Networking Working Group
Internet-Draft
Intended status: Informational
Expires: September 24, 2010

J. Tripathi, Ed.
J. de Oliveira, Ed.
Drexel University
JP. Vasseur, Ed.
Cisco Systems, Inc.
March 23, 2010

Performance Evaluation of Routing Protocol for Low Power and Lossy Networks (RPL)
draft-tripathi-roll-rpl-simulation-03

Abstract

This document presents a performance evaluation of the Routing Protocol for Low power and Lossy Networks (RPL). Detailed simulations are carried out to produce several routing performance metrics using a set of real-life scenarios.

Status of this Memo

Copyright Notice

Contents

1 Terminology 3

2 Introduction 3

3 Method 4

4 Simulation Setup 4

5 Metrics to evaluate RPL 6
5.1 Common Assumptions . . . 6
5.2 Path Quality . . . 6
5.3 Routing Table Size . . . 8
5.4 Delay bound for P2P Routing . . . 8
5.5 Control Packet Overhead . . . 9
5.6 Loss of connectivity . . . 10

6 RPL in a building routing scenario 15
6.1 Path Quality . . . 15
6.2 Delay . . . 17

7 References 18
7.1 Normative References . . . 18
7.2 Informative References . . . 18

1 Terminology

PDR - Packet Delivery Ratio

Please refer to additional terminology in [I-D.ietf-roll-terminology].

2 Introduction

Designing routing in low power devices and lossy link networks (LLNs) imposes great challenges, mainly due to low data rates, high probability of packet delivery failure, and strict energy constraint in nodes. The IETF ROLL Working Group has specified the Routing Protocol for Low power and Lossy Networks (RPL) in [I-D.ietf-roll-rpl].

RPL is designed to meet the core requirements specified in [I-D.ietf-roll-home-routing-reqs],[I-D.ietf-roll-building-routing-reqs],[I-D.ietf-roll-indus-routing-reqs] and [RFC5548].

This document's contribution is to provide several routing performance metrics of RPL using a decrete event simulator in various real-life deployment scenarios. Each result has been checked against several real-life deployed networks.

Simulation results are purely indicative since they may vary according to the discrete event simulator used to perform the simulations, the choice of the RPL parameter and so on. Still this document provides valuable inputs and the specific context in which these simulations were performed are explicitly indicated.

Several routing metrics are evaluated in this document:

- Path quality metrics;
- Control plane overhead;
- End to End delay between nodes.
- Ability to cope with unstable situations (link churns, node dying);
- Required resource constraints on nodes (routing table size, etc.).

Feedback from the ROLL Working Group are welcome to add new evaluation metrics of potential interest in further revisions of this document.

Although simulation cannot prove formally that a protocol operates properly in all situations, it could give a good level of confidence in protocol behavior in highly stressful conditions, if and only if real life data are used. Simulation is particularly useful especially because theoretical model assumptions may not be applicable to LLNs. Therefore, real deployed network data traces have been used to model link behaviors.

3 Method

RPL was simulated using OMNET++ [OMNETpp], a well-known discrete event based simulator written in C++ and NED. Castalia-2.2 [Castalia-2.2] has been used as Wireless Sensor Network Simulator framework within OMNET++. The output and events in the simulating are visualized with the help of the Network AniMator or NAM, which is distributed with NS (Network Simulator) [NS-2].

Note that NS or any of its versions were not used in this simulation study. Only the visualization tool was borrowed for verification purposes. As noted, real link layer data gathered from networks deployed on the field were used to compute the PDR (Packet Delivery Ratio) for each of the links in the network. By contrast with theoretical models (e.g. Markov Chains) which may have assumptions not applicable to lossy links, real-life data has been used for two aspects of the simulations:
* Link failure model: Time varying real network traces containing packet delivery probability for each link and over all channels for both indoor network deployment and outdoor network deployment were used. Thus, different types of link characteristics are used in the study.
* Topology: The topologies are gathered from real-life deployment (traces mentioned above) as opposed to random topology simulations.

4 Simulation Setup

A 45 node topology, shown in 1, gathered from a real deployment, was used in the simulations.

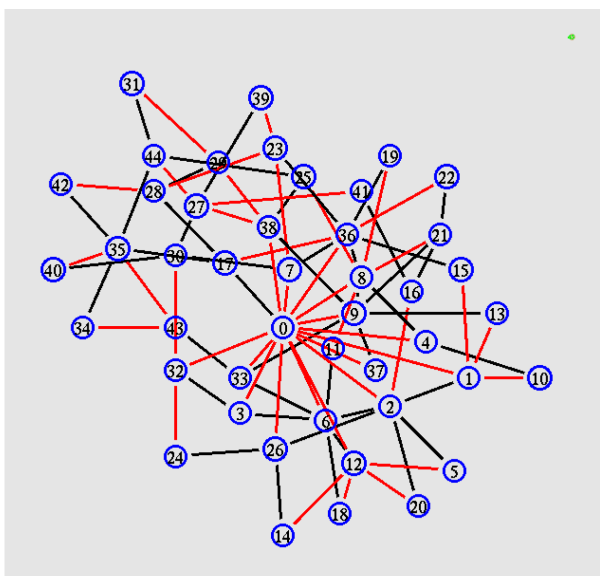

Figure 1: Network topology for preliminary simulation results.

Note that this is just a start to validate the simulation before using large scale networks.

A database of time varying link quality data, gathered from real network deployment, was created. Each link in the topology 'picks up' a link model from the database, and the link's Packet Delivery Ratio (PDR) varies according to the gathered data. Figure 2 shows some typical temporal characteristics of some links in the network for the indoor network trace used in the simulations. Packets are dropped randomly from that link with probability (1 - PDR). Each link has a PDR that varies with time (in the simulation, the new PDR is read from the database every 10 minutes). Each time a packet arrives at the Radio of a node, the module generates a random number by the Mersenne Twister Random number generation method. The random number is compared to the PDR to determine whether the packet should be dropped or not. Note that each link use a different random number generator to maintain true randomness in the simulator, and to avoid correlation between links. Also, the packet drop applies to all kinds of data and control packets (RPL) such as the DIO, DAO, DIS packets defined in [I-D.ietf-roll-rpl].

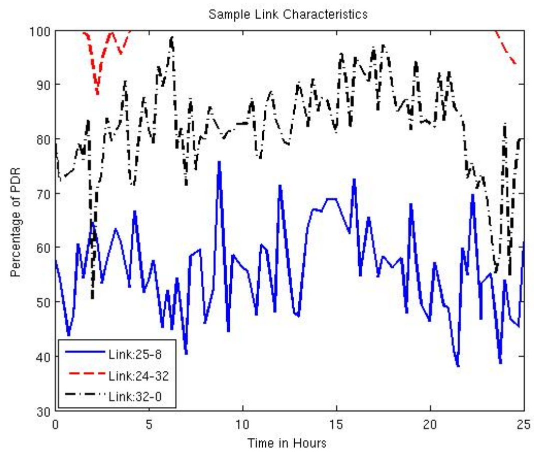

Figure 2: Example of link characteristics.

In simulating RPL, the LBR first initiates sending out DIO messages, and the DAG is gradually constructed. The trickle time interval for emitting DIO message assumes the initial value of 1 second, and then changes over simulation time as mentioned in [I-D.ietf-roll-rpl].

RPL makes use of trickle timers: I_min is initially set to 1 second and I_doubling is equal to 16, so that maximum time between two consecutive DIO emissions by a node (under a steady network condition) is 18.2 hours. Another objective of this study is to give insight to the network administrator on how to tweak the trickle values. These recommendations could then be used in applicability statement documents. Further revision of this document will include simulations for large scale networks with varied parameters and show how quickly the network will stabilize, comparing data/control traffic and studying the tradeoff between reactivity and lifetime.

Each node in the network, other than the LBR, also emits DAO messages as specified in [I-D.ietf-roll-rpl], to initially populate the routing tables with the prefixes received from children via the DAO messages in support of the Point to Point (P2P) and Point to Multipoint traffic (P2MP) in the “down” direction. In this revision of the document, it is assumed that each node is capable of storing route information for other nodes in the network. In futher revision of this document nodes without storage capability will be added to the network to see the influence of extra states on the nodes and the additional control plane overhead to propagate the route records thanks to Reverse Route Stacks in the DAO messages.

For nodes implementing RPL, as expected, the routing table memory requirement varies according to the position in the DAG. The worst-case assumption that there is no route summarization in the network is made. Thus a node closer to the DAG will have to store more routing entries. Further revision of this document will explore the influence of performing route summarization along the DAG, which could be performed thanks to a newly defined Objective Function or new address provisioning techniques. It is also assumed that all nodes have equal memory capacity to store the routing states, therefore no source routing is required.

Each node sends traffic according to a Constant Bit Rate (CBR) to all other nodes in the network over the simulation period. To simulate a more realistic scenario, 20% of the generated packets by each node are destined to the root, and the remaining 80% of the packets are uniformily assigned as destined to nodes other than the root. Therefore the root receives a considerably larger amount of data than other nodes. These values may be revised when studying the P2P traffic so as to have a majority of traffic going to all nodes as opposed to the root. In the later part of the simulation, a typical home/building routing scenario was also simulated, and different path quality metrics were computed for that traffic pattern.

The packets are routed through the DAG built by RPL according to the mechanisms specified in [I-D.ietf-roll-rpl].

Since RPL is an IP routing protocol, no assumption is made on the link layer, thus potential gains in terms of header compression provided by 6loWPAN is not under consideration [draft-iphc].

A number of RPL parameters are used (such as Packet Rate from each source, Time Period of the LBR emitting new DAG Sequence Number) to observe their effect on the RPL performance metric of interest.

5 Metrics to evaluate RPL

5.1 Common Assumptions

Routing Table Size: as the DAO messages are used to feed the routing tables in the network, routing table size for each node are recorded. Currently, the routing table size is not expressed in terms of Kbyte of memory usage but measured in terms of number of entries for each node. Each entry has next hop node and path cost associated with the destination node. In further revision of this document, a single full 128-bit address per leaf plus a few bits to store other information and flags will be used.

The link ETX (Expected Transmission Count) metric is used to build the DAG as specified in [I-D.ietf-roll-routing-metrics]. Further revisions of this document will include other metrics and constraints such as the Hop count.

5.2 Path Quality

Number of Hops: For each pair of source and destination, the average number of hops for both RPL and shortest path routing is computed. Shortest path routing refers to an hypothetical ideal routing protocol that would always provide the shortest path in term of Total path cost ETX (or whichever metric is used) in the network. The Cumulative Distribution Function (CDF) of hop distance for all paths (which is equal to n*(n-1) in an n node network) in the network with respect to number of hops is plotted in Figure 3 for both RPL and shortest path routing. One can observe that the CDF corresponding to 4 hops is around 55% for RPL and 90% for shortest path routing. In other words, for the given topology, 90% of paths will have path length of 4 hops or less with an ideal shortest path routing methodology, whereas in RPL

Point-to-Point (P2P) routing, 90% of paths will have a length shorter or equal to 5 hops. This result shows that despite having a non optimized P2P routing scheme, the path quality of RPL is not much worse than an optimized one. Another reason may be, the sink is at the center of the network, so routing through the sink is often close to an optimal (shortest path) routing. This result may be different in a topology where the sink is located at one end of the network.

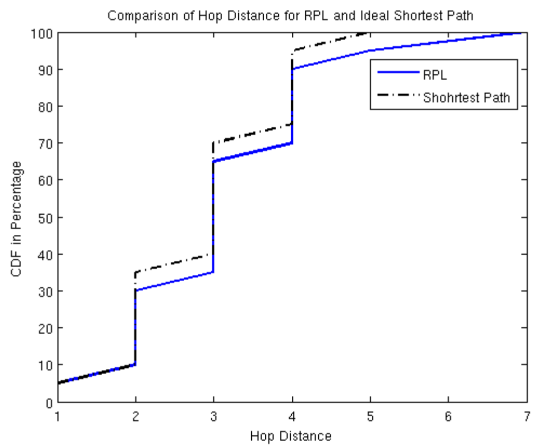

Figure 3: CDF: hop distance versus number of hops.

Path Cost ETX: When optimizing the path using link ETX metric, the path cost ETX of the path is computed for each pair. Figure 4 shows the CDF of the total number of ETX to deliver a packet from a source to any destination node with respect to total ETX of the path from each source to each destination in the network, for both RPL, and a shortest path routing. Here also one observes that total path cost ETX along the path from all source to all destination is close to that of a shortest path routing for the network in the simulation.

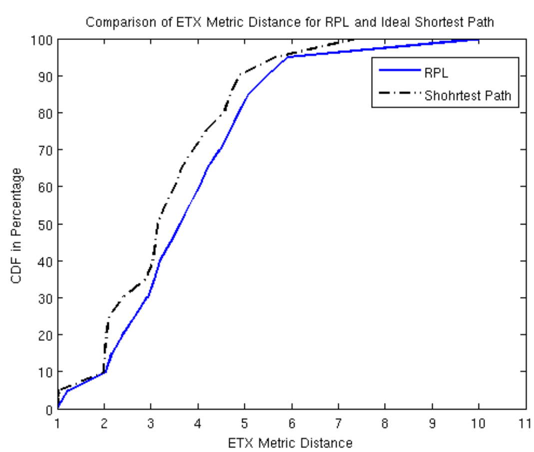

Figure 4: CDF: Total ETX along path versus ETX value.

Path Stretch: In this simulation, the path stretch is also calculated for each packet that traversed the network. The path stretch is determined as the difference between the number of hops taken by a packet while following a route built via RPL and the number of hops taken by shortest path routing (by using link ETX as the metric). Once again, the CDF of path stretch is plotted against the value of path stretch for different packets in Figures 5 and 6 for hop count stretch and ETX metric stretch respectively.

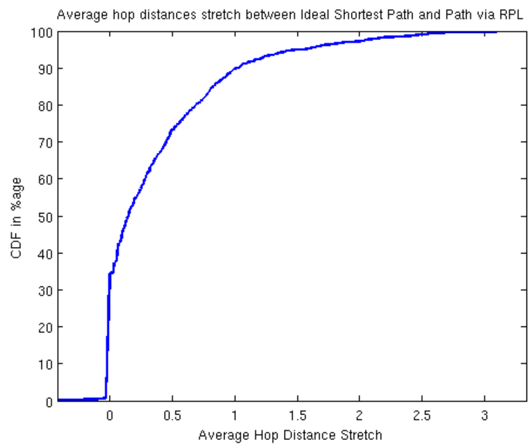

Figure 5: CDF: Hop count stretch versus hop count of a packet.

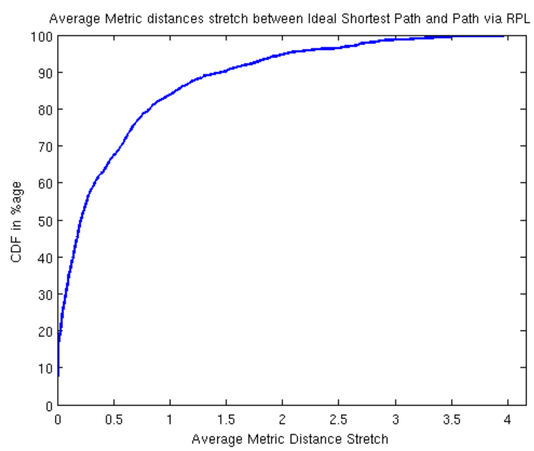

Figure 6: CDF: ETX metric stretch versus ETX value.

5.3 Routing Table Size

The objective of this metric is to observe the distribution of the number of entries per node. Figure 7 shows the CDF of required number of routing table entries for all nodes. One can see, that 90% of the nodes need to store less than 10 entries in their routing cache.

5.4 Delay bound for P2P Routing

For delay sensitive applications, such as home and building routing, etc., it is also important to limit the end-to-end delay. Figure 8 shows the upperbound and distributions of delay in P2P routing between any two given nodes when RPL is employed for different hop counts between source and destination. Here, the hop count refers to the hop distance when RPL is employed and not shortest path distance between two

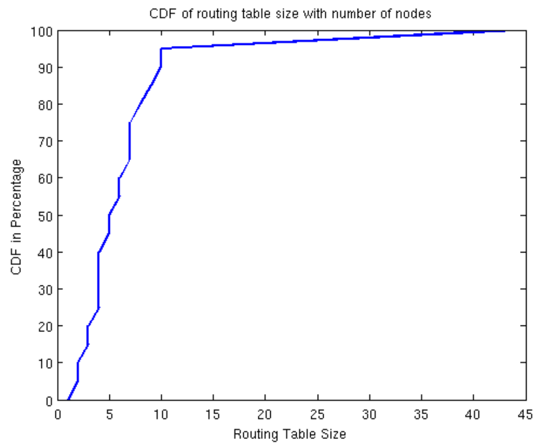

Figure 7: CDF of routing table size with respect to number of nodes.

nodes. Each packet has a length of 127 bytes, with a 240 kbps radio, which makes the transmission time to be approximately 4 ms.

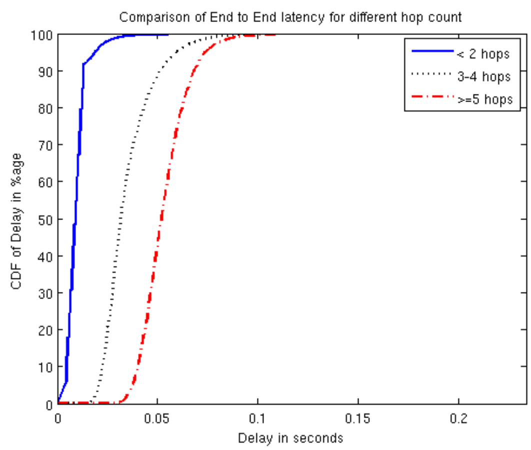

Figure 8: Comparison of packet latency for different hop count in RPL.

5.5 Control Packet Overhead

The control plane overhead is an important routing metric in Low power and Lossy Networks (LLNs). Indeed, it is imperative to bound the control plane overhead. One of the distinctive charateristics of RPL is that it makes use of trickle timers so as to reduce the number of control plane packets by eliminating redundant messages. The aim of this metric is thus to analyse the control plane overhead in stable condition (no network element failure overhead) and in the presence of failures.

Data and control plane traffic comparison for each node: Figure 9 shows the comparison of the amount of data packets transmitted (including forwarded) and control packets (DIO and DAO messages) transmitted for each node when minimizing ETX is used by the OCP along the DAG. Here one can observe that considerable amount of traffic is routed through the sink itself. And also the fact that the amount of control traffic is really negligible in the protocol is reinforced. As expected, the nodes closer to sink and that act as forwarders handle much more data packet transmission than other nodes. The leaf nodes have comparable amount of data and control packet transmission, as they do not take part in routing the data.

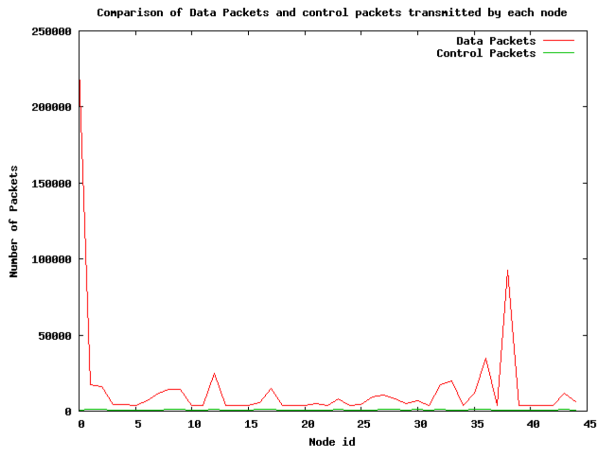

Figure 9: Amount of data and control packets transmitted for each node when minimizing ETX is used OCP along the DAG.

Data and Control Packet Transmission with respect to time: In Figures 10, 11 and 12, the amount of data and control packets transmitted for node 12 (low rank in DAG, closer to the root), node 43 (in the middle) and node 31 (leaf node)are shown, respectively. These values stand for number of packets transmitted for each 10 minutes intervals, to help understand what is the density of data and control packet exchange in the network. One can observe as the node is closer to the sink, the amount of data is larger, and the amount of control traffic is negligible in comparison to the data traffic. Also, the variation in data traffic is much larger for a node closer to sink, because the destination of the packets varies over time, and 20% of the packets are destined to sink only. For the nodes that are further away from sink , the variation in data traffic becomes lesser, and the amount of data traffic is also smaller.

The control traffic for the nodes has a wave-like pattern. The amount of control packets for each node drops quickly as the DAG stabilizes due to the effect of trickle timer. However, as a new DAG Sequence is advertised, the trickle timers are reset and the nodes start emiting DIO frequently again to stabilize the DAG. One can see, for a node closer to sink, the data packet amount is much higher than control packet, and somewhat oscillatory around a mean value. The control packet amount exhibits a 'saw-tooth' behavior, mainly because as the ETX link metric was used, and as when PDR changes, ETX path cost for a child node to its parent changes, which results in changing DAG rank of the child. This event resets the trickle timer and emit new DIO. Therefore, one can observe that the number of control packets attains a high value for one interval, and the amount comes down to lower values for subsequent intervals. Also, for leaf nodes the amount of control packets are more than data packets, as leaf nodes are more prone to face changes in their DAG rank as opposed to nodes closer to sink when the link PDR in the topology changes dynamically.

5.6 Loss of connectivity

Upon link failures, a node may loose his parents: preferred and backup (if any) and its sibling (if any). In this case, if a packet has to be sent and the routing table does not contain an entry for the corresponding destination the packet is dropped. RPL proposes two mechanims for DAG repairs, known as Global Repair

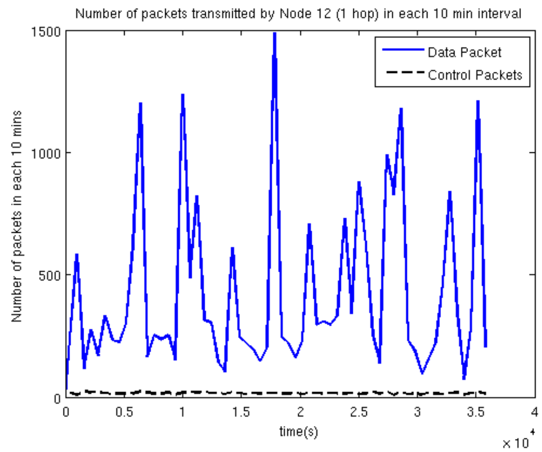

Figure 10: Amount of data and control packets transmitted for node 12.

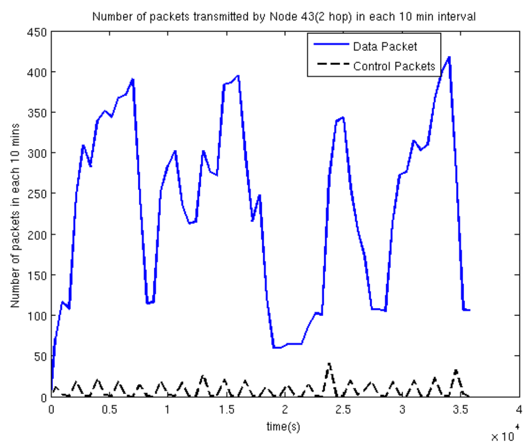

Figure 11: Amount of data and control packets transmitted for node 43.

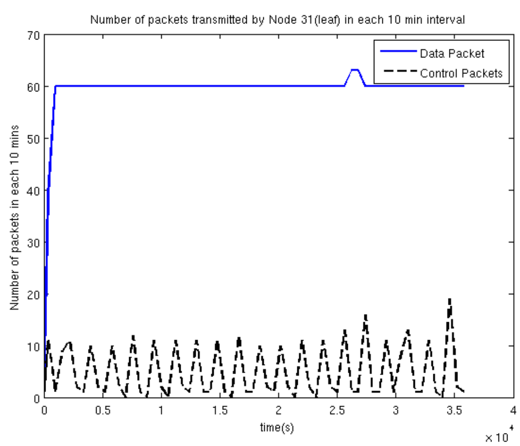

Figure 12: Amount of data and control packets transmitted for node 31.

and Local Repair. In this version of the document, simulation results are presented to evaluate the amount of time packets are lost because of loss of connectivity for two cases: *a*) when only Global Repair mechanism is implemented (i.e. of periodic emission of new DAG Sequence number by the DODAG root), and *b*) when poisoning the sub-DAG is used in case of unreachability of any parent or sibling node to forward data along with Global Repair mechanism. The idea is to tune the frequency at which new DAG Sequence Numbers are generated by the DAG root that are used for Global Repair, and also to observe the effect of the same when local repair is used in conjunction. It is expected that a higher frequency will lead to shorter duration of connectivity loss at a price of a higher rate of control packet in the network. For local repair, the simulation results show the trade-off in amount of time that a node may remain without service and total number of control packets for extra bit of signalling.

Figure 13 shows the CDF of time spent by any node without any service, when the packet rate from the sources is a packet each 10 seconds, and new DAG Sequence Number is issued every 10 minutes. This plot reflects the property of Global Repair without any Local Repair scheme. When all the parents (and siblings) are temporarily unreachable from a node, the time before it hears a DIO from another node is recorded, which gives the time without service. In some cases, this value might go up to the DAG Repair Timer value, because untill a DIO is heard, there is a lack of connectivity.

The effect of the DAG Repair Timer on time without service is plotted in Figure 14, where the source rate is 20 seconds/packet and in Figure 15, where the source sends a packet every 10 seconds.

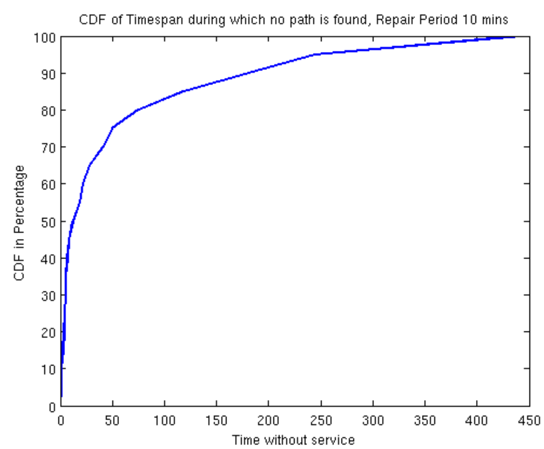

Figure 13: CDF: Loss of connectivity.

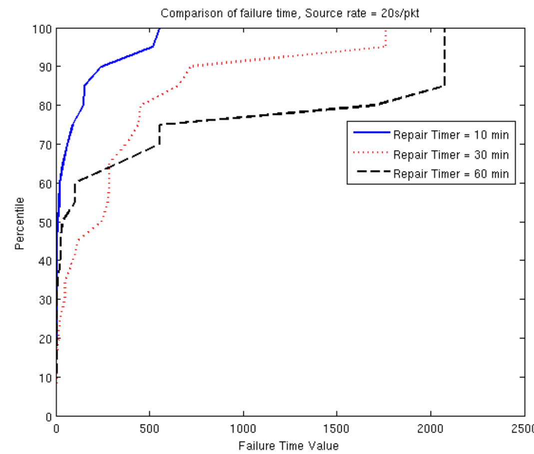

Figure 14: CDF: Loss of connectivity for different global repair period, packet rate 20/s.

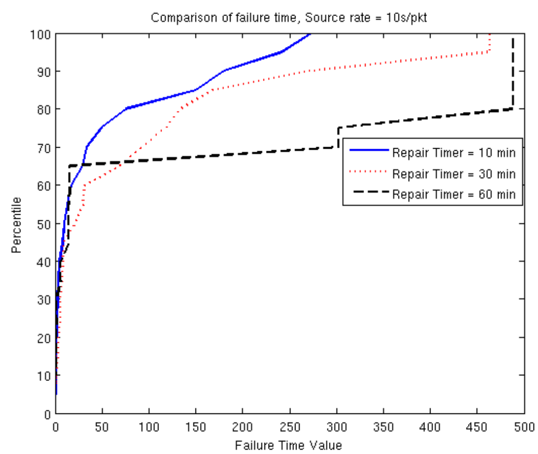

Figure 15: CDF: Loss of connectivity for different global repair period, packet rate 10/s.

Figure 16 shows effect of DAG Global Repair Timer period on control traffic. As expected, as the frequency at which new DAGSequenceNumber are generated increases, the amount of control traffic also decreases because the trickle interval gets larger for each node, which is pretty intuitive. However this smaller amount of control traffic comes at a price of increased time for loss of connectivity.

The effect of the DAG Repair Timer on time without service, when Local Repair is present, is plotted in 17, where the source rate is 20 seconds/packet. A comparison of the CDF of loss of connectivity for Global Repair Mechanism and Global + Local Repair Mechanism is shown in Figures 18 and 19 (semilog plots), where the source generates a packet every 10 seconds and 20 seconds respectively. In the plots, one can observe that using the method of poisoning the sub-DAG greatly reduces the time without connectivity.

A comparison between the amount of control overhead used for global repair only and global plus local Repair mechanism is shown in Figure 20, which highlights the improved performance of RPL in terms of convergence time at very little extra overhead.

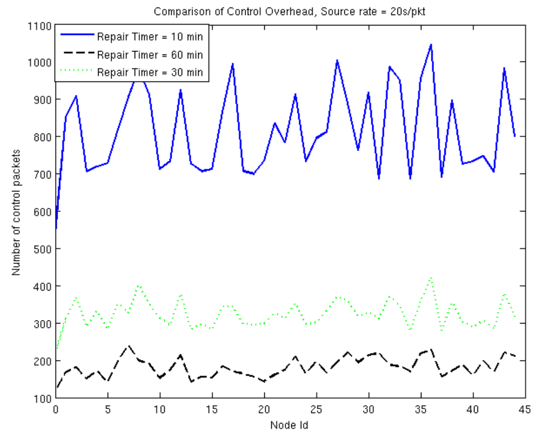

Figure 16: Amount of control traffic for different global repair timer period.

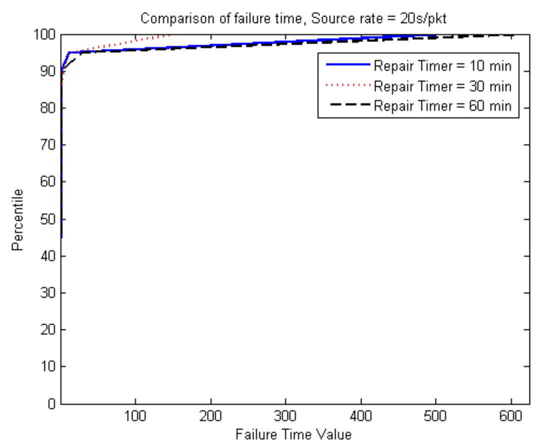

Figure 17: CDF: Loss of connectivity for different global repair period with poisoning, packet rate 20/s.

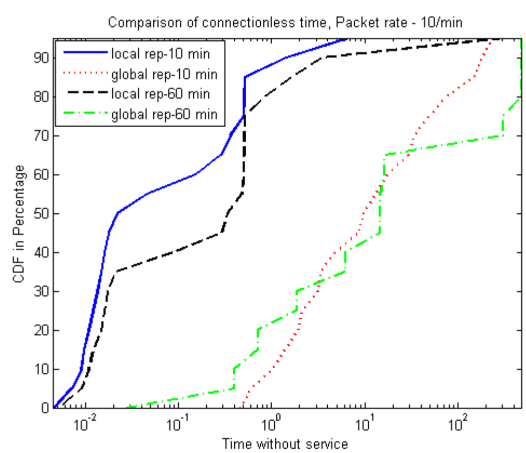

Figure 18: CDF: Comparing loss of connectivity for global repair and poisoning, packet rate 10/s.

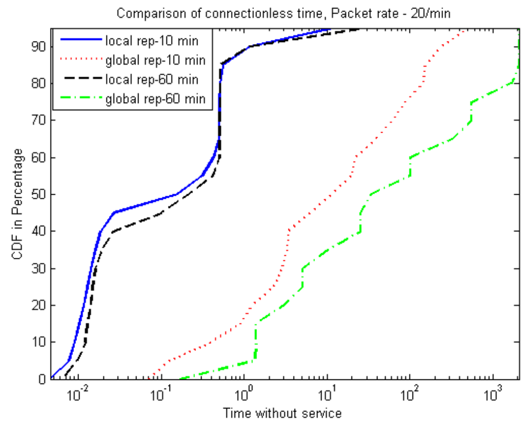

Figure 19: CDF: Comparing loss of connectivity for global repair and poisoining, packet rate 20/s.

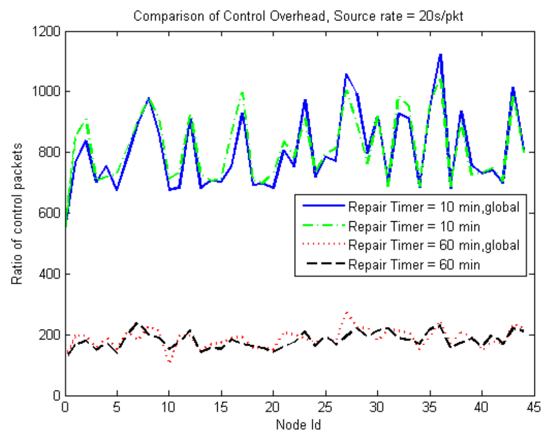

Figure 20: Number of control packets for different DAG Seq Number period, for both global repair and poisoning.

6 RPL in a building routing scenario

Unlike the previous traffic pattern, where a majority of the total traffic generated by any node is destined to the root, this section considers a different traffic pattern, which is more proominent in home or building routing scenario. A node sends 60% of its total generated traffic to its physically 1-hop distant nodes, 20% of traffic to its 2-hop distant nodes. Rest of the traffic is once again distributed among all other nodes in the network. The CDF of average hop distance path stretch in terms of hop distance, ETX path cost and delay for P2P routing for all pair of nodes is calculated. The delay bound is more important in this scenario, as the applications in home and building routing has typically low delay tolerance.

6.1 Path Quality

Figure 21 shows the CDF of number of hops for both RPL and ideal shortest path routing for the traffic scenario described above. Figure 22 shows CDF of the expected number of transmission count for each packet to reach destination. Figures 23 and 24 show CDF of the stretch factor for these two metrics.

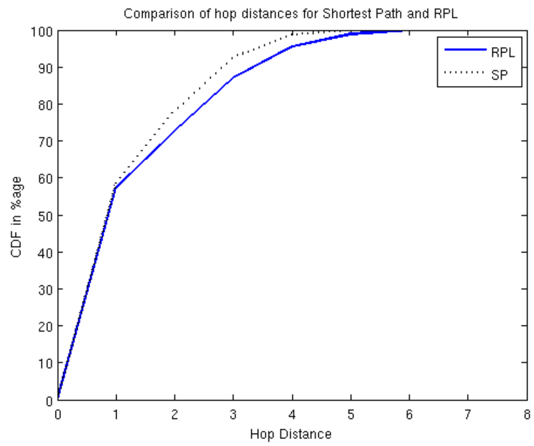

Figure 21: Comparison of end-to-end hop distancce for RPL and ideal shortest path in home routing.

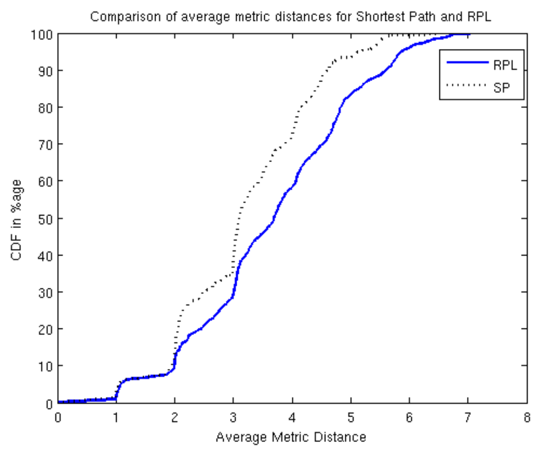

Figure 22: Comparison of link ETX metric for RPL and ideal shortest path in home routing.

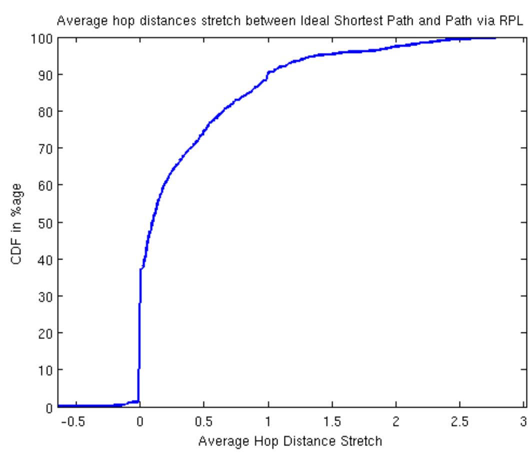

Figure 23: Stretch factor for node hop distance with ideal shortest path.

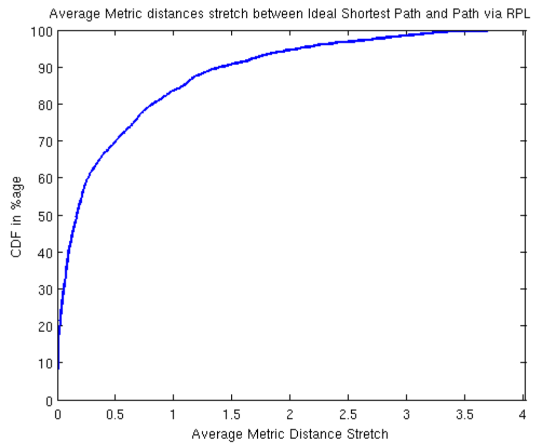

Figure 24: Stretch Factor of link ETX Metric with ideal shortest path.

6.2 Delay

To get an idea of maximum observable delay in the mentioned traffic pattern, the delay for different number of hops to the destination for RPL is considered. Figure 25 shows how the end-to-end packet latency is distributed for different packets with different hop counts in the network.

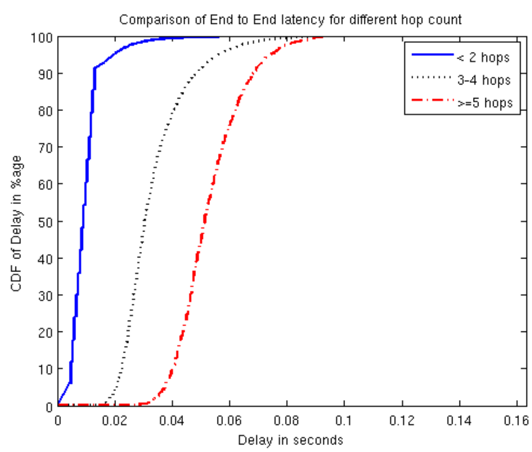

Figure 25: Comparison of packet latency for different hop count in RPL.

The next figures compare RPL to shortest path routing in terms of packet latency in a home/building routing scenario. Figure 26 shows that RPL almost meets the lowest delay bound set by the ideal shortest path route. This is interesting, because despite being a non-optimal algorithm, for this specific home routing scenario, the performance of RPL is very close to the optimal. Of course, for different topology this outcome may vary.

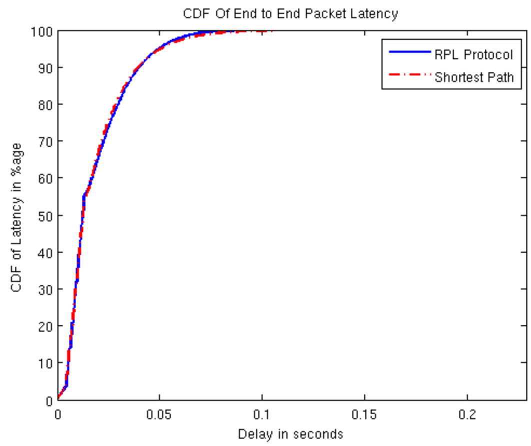

Figure 26: Comparison of packet latency in RPL and in ideal shortest path.

7 References

7.1 Normative References

[RFC2119] Bradner, S.,"Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

7.2 Informative References

[I-D.ietf-roll-building-routing-reqs] Martocci, J., Riou, N., Mil, P., and W. Vermeylen, "Building Automation Routing Requirements in Low Power and Lossy Networks, draft-ietf-roll-building-routing-reqs-07 (work in progress)", September 2009.

[I-D.ietf-roll-home-routing-reqs] Brandt, A., Buron, J., and G. Porcu, "Home Automation Routing Requirements in Low Power and Lossy Networks, draft-ietf-roll-home-routing-reqs-08 (work in progress)", September 2009.

[I-D.ietf-roll-indus-routing-reqs] Pister, K., Thubert, P., Dwars, S., Phinney, T., "Industrial Routing Requirements in Low Power and Lossy Networks, draft-ietf-roll-indus-routing-reqs-06 (work in progress)", June 2009.

[I-D.ietf-roll-routing-metrics] Vasseur, JP, Kim, M., Pister, K., Chong, H., "Routing Metrics used for Path Calculation in Low Power and Lossy Networks, draft-ietf-roll-routing-metrics-04 (work in progress)", December 2009.

[I-D.ietf-roll-rpl] Winter, T., Thubert, P., et al., "RPL: Routing Protocol for Low Power and Lossy Networks, draft-ietf-roll-rpl-04 (work in progress)", November 2009.

[I-D.ietf-roll-terminology] JP Vasseur, "Terminology in Low power And Lossy Networks, draft-ietf-roll-terminology-02 (work in progress)", May 2009.

[NS-2] "The Network Simulator-2".

[OMNETpp] Varga, "The OMNeT++ Discrete Event Simulation System", June 2001.

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

[draft-iphc] J. Jurski, "Limited IP Header Compression over PPP, draft-jurski-pppext-iphc-02.txt (work in progress)", March 2007.

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@ece.drexel.edu

JP Vasseur (editor)
Cisco Systems, Inc.
1414 Massachussetts Avenue
Boxborough, MA 01719
USA

Email: jpv@cisco.com