

Networking Working Group
Internet-Draft
Intended status: Informational
Expires: July 10, 2014

J. Tripathi, Ed.
J. de Oliveira, Ed.
Drexel University
C. Chauvenet, Ed.
Watteco.
JP. Vasseur, Ed.
Cisco Systems, Inc.
January 6, 2014

Why Reactive Protocols are Ill-Suited for LLNs
draft-tripathi-roll-reactive-applicability-02

Abstract

This document describes serious issues and shortcomings regarding the use of reactive routing protocols in low power and lossy networks (LLNs). Routing requirements for various LLN deployments are discussed in order to judge how reactive routing may or may not adhere to the necessary criteria.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on July 10, 2014.

Copyright Notice

Copyright (c) 2014 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	2
2.	Terminology	3
3.	Non-compliance with routing requirements in LLNs	3
3.1.	Inability to support P2MP traffic pattern	3
3.2.	Inability to support MP2P traffic pattern	4
4.	Dependency of control overhead on application module	5
5.	Flooding issues in LLNs	6
5.1.	Impact of flooding in battery operated nodes	7
6.	Impact on memory	7
7.	Lack of support for routing based on node capability	8
8.	High delay for emergency traffic	9
9.	Acknowledgements	10
10.	Informative References	10
	Authors' Addresses	12

[1.](#) Introduction

In 2008, the IETF chartered a Working Group called ROLL (Routing Over Low power and Lossy networks) with the objective of specifying routing requirements of Low power and Lossy Networks (LLN), and then either select an existing routing protocol or specify and design a new routing protocol in light of the unique requirements of these networks. This led to the specification of a new routing protocol called RPL (see [[RFC6550](#)]) along with other specifications related to RPL.

Despite the existence of a standard track routing protocol for LLN, discussions have been taken place as to whether other routing approaches could be suitable, such as deploying reactive routing protocols in LLNs, such as in smart metering networks, industrial automation, water management networks, etc.. The aim of this document is not to discuss a specific reactive routing protocol but why reactive routing protocols in general are ill-suited for LLNs. For the sake of illustration, we will refer to a reactive protocol called LOADng ([[I-D.clausen-lln-loadng](#)]) which was introduced at the IETF, with results seeming to indicate performance improvement over the standard AODV protocol in the context of LLNs ([[clausen-load-vtc](#)]). The aim of this document is to highlight the number of shortcomings and technical issues in using reactive routing in such networks. Note that reactive protocols are perfectly applicable to several types of mobile ad-hoc networks (MANETs), which are not LLNs.

Specifically, LLNs exhibit constrained characteristics such as very low data rate (a few Kbps), lossy physical links where the packet delivery ratio can be poor (below 50%) and the links exhibit highly dynamic properties (high variation of Packet Delivery Ratio (PDR) with respect to time).

2. Terminology

Please refer to [[ROLL-TERMS](#)] for terminology. Additionally, the following terms are used throughout the document.

U-LLN: Urban Low-Power Lossy Network.

PDR: Packet Delivery Ratio.

RREQ Route Request.

RREP Route Reply.

AMI Advanced Metering Infrastructure.

3. Non-compliance with routing requirements in LLNs

3.1. Inability to support P2MP traffic pattern

An LLN often has a more demanding traffic pattern than simple traffic between two peer nodes (P2P traffic). The nature of deployments and applications running over an LLN often requires a given node to send the same data to multiple recipients, requiring multicasting or point-to-multipoint (P2MP) support from the routing protocol. These destinations may be several hops away, therefore necessitating an efficient dissemination method. Examples of such traffic include:

- o management information multicasted to all nodes belonging to a certain region in a landscape, certain part of the manufacturing pipeline in an industrial automation setting, upgrade of a new firmware,
- o new tariff notification to a group of users in a smart grid deployment,
- o a single node providing sensed data to multiple servers.

Indeed, [[RFC5548](#)] necessitates the support of P2MP traffic for a protocol to be suitable for U-LLN deployment. It describes - "Thus, the protocol(s) should be optimized to support a large number of unicast flows from the sensing nodes or sensing clusters towards a LBR, or highly directed multicast or anycast flows from the nodes

towards multiple LBRs." and "A U-LLN may also need to support efficient large-scale messaging to groups of actuators.", which represents a massive critical need for a reliable P2MP mechanism. This requirement also necessitates a valid MP2P traffic support as well, which will be described in the next section.

While a reactive routing protocol may be altered to support an on-demand bi-directional route between any two nodes in the network as provided by LOADng (see [[I-D.clausen-lln-loadng](#)]), the protocol does not define a way to build and maintain an always usable multicast topology. It is worth mentioning here that the AODV protocol in [[RFC3561](#)] mentions provision for nodes to join a multicast tree and RREQs to use multicast IP address as their destination address. However, no such methods or provisions are available in the lightweight versions of the AODV protocol, namely AODVv2 (formerly DYMO) [[I-D.Perkins-AODVv2](#)] or LOADng ([[I-D.clausen-lln-loadng](#)]). A naive and very costly solution would be to create a copy of the same message/application data to send for each destination. Also, to find the route to each destination, the node may have to create separate route request (RREQ) messages and broadcast them. This broadcast event creates a huge control overhead. Number of intended destinations for this P2MP traffic can reach an order of hundreds or thousands, which is very common in an LLN deployed in urban area, or for a number of Advanced Metering Infrastructure (AMI) meters in a particular region in a smart grid deployment. If such is the case, the protocol does not scale well with the network size in terms of control overhead. Hence, reactive routing protocols in general may become unsuitable to be deployed in a large scale LLN such as a U-LLN where P2MP traffic needs to be supported and be maintained, even if for 1-2 times a day.

Simulation studies, such as the on in [[Tripathi-CISS](#)] have verified that a reactive protocol suffers from high (over hundred seconds) end-to-end delay and large control overhead when P2MP traffic exists in the network.

3.2. Inability to support MP2P traffic pattern

Likewise P2MP traffic, MP2P traffic needs to be supported by a routing protocol intended for an LLN. This traffic pattern arises when multiple nodes may try to send data to a single sink at the same time. As [[RFC5548](#)] describes, "A U-LLN should support occasional large-scale traffic flows from sensing nodes through LBRs (to nodes outside the U-LLN), such as system-wide alerts." By nature, this kind of traffic may be time-critical, and the alerts may need to be delivered within a small time-frame. Similar traffic may include critical scenarios in an industrial automation LLN deployment, where

all nodes in one area may need to report malfunction, fire, or other emergency in that part of the manufacturing plant.

Since, by default, reactive routing does not provide any efficient MP2P routing paradigm, all nodes will create their own route request (RREQ) in order to find the route to the LBR or server, if the route is expired or not cached. This situation will create separate broadcast control messages from each node, which may lead to a broadcast storm in the network, similar to the P2MP traffic scenario. The broadcast storm may result in individual RREPs reaching the initiator node much later, yielding a high delay for the time-critical alert packets.

Routing for MP2P traffic is therefore not efficient or optimized, and does not follow an aggregation tree neither any such hierarchy is maintained in reactive protocols. As a result, it is difficult to devise a data aggregation algorithm compatible to AODV or any of its derivatives.

4. Dependency of control overhead on application module

An LLN that is provisioned to be used for data gathering purpose today may include additional application layer modules in the future. Smart grid deployments may need to implement new modules of management traffic from the base stations to AMI meters, in addition to what is envisioned at present. LLNs are evolving and therefore it is expected that new applications and requirements will be part of its future (an LLN may also be re-purposed).

Reactive protocols discover route to destination on an on-demand basis. If communication between same source-destination pair are spaced far apart in time, the protocol tends to discover the routes every time communication is requested by the application layer. With reactive routing protocols, adding a new application module which requires more data communication in between the existing data traffic pattern will incur control overhead. Hence, if a network is designed to operate within bounds in terms of maximum control overhead load, adding new application modules may well force the control overhead to surpass the designed maximum limit. For example, a deployment requiring both MP2P and P2P application may incur more overhead than a deployment which is currently working with only data aggregation. Since LLNs will undoubtedly require more application modules and management modules to be augmented in future, a suitable routing protocol should be able to cope with the added traffic. For the sake of illustration, many smart grid networks, which were originally designed for the purpose of advanced metering, now require a multi-service networks in support of a variety of applications including meter reading, use of meters of alarms, distribution automation and

electric vehicles, leading to a variety of traffic patterns each with different quality of service requirements.

As we observed in [[Tripathi-CISS](#)], even for a smaller network with less than 100 nodes, if the application data rate is increased from once in 12 hours to once an hour, the control overhead for a reactive protocol such as LOADng gets magnified by about 10 times.

5. Flooding issues in LLNs

A reactive protocol is well-suited for a traffic pattern where data transfer is not very frequent. However, if the traffic pattern includes periodic data reporting, even as low as a few times in a day, the traffic pattern will induce periodic broadcast of route request throughout the network. A simple example scenario can clarify this: assume an application in a U-LLN requires periodic data reporting every 6 hours or 4 times in a day - morning, noon, evening and night. If the network consists of 2000 nodes, which is a very conservative number in a typical U-LLN, the application alone will create a route request broadcast for each sensor node every 11 seconds, on average. Thus, over the life of the sensor network, a reactive protocol will use more control overhead than a proactive protocol.

The amount of flooding to discover routes may also be controlled via tweaking the route expiry time or route validity time. If a route is active, nodes should not waste network resources trying to find out the route to the same destination. Keeping a high expiry time for the routes, on the other hand may prevent flooding the next time data is generated for the same destination. However, the path may well have been invalid by the route expiry time. Considering LLN link characteristics, link flapping is a very frequent event. Hence, high route expiry time may lead a node to find out invalidity of a path, thus forcing to flood the network again for route discovery. Thus, increasing route expiry time or route validity time for an AODV-based reactive protocol may not prove to control flooding in LLN.

Proactively choosing a back-up path proves to be an effective way to ensure valid routing path in presence of link flaps, whereas reactive routing approaches are not able to cope with link dynamics, thus preventing its usage for LLNs. Furthermore, if traffic is sent along a broken path, a new request would consequently be generated, thus increasing the control traffic load, in addition to incurring additional delays for the user data.

5.1. Impact of flooding in battery operated nodes

Note that there is a lot of evidences supporting the claim that using flooding or scoped flooding to discover routes is ill-suited to power constrained Low-power and Lossy Networks (LLNs) in general. This is due to the low-power requirement. In low-power wireless networks, broadcast packets usually cost much more energy by the hardware to transmit than unicast ones due to implementations of sleeping mechanisms. As pointed out in [Sun-SenSys], supporting broadcast transmission is difficult while implementing Low Power Listening (LPL) due to asynchronous duty-cycles of nodes. As the wake up time for each node in a neighborhood is independent, often multiple transmissions of a single frame are required to emulate one broadcast transmission. These multiple transmissions may consist of unicast transmissions to each of the neighbor (RI-MAC, see [Sun-RIMAC]), or repeated transmission of the frame during the whole sleep interval as done in X-MAC-UPMA ([Buettner-SenSys]). Otherwise a really long preamble would be required, thus increasing power consumption more for broadcasts packets in either case. Ad-hoc networks which are normally always on, will not face this problem. Hence, reactive routing methods that use flooding mechanism for route discovery often, are more suited for Mobile Ad-hoc networks, but not for LLNs in general. LLNs should limit use of flooding or broadcasting packets as much as possible via algorithms such as Trickle [RFC6206]. However, at this point, we would exclude consideration of such hardware dependent energy expense from our discussion.

6. Impact on memory

Reactive routing protocols usually rely on route caching for discovered destination. Hence, if any node participates in multiple active flows in the network, the node needs to store next hop (and validity) information for each source and destination node in the network. Thus, depending on the user traffic, some nodes tend to increase their routing table size proportional to the number of flows passing through themselves. Hence, the nodes require more memory storage to operate successfully in these networks, depending on the traffic pattern. However, characteristics of LLNs never guarantee enough storage space in any node for storing routing tables. In LLNs, thus it is always advantageous to store route to a subset of nodes and still be able to find a path to any destination in the network. Reactive routing protocols, in their current specification or any variant, fail to achieve low memory requirement in nodes in the context of large scale LLNs.

Destination oriented flooding in LLN, tends to worsen this situation. Multiple route requests may reach the same node at the same time for different destinations. Even though the destination may never be

reached through the concerned node, the node still has to process and re-broadcast each request, along with its neighbors. Vital memory is consumed in RREQ/RREP processing and buffering in this method.

As we observed during simulation study and in [[Tripathi-CISS](#)], the memory requirement in a reactive protocol such as LOADng grows with the network scale considerably. More specifically, they depend on number of active flows through a node, or number of destinations a node serves for which the routes have not expired yet. For a large network of over 2000 nodes, the total memory occupancy for a node, including route tables and buffered packets, for a reactive protocol may reach over 200 KB, as observed during simulation. Even for a small home automation network, the authors in [[Vucinic-WCNC](#)] pointed out, that the number of route entries and thus memory consumption depends on route expiry time for a reactive protocol. With increase in the route expiry time or route hold time, average number of route entries in a node increases. If the route expiry time is higher, the number of entries increases more sharply with increase in network size. However, at the same time, the study in [[Vucinic-WCNC](#)] agrees with the study conducted in [[Tripathi-CISS](#)] that with decrease in route expiry time in order to save memory, control overhead increases as predicted in [section 5](#).

7. Lack of support for routing based on node capability

Apart from providing a route between any two nodes in the network, a routing protocol suitable for LLN should be able to handle additional constraints. An LLN mainly consists of constrained devices, both functionality and memory-wise, and inherently heterogeneous in nature. Hence, any routing protocol suitable for LLNs should support node constraint based routing. This requirement is mandated in [[RFC5548](#)] as follows: "the routing protocol MUST be able to advertise node capabilities that will be exclusively used by the routing protocol engine for routing decision". For example, the routing protocol should avoid a node with less battery power while routing to reach a server. Similarly, for industrial automation requirements, [[RFC5673](#)] also needs a routing protocol to provide device-aware routing, as it describes "The routing algorithm MUST support node-constrained routing (e.g., taking into account the existing energy state as a node constraint). Node constraints include power and memory, as well as constraints placed on the device by the user, such as battery life". For home routing automation, [[RFC5826](#)] specifies, "The routing protocol SHOULD route via mains-powered nodes if possible. The routing protocol MUST support constraint-based routing taking into account node properties (CPU, memory, level of energy, sleep intervals, safety/convenience of changing battery)".

Clearly, recognizing a node's capability and routing accordingly is an important aspect for any routing protocol designed to be suitable for LLNs. However, any AODV-based protocol (such as AODVv2, formerly DYMO [[I-D.Perkins-AODVv2](#)] and LOADng), in their current specification, fail to provide routes based on any such constraint. Currently known reactive routing protocols do not have any provision to determine whether the next-hop node in a route has enough battery power to sustain the route, or whether the next-hop node is main powered or provides a particular functionality. Thus, these protocols fail to provide requirements mandated by [[RFC5548](#)], [[RFC5673](#)] and [[RFC5826](#)] for routing in an LLN.

8. High delay for emergency traffic

Some data in an LLN are delay sensitive by nature. While data generated for periodic reporting can be delivered even up to few seconds later, emergency alarms, fault notification and alert packets need to be delivered as quickly as possible. According to [[RFC5673](#)], in an industrial automation setting, "Non-critical closed-loop applications have a latency requirement that can be as low as 100 milliseconds but many control loops are tolerant of latencies above 1 second". Clearly, these types of alert packets need a path to the destination beforehand as soon as they are generated. However, reactive protocols depend on finding a path when data is generated. Since the receipt of an RREP packet can take up to the network traversal time, for large networks the delivery of an emergency packet may take more than few hundred milliseconds. Also, if this emergency situation initiates a system wide alert from all nodes in the region to one or multiple servers outside the LLN, each node will create their own broadcast to reach the destination. Similar to the condition in [section 3](#), the broadcast storm may lead to huge delay in receipt of RREP packets, and thus delay in delivering the emergency packet. The broadcast will lead to high control overhead, clogging the network, as well as loss of some RREQ/RREP packets.

Indeed, reactive protocols provide a very high unreliability in delay bound for any number of hop distance to the destination. The delay is small when an active path is available, but high when the source has to issue a route request. The delay is a few magnitudes larger when multiple route requests/replies are simultaneously active. As observed during simulation and in [[Tripathi-CISS](#)], a reactive protocol as LOADng may take from tens of seconds to hundreds of seconds to deliver the data, depending on the network size and whether multiple RREQs/RREPs are active. Unreliability in delay bound has been found to be as high as 95% confidence interval in average end-to-end delay for the data packets. Also in [[Vucinic-WCNC](#)], the authors pointed out how reactive protocols, even

for a small home automation network presents large delay in order of a few seconds which increases with decrease in route expiry time.

9. Acknowledgements

TBD

10. Informative References

[Buettner-SenSys]

Buettner, M., Ed., Yee, G., Ed., Anderson, E., Ed., and R. Han, Ed., "X-MAC: A Short Preamble MAC Protocol for Duty-Cycled Wireless Sensor Networks", Proceedings of the 4th ACM Conference on Embedded Networked Sensor Systems (SenSys-06) 2006, November 2006.

[I-D.Perkins-AODVv2]

Perkins, C., Ed., Ratliff, S., Ed., and J. Dowdell, Ed., "Dynamic MANET On-demand (AODVv2) Routing", Work in Progress, July 2013.

[I-D.clausen-lln-loadng]

Clausen, T., Ed., Yi, J., Ed., Niktash, A., Ed., Igarashi, Y., Ed., Satoh, H., Ed., Herberg, U., Ed., Lavenu, C., Ed., Lys, T., Ed., and J. Dean, Ed., "The Lightweight On-demand Ad hoc Distance-vector Routing Protocol - Next Generation (LOADng) ([draft-clausen-lln-loadng-10](#))", Work in Progress, October 2013.

[RFC3561] Perkins, C., Belding-Royer, E., and S. Das, "Ad hoc On-Demand Distance Vector (AODV) Routing", [RFC 3561](#), July 2003.

[RFC5548] Dohler, M., Watteyne, T., Winter, T., and D. Barthel, "Routing Requirements for Urban Low-Power and Lossy Networks", [RFC 5548](#), May 2009.

[RFC5673] Pister, K., Thubert, P., Dwars, S., and T. Phinney, "Industrial Routing Requirements in Low-Power and Lossy Networks", [RFC 5673](#), October 2009.

[RFC5826] Brandt, A., Buron, J., and G. Porcu, "Home Automation Routing Requirements in Low-Power and Lossy Networks", [RFC 5826](#), April 2010.

[RFC5867] Martocci, J., De Mil, P., Riou, N., and W. Vermeylen, "Building Automation Routing Requirements in Low-Power and Lossy Networks", [RFC 5867](#), June 2010.

- [RFC6206] Levis, P., Clausen, T., Hui, J., Gnawali, O., and J. Ko, "The Trickle Algorithm", [RFC 6206](#), March 2011.
- [RFC6550] Winter, T., Thubert, P., 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](#), March 2012.
- [ROLL-TERMS]
Vasseur, JP., "Terminology in Low power And Lossy Networks", Work in Progress, September 2011.
- [Sun-RIMAC]
Sun, Y., Ed., Gurewitz, O., Ed., and D. Johnson, Ed., "RIMAC: A Receiver-Initiated Asynchronous Duty Cycle MAC Protocol for Dynamic Traffic Loads in Wireless Sensor Networks", Proceedings of the 6th ACM Conference on Embedded Networked Sensor Systems (SenSys-08), 2008, November 2008.
- [Sun-SenSys]
Sun, Y., Ed., Gurewitz, O., Ed., Du, S., Ed., Tang, L., Ed., and D. Johnson, Ed., "ADB: An Efficient Multihop Broadcast Protocol Based on Asynchronous Duty-cycling in Wireless Sensor Networks", Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems (SenSys-09), 2009, November 2009.
- [Tripathi-CISS]
Tripathi, J., Ed. and J. de Oliveira, Ed., "Proactive versus Reactive Revisited: IPv6 Routing for Low Power Lossy Networks", 47th Annual Conference on Information Science and Systems (CISS), 2013, March 2013.
- [Vucinic-WCNC]
Vucinic, M., Ed., Tourancheau, B., Ed., and A. Duda, Ed., "Performance Comparison of the RPL and LOADng Routing Protocols in a Home Automation Scenario", IEEE WCNC - Wireless Communications and Networking Conference 2013 2013, June 2013.
- [clausen-load-rtc]
Clausen, T., Ed., Yi, J., Ed., and A. de Verdiere, Ed., "LOADng: Towards AODV Version 2", Vehicular Technology Conference (VTC Fall), 2012 IEEE, September 2012.

Authors' Addresses

Joydeep Tripathi (editor)
Drexel University
3141 Chestnut Street 7-313
Philadelphia, PA 19104
USA

Email: jt369@drexel.edu

Jaudelice C. de Oliveira (editor)
Drexel University
3141 Chestnut Street 7-313
Philadelphia, PA 19104
USA

Email: jau@coe.drexel.edu

C. Chauvenet (editor)
Watteco.
1766, Chemin de la Planquette
La Garde. 83130
France

Email: c.chauvenet@watteco.com

JP Vasseur (editor)
Cisco Systems, Inc.
11, Rue Camille Desmoulins
Issy Les Moulineaux 92782
France

Email: jpv@cisco.com

