

Redundant Border Routers for Mission-Critical 6LoWPAN Networks

Laurent Deru¹, Sébastien Dawans¹, Mathieu Ocaña¹, Bruno Quoitin², and Olivier Bonaventure³

¹ CETIC, Rue des Frères Wright 29/3, B-6041 Charleroi, Belgium

² Université de Mons, 20, place du Parc, B-7000 Mons, Belgium

³ Université Catholique de Louvain, B-1348 Louvain-la-Neuve, Belgium
first.last@{cetic.be, umons.ac.be, uclouvain.be}

Abstract. Sensor networks are gradually moving towards full-IPv6 architectures and play an important role in the upcoming Internet of Things. Some mission-critical applications of sensor networks will require a level of reliability that excludes the presence of single points of failure, as it is often the case today for the gateways connecting sensor networks to the Internet. In this paper, we introduce RPL-6LBR, a 6LoWPAN border router that addresses mission-critical deployments through redundancy. The paper discusses how existing standards may be leveraged to enable redundant border router synchronization, while identifying certain of their shortcomings. We also propose innovative network architectures incorporating multiple border routers, which deal with redundancy and node mobility without requiring any synchronization among the border routers. We implement the proposed RPL-6LBR in the Contiki operating system and report on this implementation through trials on a small-scale testbed and simulator. Our results open new possibilities for real-world wireless sensor networks requiring reliable border routers, and guide further standardization efforts in emerging technologies in support of the Internet of Things.

1 Introduction

In the foreseeable future, billions of low power wireless devices will be connected to the Internet. To make this Internet of Things (IoT) vision a reality, constrained devices are adopting standards-based solutions [4]. The IoT embraces the IPv6 networking architecture to fulfill these requirements. IPv6-based Wireless Sensor Networks (WSN) enable a diverse range of application domains from smart cities to building and home automation. As the IoT-enabling technologies mature, so do the requirements on interconnectivity, reliability and fault-tolerance.

A key element of any deployed IoT network is the border router (BR) that connects it to the Internet. The BR is a gateway between two different link-layer technologies, typically 802.15.4 on the WSN side, and Ethernet or Wi-Fi on the other. The BR also separates the control planes of two different routing domains. In the majority of current deployments, the BR is unique.

When deploying large IoT networks, several issues appear such as the high solicitation of nodes close to the BR, the single point of failure (SPoF) that the BR is, and interconnecting seamlessly with existing IPv6 networks. We argue that for reliable mission-critical networks, redundant BRs need to be deployed. With multiple BRs, the packets can be load-balanced across these BRs and their neighboring nodes if the BRs are deployed at different physical locations, spreading the energy consumption across the network and reducing the likelihood of bottlenecks during traffic bursts. Using multiple BRs also reduces the risk of unreachability of nodes if network partitioning occurs due to failures within the WSN. However, using multiple BRs raises new problems such as the need to synchronize them, how to recover from failures, how to handle node mobility and how to integrate them seamlessly into an existing IPv6 network.

In this paper, we analyze how large IoT networks attached via redundant BRs can be deployed. Our work focuses on 6LoWPAN networks running RPL, the Routing Protocol for Low-Power and Lossy Networks (LLNs), as it is the primary routing protocol for IPv6 LLNs [11]. We analyze in detail how RPL networks support multiple BRs both in theory and in practice. We design a RPL 6LoWPAN Border Router (RPL-6LBR) solution supporting multiple BRs with node mobility and fault-tolerance without compromising the energy-efficient control mechanisms provided by RPL. Finally, we implement RPL-6LBR in the Contiki operating system and evaluate its performance.

2 Background

6LoWPAN and IPv6 operate on link layers that are fundamentally different: IPv6 link layers such as Ethernet are highly-available, provide high throughputs with low latencies and are more reliable than LLNs. As a consequence, control-plane protocols for IPv6 and 6LoWPAN networks do not share the same constraints. In this section, we describe how RPL allows for low-overhead networking in support of LLNs, and we further discuss how existing standards propose to interconnect IPv6 networks and 6LoWPAN.

2.1 RPL-based 6LoWPAN Networks

The Routing Protocol for LLNs (RPL) is a distance-vector routing protocol addressing the problem of scalable any-to-any routing in low-power IPv6 networks [11]. RPL routes upwards and downwards along a tree-like topology called Destination-Oriented Directed Acyclic Graph (DODAG), maintained via an adaptive beaoning mechanism. RPL provides good scalability with the network size because all the control traffic is focused in optimizing a single topology towards a single destination, rather than one route between all pairs of nodes in the network. It also favors data-driven link estimations, meaning it will not proactively waste energy attempting to optimize its DODAG unless data packets are routed, which is attractive in terms of power savings. The main innovation of RPL is that it brings bidirectional traffic patterns over lossy, multihop networks.

2.2 Interconnecting LLNs with IPv6

Figure 1 represents a typical network architecture integrating 6LoWPAN and IPv6 networks. The BR may act as a router and separate the IPv6 and 6LoWPAN networks in different subnets, or act as a bridge, handling both IPv6 and 6LoWPAN networks as a common subnet.




Fig. 1: Border Router in a typical IPv6-6LoWPAN environment

These architectures support end-to-end IP communication between any host and a sensor on the LLN. To achieve this goal, 6LoWPAN transmits IPv6 packets over IEEE 802.15.4 frames by adding an adaptation layer for header compression and fragmentation [5]. For IPv6 control plane mechanisms, like the Neighbor Discovery Protocol (NDP) [6], to operate on energy constrained and asymmetric networks, the IETF proposes the Neighbor Discovery Optimization for LLNs (6LoWPAN-ND) [8].

6LoWPAN-ND aims to support both proposed approaches to routing in 6LoWPAN networks: route-over and mesh-under modes. However, 6LoWPAN-ND does not fully encompass the problem of interconnecting 6LoWPAN-based WSNs to the Internet, only interconnection by routing is explored.

To allow transparent communication as if the LLN is part of the IPv6 subnet, the Neighbor Discovery Proxy Gateway for 6LoWPAN-based WSNs (6LP-GW) provides bridging capabilities and allows an external host to discover and connect to sensors using NDP and protocol conversion between NDP and 6LoWPAN-ND [1]. However, that proposal does not cover the 6LoWPAN route-over scenario and its interaction with RPL.

2.3 Dealing with Mobility and Redundancy

Dynamic routing protocols like RPL are designed to adapt their routing strategies against varying channel conditions. As a side effect, an immobile sensor can appear to be mobile. RPL has built-in mechanisms to support inter-LoWPAN mobility thanks to the multi-instance support and the capability to switch from one DODAG to another in the same instance.

6LoWPAN-ND does not have any requirement on inter-LoWPAN mobility in route-over networks or on multi-BR support. This is out of scope of [8]. However, [8] requires that either sensors register their address towards all of the known 6LBRs or that an out-of-scope mechanism synchronize those 6LBRs. Again, no

current 6LoWPAN-ND implementation proposes any synchronization mechanisms. Thus, inter-LoWPAN mobility and redundancy is severely constrained as a WSN node moving from one DODAG to another will see the subnet prefix changing and therefore will get a new global address. If each LLN is configured with its own prefix, Mobile IPv6 or one of its variants could be implemented [7].

3 Design of a RPL-compatible 6LBR

In order to tackle the shortcomings identified in the previous sections, we propose to define the routing and mobility requirements for multi-BR RPL-based networks as follows:

1. The BR must be compatible with RPL in a route-over mode.
2. Multiple BRs must be able to co-exist in the same infrastructure, handling contiguous or disjoint LLNs.
3. Only one BR at a time should be responsible for a given 6LoWPAN host.
4. LLNs connected to different BRs should be able to share the same prefix.
5. A 6LoWPAN host should be able to maintain its global IPv6 address when switching between BRs that share a common prefix.
6. Multiple WSNs and an adjacent Ethernet segment can share the same prefix.

These requirements form the basis of our RPL-6LBR design, described below.

3.1 ND Proxy for RPL-6LBR

If the LLNs share the same prefix, one could naïvely imagine that it is enough to forward packets between the two interfaces. However, interoperability problems arise as the IPv6 hosts will use NDP to resolve addresses whereas the hosts on the WSN side will use routes and default routes provided by RPL to communicate. The RPL-6LBR must provide an adaptation mechanism, handling NDP requests and converting NDP configuration options to RPL options. The concept of ND-Proxy is not a new one, a similar concept has been defined in RFC 4389 ND Proxy [9], but it is mainly targeted at proxifying wireless or PPP based bridges and not applicable as such for RPL-6LBR. ND Proxy has also been implemented for adapting 6LoWPAN-ND with NDP [1], and proposed for RPL [10]. In our case we propose to bridge a RPL network with a NDP network; the information gathered through RPL messages is used for the ND proxy functionalities and the configuration parameters received through NDP configure the RPL network.

3.2 Mobility Support for Redundant 6LBRs

Aggregating several WSNs and allowing sensor mobility across them is not currently defined in 6LoWPAN-ND or RPL. Following the description of the possible LLNs topologies, aggregation of LLNs can be done at different levels, depicted in Figure 2. It should be noted that redundancy can be seen as an extreme case of WSN aggregation, where the two WSNs are identical.

Aggregating WSNs by synchronizing RPL-6LBR. Aggregating WSNs using ND Proxy based RPL-6LBR sharing the same prefix is feasible but requires the synchronization of the different RPL-6LBRs to allow sensor mobility. Otherwise, several RPL-6LBRs would perform ND Proxy for the same WSN node, which is undesirable. In this paper we rely on the Neighbor Advertisements to synchronize the ND Proxies of the connected RPL-6LBR.




Fig. 2: Left: Routing level aggregation, each BR is a RPL Root with its own DODAG. Right: RPL level aggregation, the RPL Root runs on the backbone.

Aggregating WSN using an External DODAG Root. RPL integrates sensor mobility deeply, so an innovative solution we propose is to create a unique DODAG across multiple BRs encompassing all the WSN nodes. This unique DODAG is managed by an External DODAG Root connected to the BRs through an Ethernet backbone.

4 Implementation

We provide an implementation of our RPL-6LBR in the latest development version of the Contiki operating system⁴. We support multiple types of RPL-6LBR platforms: the RaspberryPi with a custom 802.15.4 radio, the BeagleBone, the Econotag running Contiki natively and a standard Linux PC.

To enable multiple BR support in Contiki, we implement an ND Proxy as designed in Section 3.1 to adapt IPv6/NDP and 6LoWPAN/RPL. We also implement prefix distribution by leveraging the Prefix Information Option (PIO) of the RPL DIO broadcast messages combined with a global reset of the DIO Trickle timer. Due to space limitations we refer the reader to the online documentation for implementation details.

5 Evaluation

We evaluate RPL-6LBR in different scenarios with a Python-based unit test framework, provided as part of the RPL-6LBR open-source repository.

As our implementation targets real deployments, the results in this paper are partially made up of tests on the RaspberryPi, interfaced with real sensor

⁴ <http://cetic.github.com/6lbr>. tag 6lbr-1.1.0. 2013 (licence: BSD)

nodes of our testbed: 20 nodes spread over a $500m^2$ office space (10 TelosB, 10 Zolertia Z1). In order to explore certain parameters in a controlled manner, we also parallelized a large number of test instances on Linux virtual machines, interfacing the same RPL-6LBR with a virtual WSN using the COOJA simulator with TelosB emulated sensor devices.

All sensor modules run a Contiki application over the ContikiRPL implementation using the most commonly used configuration: the Minimum Rank Hop Objective Function (MRHOF) with ETX [2] as the routing metric and with RPL’s storing mode for downwards routing. We control the network size, topology, RPL DAG stabilization times and types and amount of cross-traffic. All runs consist in analyzing the effects of cross-traffic, node mobility and prefix re-configuration on bidirectional application traffic between a host on the Ethernet network and a constituent of the WSN network, as far down the RPL DAG as possible. Each testcase was executed 10 times, with different random seeds for the simulated ones.

Prefix Distribution We measure the lapse of time during which a particular host on the WSN is unreachable on a new prefix while this prefix is propagated on a running network. Figure 3 shows the average duration of unreachability versus the distance from the RPL Root for different traffic patterns. The main result is that there is little effect of the cross-traffic, prefix distribution being only dependant of the beaconing mechanisms of RPL.

Inter-WSN routing We next evaluate inter-WSN routing between two RPL-6LBRs set up in both SmartBridge and Switch modes, handling physically disjoint WSNs. Ping messages are exchanged between elements of disjoint DODAGs (garanteed in simulation), each roughly 3 hops away from its root. In SmartBridge mode, the Round Trip Times (RTT) of pings average $0.62s \pm 0.12s$, which is slightly shorter than the ones in the Switch mode measured as $0.67s \pm 0.12s$ due to the extra hop that is induced in Switch mode.

Node mobility On Figure 4, we observe the behavior of RPL-6LBR in the presence of a mobile node switching between DODAGs. We see that the performance is dependant of RPL’s link metric estimation. High data collection rates update the mobile node’s ETX rapidly enough to allow it to select a member of the secondary DODAG as preferred parent in all cases. On the other hand, low data collection rates do not all result in the reachability of the node within a 600s limit we imposed, and the trials that do converge exhibit a longer convergence time on average.

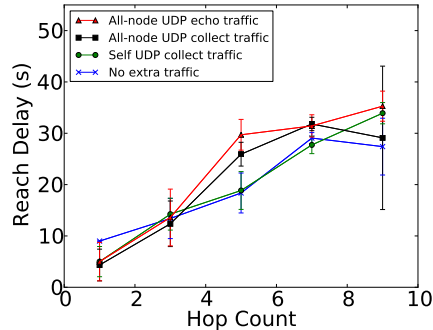


Fig. 3: Time of propagation of new prefix vs. hop count




Fig. 4: Mobile node trials. The success rate and unreachability duration is independent of whether or not the other nodes of the network update their link metrics with periodic unicasts.




Fig. 5: BR outage trials. The success rate and unreachability delay are improved when all elements of the network accelerate their ETX updates as they emit periodic UDP messages.

Border Router Outages In the case of BR outages, the effects of data-traffic is amplified (Figure 5). A BR failure orphans the entire sub-DODAG originally handled by the BR and requires nodes closest to the alternate DODAG to update their link ETX first, then join the DODAG and announce their ranks to neighboring orphans via DIOs. The remaining orphans of the previous DODAG gradually attach themselves to the new DODAG in a hop by hop manner.

A close look at our test results reveals that routing loops appeared among the orphaned elements of the failed cases, which prevents ETX from rising since unicasts keep getting acknowledged within the routing loop. This is still an open problem in ContikiRPL, as it is capable of detecting routing loops using the IPv6 Hop-by-Hop Option for RPL [3] but does not implement any local repair.

6 Discussion

Our evaluation covers functional and performance aspects of the RPL-6LBR implementation. We started by showcasing the prefix distribution functionalities, which is an important deployment aspect for real-world WSNs. Our results shows that the principle of a RPL ND-Proxy to interconnect IPv6 networks and RPL networks is sound and reliable. The synchronization mechanism, based only on existing standards and messages, is proved to be functional and allows fast transition of one sensor node from one WSN to another.

Performance-wise our measurements show a heavy influence of ETX-based link estimation on RPL hosts' reactivity to change, in both mobility and BR outages scenarios. Low-overhead link estimation is desirable for long-running battery-operated sensor networks, and it is important to note that the solution we propose does not compromise this design.

7 Conclusion

This work introduces a new 6LoWPAN border router design and implementation which leverages existing standardization efforts while complying with the reactive link estimation mechanisms of RPL. We also propose solutions for redundant BR deployments and assess RPL-6LBR in the presence of mobile nodes

and BR failures. The results show that RPL-6LBR is viable for redundant BR deployments, but suggest that a different choice of Objective Function would be needed for critical systems.

In future work, multiple RPL instances will accommodate different types of routing QoS to harness the flexibility of RPL. We will also address larger WSN deployments by extending our testbed and integrating multiple RPL-6LBRs on larger, existing testbeds and real-world deployments.

Acknowledgment

This work was partly funded by the Walloon Region under the First-DoCA funding number 1017211. The authors would also like to thank Maxime Denis from UMons for his valuable contributions.

References

- [1] L. Maqueda Ara. “Neighbor Discovery Proxy-Gateway for 6LoWPAN-based Wireless Sensor Networks”. MA thesis. KTH Information and Communication Technology, 2011.
- [2] D. S. J. De Couto et al. “A high-throughput path metric for multi-hop wireless routing”. In: *Proceedings of the International Conference on Mobile Computing and Networking (ACM MobiCom)*. San Diego, CA, USA: ACM, 2003, pp. 134–146.
- [3] J. Hui and J.P. Vasseur. *The RPL Option for Carrying RPL Information in Data-Plane Datagrams*. RFC 6553. Mar. 2012.
- [4] Isam Ishaq et al. “IETF Standardization in the Field of the Internet of Things (IoT): A Survey”. In: *Journal of Sensor and Actuator Networks* 2.2 (2013), pp. 235–287. ISSN: 2224-2708. DOI: 10.3390/jsan2020235.
- [5] G. Montenegro et al. *Transmission of IPv6 Packets over IEEE 802.15.4 Networks*. Internet proposed standard RFC 4944. Sept. 2007.
- [6] T. Narten et al. *Neighbor Discovery for IP version 6 (IPv6)*. RFC 4861. IETF, Sept. 2007.
- [7] Aleksandar Obradovic and Gorica Nikolic. “Overview of Mobility Protocols Features for 6LoWPAN”. In: *YUINFO 2012*.
- [8] Z. Shelby et al. *Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)*. Standard Track 6775. IETF, Nov. 2012.
- [9] D. Thaler, M. Talwar, and C. Patel. *Neighbor Discovery Proxies (ND Proxy)*. Experimental RFC 4389. Apr. 2006.
- [10] P. Thubert. *6LoWPAN Backbone Router*. IETF Draft. 2013. URL: <http://datatracker.ietf.org/doc/draft-thubert-6lowpan-backbone-router/>.
- [11] T. Winter (Ed.), P. Thubert (Ed.), and RPL Author Team. *RPL: IPv6 Routing Protocol for Low power and Lossy Networks*. Mar. 2012.