apply Layer 2 security to RPL messages using the distributed network key.

The joining node does not initially have access to the network security material. Therefore, it is not able to apply Layer 2 security to the packets exchanged during the authentication process. The enforcement point rules at the edge of the network ensure that the packets involved in PANA authentication are processed even though they are unsecured at the Medium Access Control (MAC) layer. The rules also ensure that any other incoming traffic that is not secured at the MAC layer is discarded and is not forwarded.

4.1.8.2.1. Authentication Stack

Authentication can be viewed as a protocol stack as a layer encapsulates the layers above it.

- o Transport Layer Security (TLS) [RFC5246] MUST be used at the highest layer of the authentication stack and carries the authentication exchange. There is one cipher suite based on a Pre-Shared Key (PSK) [RFC6655] and one cipher suite based on Elliptic Curve Cryptography (ECC) [RFC7251].
- o Extensible Authentication Protocol-TLS (EAP-TLS) [RFC5216] MUST be used at the next layer to carry the TLS records for the authentication protocol.
- o EAP [RFC3748] MUST be used to provide the mechanisms for mutual authentication. EAP requires a way to transport EAP packets between the joining node and the node on which the Authentication Server resides. These nodes are not necessarily in radio range of each other, so it is necessary to have multi-hop support in the EAP transport method. PANA [RFC5191] [RFC6345], which operates over UDP, MUST be used for this purpose. [RFC3748] specifies the derivation of a session key using the EAP key hierarchy; only the EAP Master Session Key shall be derived, as [RFC5191] specifies that it is used to set up keys for PANA authentication and encryption.
- o PANA [RFC5191] and a PANA relay [RFC6345] MUST be used at the next layer:
 - * The joining node MUST act as the PANA Client (PaC).
 - * The parent edge router node MUST act as a PANA Relay Element (PRE) according to [RFC6345], unless it is also the Authentication Server. All routers at the edge of the network MUST be capable of functioning in the PRE role.
 - * The Authentication Server node MUST act as the PANA Authentication Agent (PAA). The Authentication Server MUST be able to handle packets relayed according to [RFC6345].

This network authentication process uses link-local IPv6 addresses for transport between the new node and its parent. If the parent is not the Authentication Server, it MUST then relay packets from the joining node to the Authentication Server and vice versa, using the PANA relay mechanism [RFC6345]. The joining node MUST use a link-local address based on its EUI-64 as the source address for initial PANA authentication message exchanges.

4.1.8.2.2. Applicability Statements

The following applicability statements describe the relationship between the various specifications.

4.1.8.2.2.1. Applicability Statement for PSK TLS

[RFC6655] contains Authenticated Encryption with Associated Data (AEAD) TLS cipher suites that are very similar to [RFC5487], whose AEAD part is detailed in [RFC5116]. [RFC5487] references both [RFC5288] and the original PSK cipher suite document [RFC4279], which references RFC 2246, which was eventually replaced by [RFC5246], which defines the TLS 1.2 messages.

4.1.8.2.2.2. Applicability Statement for ECC TLS

[RFC7251] contains AEAD TLS cipher suites that are very similar to [RFC5289], whose AEAD part is detailed in [RFC5116]. [RFC5289] references the original ECC cipher suite document [RFC4492], which references RFC 2246, which was eventually replaced by [RFC5246], which defines the TLS 1.2 messages.

4.1.8.2.2.3. Applicability Statement for EAP-TLS and PANA

[RFC5216] specifies how [RFC3748] is used to package [RFC5246] TLS records into EAP packets. [RFC5191] provides transportation for the EAP packets and the network-wide key carried in an encrypted Attribute-Value Pair (AVP) as specified in [RFC6786]. The proposed Pseudorandom Function (PRF) and authentication (AUTH) hashes based on SHA-256 are represented as specified in [RFC7296] and detailed in [RFC4868].

4.1.8.2.3. Security Using RPL Message Security

If RPL is used with secured messages [RFC6550], the following RPL security parameter values SHOULD be used:

- o Counter is Time (T) flag = 0: Do not use the timestamp in the Counter field. Counters based on timestamps are typically more applicable to industrial networks, where strict timing synchronization between nodes is often implemented. Home and building networks typically do not implement such strict timing synchronization; therefore, a monotonically increasing counter is more appropriate.
- o Algorithm = 0: Use Counter with the Cipher Block Chaining Message Authentication Code (CBC-MAC Mode) (CCM) with AES-128. This is the only assigned mode at present.

- o Key Identifier Mode (KIM) = 10: Use a group key, Key Source present, and Key Index present. Given the relatively confined perimeter of a home or building network, a group key is usually sufficient to protect RPL messages sent between nodes. The use of the Key Source field allows multiple group keys to be used within the network.
- o Security Level (LVL) = 0: Use MAC-32. This is recommended, as integrity protection for RPL messages is the basic requirement. Encryption is unlikely to be necessary, given the relatively non-confidential nature of RPL message payloads.

4.1.9. P2P Communications

[RFC6997] MUST be used to accommodate P2P traffic, which is typically substantial in home and building automation networks.

4.1.10. IPv6 Address Configuration

Assigned IP addresses MUST be routable and unique within the routing domain [RFC5889].

4.2. Layer 2 Features

No particular requirements exist for Layer 2, except for those cited in the "IP-over-Foo" RFCs (see Section 2.3).

4.2.1. Specifics about Layer 2

Not applicable

4.2.2. Services Provided at Layer 2

Not applicable

4.2.3. IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Options Assumed

Not applicable

4.2.4. Mesh Link Establishment (MLE) and Other Things

Not applicable

4.3. Recommended Configuration Defaults and Ranges

The following sections describe the recommended parameter values for P2P-RPL and Trickle.

4.3.1. Trickle Parameters

Trickle is used to distribute network parameter values to all nodes without stringent time restrictions. The recommended Trickle parameter values are:

- o DIOIntervalMin 4, which translates to 16 ms
- o DIOIntervalDoublings 14
- o DIORedundancyConstant 1

When a node sends a changed DIO, this is an inconsistency and forces the receiving node to respond within I_{min} . So, when something happens that affects the DIO, the change is ideally communicated to a node that is n hops away, within n times I_{min} . Often, depending on the node density, packets are lost or are not sent, leading to larger delays.

In general, we can expect DIO changes to propagate within 1 to 3 seconds within the envisaged networks.

When nothing happens, the DIO sending interval increases to 4.37 minutes, thus drastically reducing the network load. When a node does not receive DIO messages for more than 10 minutes, it can safely conclude that the connection with other nodes has been lost.

4.3.2. Other Parameters

This section discusses the P2P-RPL parameters.

P2P-RPL [RFC6997] provides the features requested by [RFC5826] and [RFC5867]. P2P-RPL uses a subset of the frame formats and features defined for RPL [RFC6550] but may be combined with RPL frame flows in advanced deployments.

The recommended parameter values for P2P-RPL are:

- o MinHopRankIncrease 1
- o MaxRankIncrease 0
- o MaxRank 6
- o Objective Function: OF0

5. MPL Profile

MPL is used to distribute values to groups of devices. Using MPL, based on the Trickle algorithm, timeliness should also be guaranteed. A deadline of 200 ms needs to be met when human action is followed by an immediately observable action such as switching on lights. The deadline needs to be met in a building where the number of hops from seed to destination varies between 1 and 10.

5.1. Recommended Configuration Defaults and Ranges

5.1.1. Real-Time Optimizations

When the network is heavily loaded, MAC delays contribute significantly to the end-to-end delays when MPL intervals between 10 and 100 ms are used to meet the 200 ms deadline. It is possible to set the number of buffers in the MAC to 1 and set the number of back-off repetitions to 1. The number of MPL repetitions compensates for the reduced probability of transmission per MAC invocation [RT-MPL].

In addition, end-to-end delays and message losses are reduced by adding a real-time layer between MPL and MAC to throw away the earliest messages (exploiting the MPL message numbering) and favor the most recent ones.

5.1.2. Trickle Parameters

This section proposes values for the Trickle parameters used by MPL for the distribution of packets that need to meet a 200 ms deadline. The probability of meeting the deadline is increased by (1) choosing a small I_{min} value, (2) reducing the number of MPL intervals, thus reducing the load, and (3) reducing the number of MPL Forwarders to also reduce the load.

The consequence of this approach is that the value of k can be larger than 1 because network load reduction is already guaranteed by the network configuration.

Under the condition that the density of MPL repeaters can be limited, it is possible to choose low MPL repeat intervals (I_{min}) connected to k values such that $k > 1$. The minimum value of k is related to:

- o The value of I_{min} . The length of I_{min} determines the number of packets that can be received within the listening period of I_{min} .
- o The number of repeaters receiving the broadcast message from the same forwarder or seed. These repeaters repeat within the same I_{min} interval, thus increasing the c counter.

Within the first MPL interval, a limited number, q , of messages can be transmitted. Assuming a 3 ms transmission interval, q is given by $q = I_{min} / 3$. Assuming that at most q message copies can reach a given forwarder within the first repeat interval of length I_{min} , the related MPL parameter values are suggested in the following sections.

5.1.2.1. I_{min}

The recommended value is $I_{min} = 10$ to 50 ms.

When the chosen I_{min} value is much smaller, the interference between the copies leads to significant losses, given that q is much smaller than the number of repeated packets. With much larger intervals, the probability that the deadline will be met decreases with increasing hop count.

5.1.2.2. I_{max}

The recommended value is $I_{max} = 100$ to 400 ms.

The value of I_{max} is less important than the value of $max_expiration$. Given an I_{min} value of 10 ms, the third MPL interval has a value of $10 * 2 * 2 = 40$ ms. When I_{min} has a value of 40 ms, the third interval has a value of 160 ms. Given that more than three intervals are unnecessary, I_{max} does not contribute much to performance.

5.1.3. Other Parameters

Other parameters are the k parameter and the $max_expiration$ parameter.

$k > q$ (see condition above). Under this condition, and for a small I_{min} value, a value of $k = 2$ or $k = 3$ is usually sufficient to minimize the losses of packets in the first repeat interval.

$max_expiration = 2 - 4$. Higher values lead to more network load while generating copies that will probably not meet their deadline.

6. Manageability Considerations

At this time, it is not clear how homenets will be managed. Consequently, it is not clear which tools will be used and which parameters must be visible for management.

In building control, management is mandatory. It is expected that installations will be managed using the set of currently available tools (including IETF tools like Management Information Base (MIB) modules, Network Configuration Protocol (NETCONF) modules, Dynamic Host Configuration Protocol (DHCP), and others), with large differences between the ways an installation is managed.

7. Security Considerations

This section refers to the security considerations of [RFC6997], [RFC6550], and [RFC7731], as well as some attacks and countermeasures as discussed in Sections 6 and 7, respectively, of [RFC7416].

Communications network security is based on providing integrity protection and encryption to messages. This can be applied at various layers in the network protocol stack, based on using various credentials and a network identity.

The credentials that are relevant in the case of RPL are (i) the credential used at the link layer in the case where link-layer security is applied (see Section 7.1) or (ii) the credential used for securing RPL messages. In both cases, the assumption is that the credential is a shared key. Therefore, there **MUST** be a mechanism in place that allows secure distribution of a shared key and configuration of a network identity. Both **MAY** be done using (i) pre-installation using an out-of-band method, (ii) secure delivery when a device is introduced into the network, or (iii) secure delivery by a trusted neighboring device, as described in Section 4.1.8.1. The shared key **MUST** be stored in a secure fashion that will make it difficult to be read by an unauthorized party.

This document mandates that a Layer 2 mechanism be used during initial and incremental deployment. Please see the following sections.

7.1. Security Considerations during Initial Deployment

Wireless mesh networks are typically secured at the link layer in order to prevent unauthorized parties from accessing the information exchanged over the links. It is a basic practice to create a network of nodes that share the same keys for link-layer security and exclude nodes sending unsecured messages. With per-message data origin authentication, it is possible to prevent unauthorized nodes from joining the mesh.

At initial deployment, the network is secured by consecutively securing nodes at the link layer, thus building a network of secured nodes. Section 4.1.8.2 describes a mechanism for building a network of secured nodes.

This document does not specify a multicast security solution. Networks deployed with this specification will depend upon Layer 2 security to prevent outsiders from sending multicast traffic. It is recognized that this does not protect this control traffic from impersonation by already-trusted devices. This is an area for a future specification.

For building control, an installer will use an installation tool that establishes a secure communication path with the joining node. It is recognized that the recommendations for initial deployment as discussed in this section do not cover all building requirements, such as selecting -- independent of network topology -- the node to be secured.

It is expected that a set of protocol combinations will evolve within currently existing alliances of building control manufacturers. Each set satisfies the installation requirements of installers, operators, and manufacturers of building control networks in a given installation context, e.g., lighting deployment in offices, HVAC installation, incremental addition of equipment in homes, and others.

In the home, nodes can be visually inspected by the home owner. Also, a simple procedure, e.g., pushing buttons simultaneously on an already-secured device and an unsecured joining device, is usually sufficient to ensure that the unsecured joining device is authenticated securely, configured securely, and paired appropriately.

This recommendation is in line with the countermeasures described in Section 7.1 of [RFC7416].

7.2. Security Considerations during Incremental Deployment

Once a network is operational, new nodes need to be added, or nodes fail and need to be replaced. When a new node needs to be added to the network, the new node is added to the network via an assisting node in the manner described in Section 7.1.

On detection of a compromised node, all trusted nodes need to have their symmetric keys that are known to be shared with the compromised node rekeyed, and the trusted network is built up as described in Section 7.1.

7.3. Security Considerations for P2P Implementations

Refer to the security considerations of [RFC6997].

7.4. MPL Routing

The routing of MPL is determined by the enabling of the interfaces for specified multicast addresses. The specification of these addresses can be done via a Constrained Application Protocol (CoAP) application as specified in [RFC7390]. An alternative is the creation of an MPL MIB and the use of the Simple Network Management Protocol (SNMPv3) [RFC3411] or equivalent techniques to specify the multicast addresses in the MIB. For secure dissemination of MPL packets, Layer 2 security SHOULD be used, and the configuration of multicast addresses as described in this section MUST be secure.

7.5. RPL Security Features

This section refers to the structure of Section 8 ("RPL Security Features") of [RFC7416]. [RFC7416] provides a thorough analysis of security threats and proposed countermeasures relevant to RPL and MPL.

In accordance with Section 8.1 ("Confidentiality Features") of [RFC7416], RPL message security implements payload protection, as explained in Section 7 of this document. The attributes for key length and lifetime of the keys depend on operational conditions, maintenance, and installation procedures.

Sections 7.1 and 7.2 of this document recommend link-layer security to assure integrity in accordance with Section 8.2 ("Integrity Features") of [RFC7416].

The provision of multiple paths recommended in Section 8.3 ("Availability Features") of [RFC7416] is also recommended from a reliability point of view. Randomly choosing paths MAY be supported.

A mechanism for key management, as discussed in Section 8.4 ("Key Management") of [RFC7416], is provided in Section 4.1.8.2 of this document.

8. Other Related Protocols

Application and transport protocols used in home and building automation domains are expected to mostly consist of CoAP over UDP, or equivalents. Typically, UDP is used for IP transport to keep down the application response time and bandwidth overhead. CoAP is used at the application layer to reduce memory footprint and bandwidth requirements.

9. References

9.1. 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, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3748] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowitz, Ed., "Extensible Authentication Protocol (EAP)", RFC 3748, DOI 10.17487/RFC3748, June 2004, <<http://www.rfc-editor.org/info/rfc3748>>.
- [RFC4279] Eronen, P., Ed., and H. Tschofenig, Ed., "Pre-Shared Key Ciphersuites for Transport Layer Security (TLS)", RFC 4279, DOI 10.17487/RFC4279, December 2005, <<http://www.rfc-editor.org/info/rfc4279>>.
- [RFC4492] Blake-Wilson, S., Bolyard, N., Gupta, V., Hawk, C., and B. Moeller, "Elliptic Curve Cryptography (ECC) Cipher Suites for Transport Layer Security (TLS)", RFC 4492, DOI 10.17487/RFC4492, May 2006, <<http://www.rfc-editor.org/info/rfc4492>>.
- [RFC4868] Kelly, S. and S. Frankel, "Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec", RFC 4868, DOI 10.17487/RFC4868, May 2007, <<http://www.rfc-editor.org/info/rfc4868>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<http://www.rfc-editor.org/info/rfc4944>>.

- [RFC5116] McGrew, D., "An Interface and Algorithms for Authenticated Encryption", RFC 5116, DOI 10.17487/RFC5116, January 2008, <<http://www.rfc-editor.org/info/rfc5116>>.
- [RFC5191] Forsberg, D., Ohba, Y., Ed., Patil, B., Tschofenig, H., and A. Yegin, "Protocol for Carrying Authentication for Network Access (PANA)", RFC 5191, DOI 10.17487/RFC5191, May 2008, <<http://www.rfc-editor.org/info/rfc5191>>.
- [RFC5216] Simon, D., Aboba, B., and R. Hurst, "The EAP-TLS Authentication Protocol", RFC 5216, DOI 10.17487/RFC5216, March 2008, <<http://www.rfc-editor.org/info/rfc5216>>.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<http://www.rfc-editor.org/info/rfc5246>>.
- [RFC5288] Salowey, J., Choudhury, A., and D. McGrew, "AES Galois Counter Mode (GCM) Cipher Suites for TLS", RFC 5288, DOI 10.17487/RFC5288, August 2008, <<http://www.rfc-editor.org/info/rfc5288>>.
- [RFC5289] Rescorla, E., "TLS Elliptic Curve Cipher Suites with SHA-256/384 and AES Galois Counter Mode (GCM)", RFC 5289, DOI 10.17487/RFC5289, August 2008, <<http://www.rfc-editor.org/info/rfc5289>>.
- [RFC5487] Badra, M., "Pre-Shared Key Cipher Suites for TLS with SHA-256/384 and AES Galois Counter Mode", RFC 5487, DOI 10.17487/RFC5487, March 2009, <<http://www.rfc-editor.org/info/rfc5487>>.
- [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, <<http://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, <<http://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, <<http://www.rfc-editor.org/info/rfc5826>>.

- [RFC5867] Martocci, J., Ed., De Mil, P., Riou, N., and W. Vermeylen, "Building Automation Routing Requirements in Low-Power and Lossy Networks", RFC 5867, DOI 10.17487/RFC5867, June 2010, <<http://www.rfc-editor.org/info/rfc5867>>.
- [RFC6282] Hui, J., Ed., and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<http://www.rfc-editor.org/info/rfc6282>>.
- [RFC6345] Duffy, P., Chakrabarti, S., Cragie, R., Ohba, Y., Ed., and A. Yegin, "Protocol for Carrying Authentication for Network Access (PANA) Relay Element", RFC 6345, DOI 10.17487/RFC6345, August 2011, <<http://www.rfc-editor.org/info/rfc6345>>.
- [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, <<http://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, <<http://www.rfc-editor.org/info/rfc6551>>.
- [RFC6554] Hui, J., Vasseur, JP., Culler, D., and V. Manral, "An IPv6 Routing Header for Source Routes with the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6554, DOI 10.17487/RFC6554, March 2012, <<http://www.rfc-editor.org/info/rfc6554>>.
- [RFC6655] McGrew, D. and D. Bailey, "AES-CCM Cipher Suites for Transport Layer Security (TLS)", RFC 6655, DOI 10.17487/RFC6655, July 2012, <<http://www.rfc-editor.org/info/rfc6655>>.
- [RFC6786] Yegin, A. and R. Cragie, "Encrypting the Protocol for Carrying Authentication for Network Access (PANA) Attribute-Value Pairs", RFC 6786, DOI 10.17487/RFC6786, November 2012, <<http://www.rfc-editor.org/info/rfc6786>>.

- [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, <<http://www.rfc-editor.org/info/rfc6997>>.
- [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, <<http://www.rfc-editor.org/info/rfc6998>>.
- [RFC7102] Vasseur, JP., "Terms Used in Routing for Low-Power and Lossy Networks", RFC 7102, DOI 10.17487/RFC7102, January 2014, <<http://www.rfc-editor.org/info/rfc7102>>.
- [RFC7251] McGrew, D., Bailey, D., Campagna, M., and R. Dugal, "AES-CCM Elliptic Curve Cryptography (ECC) Cipher Suites for TLS", RFC 7251, DOI 10.17487/RFC7251, June 2014, <<http://www.rfc-editor.org/info/rfc7251>>.
- [RFC7296] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)", STD 79, RFC 7296, DOI 10.17487/RFC7296, October 2014, <<http://www.rfc-editor.org/info/rfc7296>>.
- [RFC7416] Tsao, T., Alexander, R., Dohler, M., Daza, V., Lozano, A., and M. Richardson, Ed., "A Security Threat Analysis for the Routing Protocol for Low-Power and Lossy Networks (RPLs)", RFC 7416, DOI 10.17487/RFC7416, January 2015, <<http://www.rfc-editor.org/info/rfc7416>>.
- [RFC7731] Hui, J. and R. Kelsey, "Multicast Protocol for Low-Power and Lossy Networks (MPL)", RFC 7731, DOI 10.17487/RFC7731, February 2016, <<http://www.rfc-editor.org/info/rfc7731>>.

- [IEEE802.15.4] IEEE, "IEEE Standard for Local and metropolitan area networks--Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs)", IEEE 802.15.4, DOI 10.1109/ieeestd.2011.6012487, <<http://ieeexplore.ieee.org/servlet/opac?punumber=6012485>>.
- [G.9959] International Telecommunication Union, "Short range narrow-band digital radiocommunication transceivers - PHY, MAC, SAR and LLC layer specifications", ITU-T Recommendation G.9959, January 2015, <<http://www.itu.int/rec/T-REC-G.9959>>.

9.2. Informative References

- [RFC3411] Harrington, D., Presuhn, R., and B. Wijnen, "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks", STD 62, RFC 3411, DOI 10.17487/RFC3411, December 2002, <<http://www.rfc-editor.org/info/rfc3411>>.
- [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, <<http://www.rfc-editor.org/info/rfc3561>>.
- [RFC5889] Baccelli, E., Ed., and M. Townsley, Ed., "IP Addressing Model in Ad Hoc Networks", RFC 5889, DOI 10.17487/RFC5889, September 2010, <<http://www.rfc-editor.org/info/rfc5889>>.
- [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, <<http://www.rfc-editor.org/info/rfc6552>>.
- [RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <<http://www.rfc-editor.org/info/rfc7228>>.
- [RFC7390] Rahman, A., Ed., and E. Dijk, Ed., "Group Communication for the Constrained Application Protocol (CoAP)", RFC 7390, DOI 10.17487/RFC7390, October 2014, <<http://www.rfc-editor.org/info/rfc7390>>.

- [RFC7428] Brandt, A. and J. Buron, "Transmission of IPv6 Packets over ITU-T G.9959 Networks", RFC 7428, DOI 10.17487/RFC7428, February 2015, <<http://www.rfc-editor.org/info/rfc7428>>.
- [SOFT11] Baccelli, E., Philipp, M., and M. Goyal, "The P2P-RPL Routing Protocol for IPv6 Sensor Networks: Testbed Experiments", Proceedings of the 19th Annual Conference on Software Telecommunications and Computer Networks, Split, Croatia, September 2011.
- [INTEROP12] Philipp, M., Baccelli, E., Brandt, A., Valev, H., and J. Buron, "Report on P2P-RPL Interoperability Testing", INRIA Research Report RR-7864, January 2012.
- [RT-MPL] van der Stok, P., "Real-Time multicast for wireless mesh networks using MPL", White paper, April 2014, <<http://www.vanderstok.org/papers/Real-time-MPL.pdf>>.
- [OccuSwitch] Philips lighting Electronics, "OccuSwitch Wireless (brochure)", May 2012, <http://www.philipslightingcontrols.com/assets/cms/uploads/files/osw/MK_OSWNETBROC_5.pdf>.
- [Office-Light] Clanton and Associates, Inc., "Wireless Lighting Control - A Life Cycle Cost Evaluation of Multiple Lighting Control Strategies", February 2014, <http://www.daintree.net/wp-content/uploads/2014/02/clanton_lighting_control_report_0411.pdf>.
- [RTN2011] Holtman, K. and P. van der Stok, "Real-time routing for low-latency 802.15.4 control networks", 23rd Euromicro Conference on Real-Time Systems, Porto, Portugal, July 2011.
- [MEAS] Holtman, K., "Connectivity loss in large scale IEEE 802.15.4 network", Private Communication, November 2013.

[BC-Survey]

Kastner, W., Neugschwandtner, G., Soucek, S., and H. Newmann, "Communication Systems for Building Automation and Control", Proceedings of the IEEE, Vol. 93, No. 6, DOI 10.1109/JPROC.2005.849726, June 2005.

[ZigBeeIP]

ZigBee Alliance, "ZigBee IP specification", ZigBee document 095023r34, March 2014, <<http://www.zigbee.org/>>.

Appendix A. RPL Shortcomings in Home and Building Deployments

A.1. Risk of Undesirable Long P2P Routes

The DAG, being a tree structure, is formed from a root. If nodes residing in different branches need to communicate internally, DAG mechanisms provided in RPL [RFC6550] will propagate traffic towards the root, potentially all the way to the root, and down along another branch [RFC6998]. In a typical example, two nodes could reach each other via only two router nodes, but in some unfortunate cases, RPL may send traffic three hops up and three hops down again. This leads to several undesirable phenomena, as described in the following sections.

A.1.1. Traffic Concentration at the Root

If many P2P data flows have to move up towards the root to get down again in another branch, there is an increased risk of congestion the nearer to the root of the DAG the data flows. Due to the broadcast nature of radio frequency (RF) systems, any child node of the root is not only directing RF power downwards in its sub-tree but just as much upwards towards the root, potentially jamming other MP2P traffic leaving the tree or preventing the root of the DAG from sending P2MP traffic into the DAG because the listen-before-talk link-layer protection kicks in.

A.1.2. Excessive Battery Consumption in Source Nodes

Battery-powered nodes originating P2P traffic depend on the route length. Long routes cause source nodes to stay awake for longer periods before returning to sleep. Thus, a longer route translates proportionally (more or less) into higher battery consumption.

A.2. Risk of Delayed Route Repair

The RPL DAG mechanism uses DIO and DAO messages to monitor the health of the DAG. On rare occasions, changed radio conditions may render routes unusable just after a destination node has returned a DAO indicating that the destination is reachable. Given enough time, the next Trickle timer-controlled DIO/DAO update will eventually repair the broken routes; however, this may not occur in a timely manner appropriate to the application. In an apparently stable DAG, Trickle timer dynamics may reduce the update rate to a few times every hour. If a user issues an actuator command, e.g., light on in the time interval between the time that the last DAO message was issued the destination module and the time that one of the parents sends the next DIO, the destination cannot be reached. There is no

mechanism in RPL to initiate the restoration of connectivity in a reactive fashion. The consequence is a broken service in home and building applications.

A.2.1. Broken Service

Experience from the telecom industry shows that if the voice delay exceeds 250 ms, users start getting confused, frustrated, and/or annoyed. In the same way, if the light does not turn on within the same period of time, a home control user will activate the controls again, causing a sequence of commands such as `Light{on,off,off,on,off,...}` or `Volume{up,up,up,up,up,...}`. Whether the outcome is nothing or some unintended response, this is unacceptable. A controlling system must be able to restore connectivity to recover from the error situation. Waiting for an unknown period of time is not an option. Although this issue was identified during the P2P analysis, it applies just as well to application scenarios where an IP application outside the LLN controls actuators, lights, etc.

Appendix B. Communication Failures

Measurements of connectivity between neighboring nodes are discussed in [RTN2011] and [MEAS].

The work is motivated by the measurements in literature that affirm that the range of an antenna is not circle symmetric but that the signal strength of a given level follows an intricate pattern around the antenna, and there may be holes within the area delineated by a polar plot. It is reported that communication is not symmetric: reception of messages from node A by node B does not imply reception of messages from node B by node A. The quality of the signal fluctuates over time, and also the height of the antenna within a room can have consequences for the range. As a function of the distance from the source, three regions are generally recognized: (1) a clear region with excellent signal quality, (2) a region with fluctuating signal quality, and (3) a region without reception. Installation of meshes with neighbors in the clear region is not sufficient, as described below.

[RTN2011] extends existing work by:

- o Observations over periods of at least a week,
- o Testing links that are in the clear region,
- o Observation in an office building during working hours, and
- o Concentrating on one-hop and two-hop routes.

Eight nodes were distributed over a surface of 30 square meters. All nodes are at a one-hop distance from each other, and all are situated in each other's clear region. Each node sends messages to each of its neighbors and repeats the message until it arrives. The latency of the message was measured over periods of at least a week. It was noticed that latencies longer than a second occurred without any apparent reason, but only during working days and never during the weekends. Bad periods could last for minutes. By sending messages via two paths -- (1) a one-hop path directly and (2) a two-hop path via a randomly chosen neighbor -- the probability of delays larger than 100 ms decreased significantly.

The conclusion is that even for one-hop communication between not-too-distant "line of sight" nodes, there are periods of low reception in which communication deadlines of 200 ms are exceeded. It pays to send a second message over a two-hop path to increase the reliability of timely message transfer.

[MEAS] confirms that temporary bad reception by close neighbors can occur within other types of areas. Nodes were installed on the ceiling in a grid with a distance of 30-50 cm between them. Two hundred nodes were distributed over an area of 10 m x 5 m. It clearly transpired that with increasing distance the probability of reception decreased. At the same time, a few nodes furthest away from the sender had a high probability of message reception, while some close neighbors of the sender did not receive messages. The patterns of nodes experiencing good reception evolved over time.

The conclusion here is that even for direct neighbors reception can temporarily be bad for periods of several minutes. For reliable and timely communication, it is imperative to have at least two communication paths available (e.g., two-hop paths next to the one-hop path for direct neighbors).

Acknowledgements

This document reflects discussions and remarks from several individuals, including (in alphabetical order) Stephen Farrell, Mukul Goyal, Sandeep Kumar, Jerry Martocci, Catherine Meadows, Yoshihiro Ohba, Charles Perkins, Yvonne-Anne Pignolet, Michael Richardson, Ines Robles, Zach Shelby, and Meral Sherazipour.

Authors' Addresses

Anders Brandt
Sigma Designs

Email: anders_Brandt@sigmadesigns.com

Emmanuel Baccelli
INRIA

Email: Emmanuel.Baccelli@inria.fr

Robert Cragie
ARM Ltd.
110 Fulbourn Road
Cambridge CB1 9NJ
United Kingdom

Email: robert.cragie@arm.com

Peter van der Stok
Consultant

Email: consultancy@vanderstok.org