Comparative Study of RPL-Enabled Optimized Broadcast in Wireless Sensor Networks

Thomas Clausen, Ulrich Herberg
Laboratoire d'Informatique (LIX) – Ecole Polytechnique, France
Thomas@ThomasClausen.org, Ulrich@Herberg.name

***Abstract*—Recent trends have suggested convergence to Wireless Sensor Networks (WSNs) becoming IPv6-based. To this effect, the Internet Engineering Task Force has chartered a Working Group to develop a routing protocol specification, enabling IPv6-based multi-hop WSNs. The current effort of this working group is development of a unicast routing protocol denoted RPL. RPL constructs a "DAG-like" logical structure with a single root, at which the majority of the traffic flows terminate, and assumes restrictions on network dynamics and traffic generality, in order to satisfy strict constraints on router state and processing.**

This paper investigates the efficient network-wide broadcast mechanisms in WSNs, using the logical structure already established by RPL. The aim hereof is to impose minimal additional state requirements on WSN routers, beyond that already maintained by RPL. This paper presents a selection of such broadcast mechanisms for RPL routed WSNs, and evaluates their performances. As part of this evaluation, the paper compares with MPR Flooding – an established efficient flooding optimization, widely used in MANETs.

I. INTRODUCTION

The general context for routing in Wireless Sensor Networks (WSNs) is small, cheap devices whose primary function is data acquisition, and for which communications capabilities are a "commodity to their primary function" – a necessary, but in preference unobtrusive, functionality, specifically targeted to the precise goal which the WSN is deployed to satisfy. As an example, a WSN deployed for environmental monitoring might contain a set of temperature sensors, sending "notifications" to a central controller when the temperature exceeds certain thresholds – and occasional "keepalive" messages otherwise, to let the controller know that the sensors are still operational. Traffic from the controller to the individual sensors may be limited to "setting the thresholds" – possibly rarely, such as at system deployment.

A. Routing Protocol for Low power and Lossy Networks

Recent trends have suggested convergence to WSNs becoming IPv6-based. To this effect, the ROLL working group of the IETF is currently specifying an IPv6-based unicast routing protocol for WSNs, denoted RPL (*"IPv6 Routing Protocol for Low power and Lossy Networks"* [1]). The basic construct in RPL is the DODAG — a destination oriented DAG, rooted in a "controller". In the converged state, each WSN router has identified a stable set of parents, on a path towards the "root" of the DODAG, as well as one among these as its *preferred parent*. Each router, which is part of a DODAG (*i.e.* has selected parents) will emit *DODAG Information Object* (DIO) messages, using link-local multicasting, indicating its respective *Rank* in the DODAG (*i.e.* its distance from the root according to some metric(s), in the simplest form hop-count). Upon having received a (number of such) DIO messages, a router will calculate its own rank such that it is greater than the rank of each of its parents, and will itself start emitting DIO messages. Thus, the DODAG formation starts at the root, and spreads gradually to cover the whole network. The root can trigger "global recalculation" of the DODAG by way of increasing a sequence number in the DIO messages.

1) RPL Operational Requirements: The minimal set of in-router state required in a WSN router running RPL is, (i) the identifier of the DODAG root, (ii) the address and rank of the preferred parent, (iii) the configuration parameters shared by the DODAG root and (iv) the maximum rank that the WSN router has itself advertised. For redundancy, a WSN router running RPL can maintain information describing additional parents, which may allow rapidly changing its preferred parent in case the former preferred parent becomes unreachable.

RPL control message generation is timer-based, with the root able to configure suitable back-off of message emission intervals using *trickle timers* [2].

2) RPL Traffic Patterns: "Upward paths" or "multipoint-to-point paths" from the sensors towards the controller are supported by installing the "preferred parent" in each WSN router as the next hop on the path towards the DODAG root. The DODAG root may in its DIO messages have advertised a set of *destination prefixes*, to which it provides connectivity. These prefixes can be used to populate the routing table in the WSN routers in the network, or a default-route via the preferred parent and the DODAG root can be installed.

"Downward paths" or "point-to-single-sensor paths" are supported by having WSN routers, which wish to be reachable, issue *Destination Advertisement Object* (DAO) messages. These propagate via parents towards the DODAG root, and describe which prefixes can be reached via which WSN router. Each intermediate WSN router, forwarding a DAO message towards the root, adds its address to a *reverse routing stack* in the DAO message, thereby providing the DODAG root with the ability to do source routing for reaching destinations in the WSN.

"Sensor-to-sensor paths" are as default in [1] supported by having the source sensor transmit via its default route to the root, which will add a source-route to the received data for reaching the destination sensor.

B. Problem Statement

RPL, as *currently* specified in [1], supports only unicast traffic. RPL does, however, not explicitly provide support for any form of "optimized broadcasting" – delivery of the same data packet to all routers in the WSN. One important application of broadcasting in a WSN is for a controller to request that all sensors in the WSN transmit their sensor information – *e.g.* to verify if an alarming condition, signaled by a single sensor, is confirmed by other sensors in the WSN.

While such a "broadcast" could be accomplished by the DODAG root performing "bulk-unicast" to all sensors in the network, this is hardly efficient due to redundant transmissions of the same packets. Thus, this paper investigates ways of providing a reasonable optimized broadcast capability for an RPL routed network.

C. Paper Outline

The remainder of this paper is organized as follows: section II suggests a selection of different mechanisms for, by way of using the data structures and topologies already maintained by RPL, providing support for broadcast traffic in a WSN, and also briefly presents MPR Flooding. Section III provides a comparative performance study of the suggested broadcast mechanisms. Section IV concludes this paper.

II. DATA BROADCASTING IN RPL

This section suggests mechanisms for exploiting the DODAG, as constructed by RPL, in order to undertake better-than-classic-flooding for WSN-wide broadcasting. The fundamental hypothesis for these mechanisms is that all broadcast operations are launched from the root of the DODAG. If a sensor needs to undertake a network-wide broadcast, the assumption is that this broadcast is sent to the root using unicast, from where the DODAG root will launch the broadcast operation.

A. Classic Flooding (CF)

A common baseline for broadcast operations is that of classic flooding: each router relays a broadcast packet upon its first receipt by that router; subsequent receipts of the same packet are suppressed and do not cause retransmissions. This has to its merit that no control traffic is required – however also entails (i) that each data packet must be uniquely identifiable (*e.g.* by embedding a unique sequence number for a given source), (ii) that each router must maintain state for each already received and relayed data packet so as to enable suppression of duplicates, and (iii) each data packet is retransmitted by each router in the network – often with a large degree of redundant transmissions as consequence.

Redundant retransmissions cause increased battery drain, both when transmitting and receiving (and discarding) the redundant packets, and increase contention on the wireless media, increasing the probability of data loss due to collisions. CF is, for these reasons, not suggested as a mechanism for data broadcast in WSNs, but is described here as a baseline for data broadcasting in RPL.

B. MultiPoint Relay Flooding (MPRF)

A common improvement over Classic Flooding is for each router to select and designate a subset of its neighbors (Multi-Point Relays – MPRs [3]) for relaying broadcast transmissions, thereby reducing the number of redundant retransmissions of each packet. This has been shown to offer dramatic reductions in the network load (fewer transmissions), as well as a dramatic reduction in data loss due to collisions [4].

In order for MPRF to work, a router must select its MPRs such that a message relayed by these MPRs will be received by all routers two hops away. To this end, each router must maintain, at a minimum, state describing both its neighbor routers, as well as its 2-hop neighbors. MPRF – as CF – requires identification of each broadcast packet, and maintenance of state allowing elimination of duplicate packets.

MPRF is a common approach in wireless ad hoc networks, where it is used *e.g.* for network-wide broadcast of routing protocol control traffic by [5], [6] and [7] – as well as for network-wide data broadcast [8]. Comparing RPL-specific broadcast mechanisms with MPRF is therefore, to a certain extend, a comparison with "state of the art" of broadcasting in wireless multi-hop networks.

C. Parent Flooding (PF)

Admitting the RPL "philosophy" of data transmission to sensors originating at (or relaying via) the DODAG root, RPL lends itself to a first and simple broadcast optimization: restricting a RPL router to retransmit only broadcast packets received from a "parent". Logically, the basic performance hereof should be similar to that of classic flooding: with the broadcast operation initiated from the DODAG root, each router will retransmit the packet upon receipt from a parent. PF does not require any additional control traffic over that which is caused by RPL. PF may apply identification of each broadcast packet, and maintenance of state allowing elimination of duplicate packets in order to avoid multiple retransmissions of the same packet received from different parents – similar to MPRF and CF.

D. Preferred Parent Flooding (PPF)

In order to not incur any additional in-router state requirements for detecting and suppressing retransmission of duplicate packets, preferred parent flooding utilizes the existing relationship between RPL routers, in order to ensure that no router will forward a broadcast packet more than once. Each RPL router is required to select exactly one Preferred Parent. Restricting retransmissions of broadcast packets to only those received from the router's preferred parent ensures that duplicates received from other routers are ignored for retransmission.

E. Preferred Parent MPR Flooding (PPMPRF)

PPF is fundamentally a derivative of the MPRF optimization, attempting further to decrease the number of retransmissions necessary for a network wide broadcast. The idea is as follows: each router, selected as “Preferred Parent”, must designate a subset of its “selectees” (children which have selected it as preferred parent) as “Preferred Children”. These “Preferred Children” must be selected such that a message, relayed by these “Preferred Children”, will reach all its “grand children” – *i.e.* the children of its “selectees”.

Whenever a router receives a data packet that is to be broadcast throughout the network, that router will only then forward the packet if (i) at least one parent of that router has selected it as preferred child and (ii) the packet has not been previously received (as determined by a duplicated detection mechanism). It is to be noted that it is not sufficient to restrict forwarding to packets received from the preferred parent of a router, but that packets from *any* parent of that router have to be forwarded if the router has been selected as preferred child by *at least one* of its parents. The rationale is illustrated in figure 1. Assuming that the root of the network (router *0*) has selected *B* as preferred child as indicated by the downward arrow, and *B* forwards a packet originating from the root. If forwarding was restricted to packets received from the preferred parent of a router, *D* would not forward the packet from *B* (since it is no preferred parent of *D*), and thus *X* would never receive the packet.

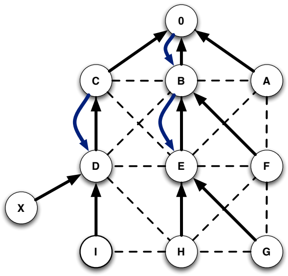

Figure 1. PPMPRF: Example showing the need to forward packets not only received by the preferred parent, but by any parent if the router is selected as preferred child by at least one of its parents. Upward arrows depict preferred parent selection, downward arrows preferred child selection.

Compared to “classic” MPR selection, the “Preferred Children selection” (i) concerns only coverage of “grand children” (*i.e.* “downward” in the DODAG as constructed by RPL) and (ii) is restricted by the preferred parent selection from RPL.

This restriction entails less liberties with respect to selecting relays for “best 2-hop coverage”. It is quite possible that the child providing the “best” coverage of a router has not selected that router as Preferred Parent, and that therefore PPMPRF will result in more relays than MPRF. In RPL, the Preferred Parent selection is intended to optimize for “best upwards paths towards the DODAG root” (possibly according to some deployment specific optimization criteria), which may not coincide with what would be optimal for “best downwards coverage”.

The PPMPRF mechanism also requires that each router knows (i) which children have selected it as Preferred Parent (*i.e.* its *selectees*), and (ii) which routers are Preferred Children of these selectees. This information can be made available through adding an option to DIO messages, emitted by all routers running RPL.

F. Optimized Preferred Parent MPR Flooding (PPMPRF-opt)

This mechanism represents a small optimization over PPMPRF, in that it provides all neighboring routers with the same rank with information, encouraging coordinated Preferred Parent selection so as to try to reduce the number of routers selected as Preferred Parent. Thus, a router will select as its Preferred Parent among its parents, the one which most of its adjacent routers also have in their parent set. Given a tie, the parent which a majority of the adjacent routers have already selected as Preferred Parent will be chosen. Thus, in addition to the information indicated for PPMPRF, PPMPRF-opt requires all parents to be advertised.

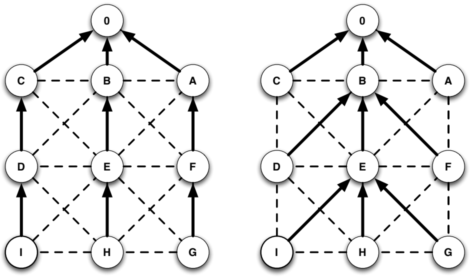

(a) (b)

Figure 2. Uncoordinated PP selection (a) and coordinated PP selection (b) in the same network. Solid arrows indicate the selection of a Parent as Preferred Parent; dotted lines the connectivity of the network.

Figure 2(a) depicts an example of Preferred Parent selection, as may happen in basic RPL: a router selects its Preferred Parent amongst all its parents with the lowest rank in an uncoordinated way. Worst case (in terms of redundant transmissions and therefore possible collisions when broadcasting), routers D, E and F all select different Preferred Parents (C, B, and A respectively). Similarly, I, H, and G may select three different Preferred Parents. For PPMPRF, this means that all routers, other than 0, will be selected as MPRs and thus retransmit a broadcast.

Figure 2(b) depicts a coordinated Preferred Parent selection. Router D will advertise all its parents (C and B) in its control messages, as will E (parents C, B and A) and F (parents B and A). D has an equal choice between parents C and B, and F has the same choice between B and A. E will select B as Preferred Parent because this is the only parent that both of its adjacent routers can also select as Preferred Parent. Once D and F receive a control message from E, advertising that B is selected as Preferred Parent, they will also select B. Thus, only routers B, E and H will be selected as Preferred Parents and therefore retransmit a broadcast.

Such coordinated Preferred Parent selection may be a double-edged sword for RPL. While it is a potential benefit for broadcast traffic from the DODAG root, unicast traffic flows towards the DODAG root via Preferred Parents. Thus, coordinated selection of Preferred Parents implies that unicast traffic is concentrated through a subset of the routers in the network, possibly increasing congestion in these routers, increasing the battery drain in these routers etc.

III. RPL BROADCAST PERFORMANCE STUDY

In order to explore the performance of RPL-enabled broadcast, simulations of MPRF, PPMPRF(-opt), PF and PPF have been performed using the Ns2 network simulator. The RPL protocol itself, providing the basic DODAG, used by PF and PPF, has been implemented in Java according to [1]. For MPRF and PPMPRF, the neighborhood discovery and MPR selection part of the Java based OLSRv2 implementation [9] has been used. The specific scenario settings are detailed in table I; for each datapoint in the results presented in this section, ten different scenarios have been simulated, with the results presented being the average from these runs.

Table I
NS2 PARAMETERS

Parameter	Value
Ns2 version	2.34
Mobility scenarios	No mobility, random uniform distribution of WSN routers
Grid size	variable
WSN router density	50 / km^2
Communication range	250m
Radio propagation model	Two-ray ground
Simulation time	200 secs
Interface type	802.11b
Radio frequency	2.4 GHz

For the purpose of this study, only a single RPL instance with a single DODAG is considered (but for each run with different, random positions of all routers). At the beginning of the simulation, only the root (which is the WSN router with the ID of 0) starts transmitting DIOs. Upon successful convergence of the DODAG, the root starts sending broadcast data with a data rate of 1280 bytes/s (64 byte long packets, sent every 50 ms).

In the following, the broadcast mechanisms presented in section II are analyzed in terms of MAC layer collisions, delivery ratio, overhead, delay, and path length. CF and PF (without duplicate detection) are not considered since their performance is expectedly much worse than any of the other mechanisms.

A. Basic Results

Figure 3 depicts the number of collisions of frames on the MAC layer, for the different broadcast mechanisms. MPRF and PPMPRF-opt yield the lowest number of collisions among the analyzed protocols, with PPMPRF causing about the same number of collisions as PF+DD (PF with duplicate packet detection). This is expected, as in MPRF, relays are explicitly selected so as to avoid redundant retransmission by topologically close routers, and the coordinated preferred parent selection in PPMPRF-opt also reduces the number of relays. PPMPRF without coordination entails more relays, as more routers in the network will be selected as preferred parents, which in turn select the relays (*i.e.* preferred children). In PPF, topologically close routers are likely to have chosen the same Preferred Parent and so will explicitly produce redundant retransmissions. Consider the example in figure 4, wherein a broadcast transmission is made by router 0 and relayed as indicated by the solid arrows. In PPF, as indicated in figure 4(a), each router will select its Preferred Parent and retransmit the packet once upon receipt from that preferred parent. Routers A, B and C all receive the transmission directly from router 0. Routers D, E and F have all chosen one of A, B and C as Preferred Parent and will thus all retransmit when receiving the transmission from their chosen preferred parent – similar for I, H, G, even though these three do not have any routers further down the network. In contrast, in figure 4(b), MPRs have been selected. Router 0 has selected B as MPR (as B "covers" D, E, F) and router B has chosen router E as MPR (as it covers all of G, H, I). As there are no further routers "below" in the network, router E has chosen no MPRs downwards. Thus, only B and E retransmit the broadcast packet from 0 – i.e. for each "level" in this simple network, only a single transmission occurs, with no collisions at each level.

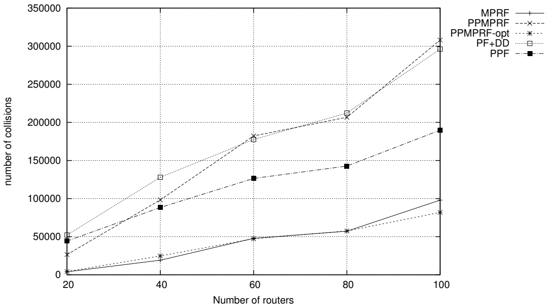

Figure 3. Broadcast: total number of MAC layer collisions

Figure 5 depicts the delivery ratio of broadcast packets. The delivery ratio of the MPRF and PPMPRF mechanisms are the highest of the compared broadcast mechanisms, with PPMPRF-opt being not much below MPR. This can be interpreted as a tradeoff between redundancy and efficiency: in relatively scarce networks (such as the simulated scenario) a higher redundancy of relays, such as in PPMPRF, can lead to a higher delivery ratio, despite of the increased number of collisions. as observed in figure 3. In dense networks, however, the large number of collisions with more redundant delays can reverse that effect and reduce the delivery ratio. A detailed analysis of the MPR relaying mechanism can be found in [3].

PF+DD has a higher delivery ratio than PPF, due to the redundancy of transmissions – when a router receives the same broadcast packet from several of its parents, chances are higher

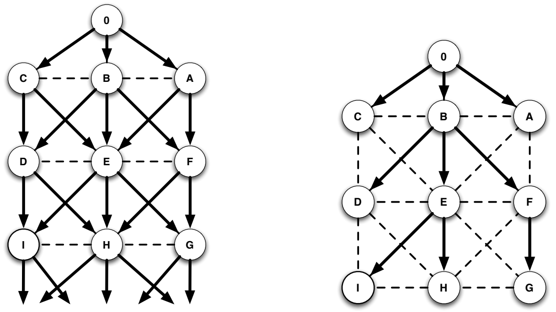

(a) PPF, originated by router 0

(b) MPRF, originated by router 0

Figure 4. PPF (a) and MPRF (b) in the same network. Solid arrows indicate transmission of a packet; dotted lines the connectivity of the network.

that at least one of the packets will reach the router, while if the one transmission from the preferred parent in PPF is lost due to a collision, the router will not forward the other incoming packets from its (non-preferred) parents. The higher delivery ratio of PF+DD is at the expense of vastly higher media load, as depicted in figure 6: the cumulative number of bytes transmitted during the simulations are significantly higher for PF+DD and for PPMPRF without the optimization. PPF incurs a lower overhead than PF+DD with MPRF still outperforming PPF by a large, and constant, margin. PPMPRF-opt has a similar overhead as MPRF.

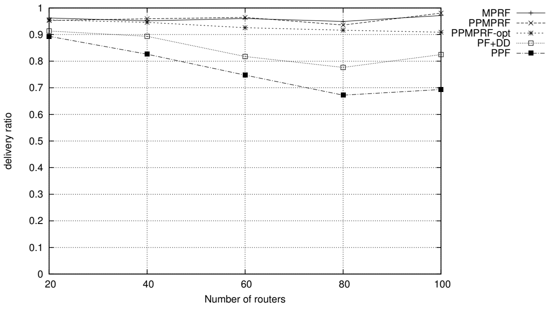

Figure 5. Broadcast: delivery ratio

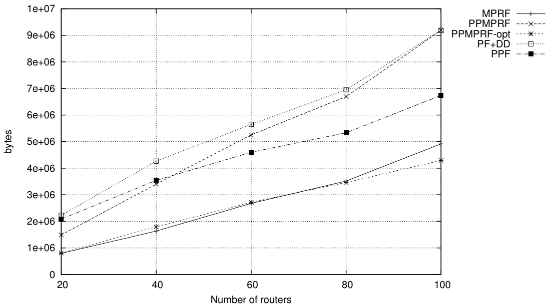

Figure 6. Broadcast: total retransmission overhead

Figure 7 depicts the average end-to-end delay for data traffic from the root to every WSN router in the network, and figure 8 depicts the average path length of successfully delivered data packets. The optimized MPR-based broadcast mechanisms incur the lowest delay of the protocols, while PPF causes a slightly lower delay than does PF+DD. The, on average, longer path lengths of MPRF are due to the data delivery ratio being higher – MPRF successfully "reaches" routers farther away from the root (as depicted in figure 9). It has been shown ([3]) that MPR leads to optimal path length. That means that every mechanism indicating a shorter path in the figure entails a lower reachability of routers further away from the broadcast source. Longer paths indicate suboptimal paths. It is worth observing that MPRF achieves the optimal path length with a lower delay still. This can in part be explained by the fact that MPRF ensures that data is flooded via shortest paths, and in part by the fact that with fewer retransmissions, less media and queue contention occurs.

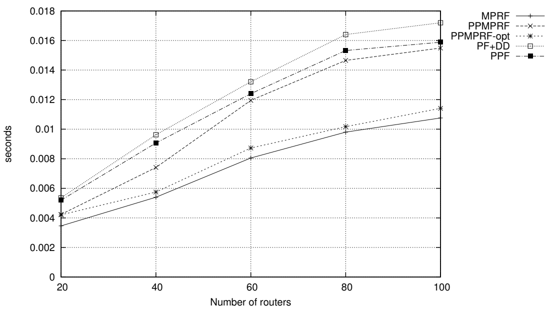

Figure 7. Broadcast: average delay

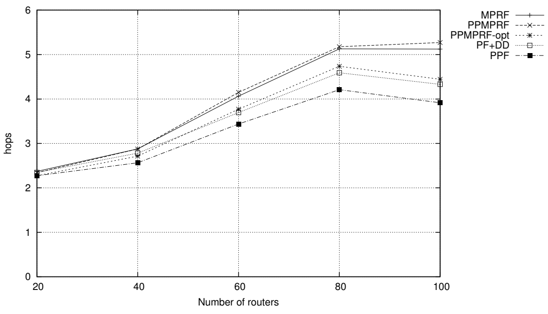

Figure 8. Broadcast: average path length

B. PPF with Jitter

In the results presented in section III-A, data traffic has been promptly forwarded by each WSN router, without explicit delay. As has been shown in [10], [11], adding a random jitter before retransmitting a broadcast packet can significantly reduce the number of collisions and, therefore, increase the delivery ratio for broadcast packets. In the following, the effect of adding jitter to PPF is investigated.

Figure 10 depicts the collision ratio of frames when using no jitter, and a random jitter uniformly distributed between 0 and 500 ms respectively. With jitter, the collision ratio is much lower than it is without. This is due to the fewer concurrent

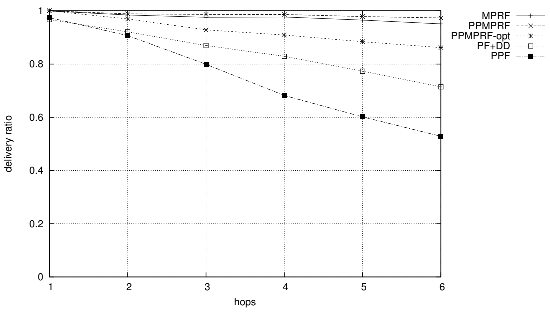

Figure 9. Broadcast: traffic delivery ratio with respect to distance from the root in hops (with 100 routers in the network)

retransmissions by adjacent WSN routers. Comparing to figure 3, PPF with jitter yields a collision ratio comparable to, or lower than, MPRF without jitter.

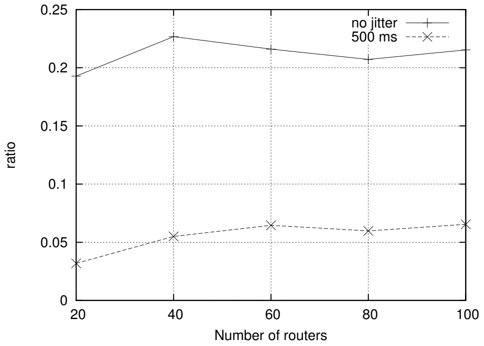

Figure 10. Collision ratio of PPF with jitter

As a consequence of the lower collision ratio, the delivery ratio of PPF with jitter is higher than it is without, as depicted in figure 11. Comparing to figure 5, the delivery ratio of PPF still remains consistently below that of MPRF, even when PFF is used with jitter.

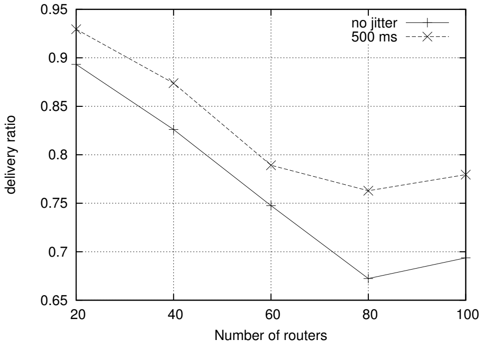

Figure 11. Delivery ratio of PPF with jitter

The drawback of using jitter is a higher end-to-end delay of packets, as depicted in figure 12. With jitter, the delay is considerably higher than it is without.

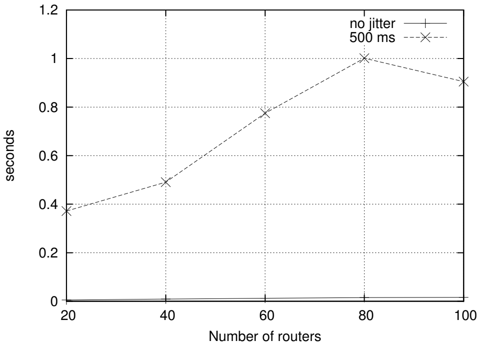

Figure 12. Average delay of PPF with jitter

IV. CONCLUSION

This paper presents a comparative study of broadcast mechanisms for RPL routed wireless sensor networks. Broadcast mechanisms, using the rooted DAG-like logical structure, maintained by the unicast routing protocol RPL, are introduced, and their performance studied. These broadcast mechanisms, denoted “Parent Flooding” (PF), “Preferred Parent Flooding” (PPF) and “Preferred Parent MPR Flooding” (PPMPRF) adhere to the “root-oriented” concept of RPL, in that all broadcast operations are to be initiated by the root of the DAG.

PF, PPF and PPMPRF are studied and compared by way of network simulations – and as a point of comparison, MPR Flooding (MPRF), known from wireless ad hoc networks, is also subjected to the same network scenarios in the simulator.

REFERENCES

[1] T. Winter, P. Thubert, RPL Author Team, “RPL: IPv6 Routing Protocol for Low power and Lossy Networks”, http://tools.ietf.org/html/draft-ietf-roll-rpl-10 (work in progress), June 2010

[2] P. Lewis, N. Patel, D. Culler, S. Shenker, “Trickle: A Self-Regulating Algorithm for Code Propagation and Maintenance in Wireless Sensor Networks”, http://csl.stanford.edu/ pal/pubs/trickle-nsdi04.pdf

[3] A. Qayyum, L. Viennot, A. Laouiti, “Multipoint relaying: An efficient technique for flooding in mobile wireless networks”, 35th Annual Hawaii International Conference on System Sciences (HICSS), 2001

[4] T. Clausen, L. Viennot, T. Olesen, N. Larsen, “Investigating Broadcast Performance in Mobile Ad-hoc Networks”, Proceedings of the IEEE conference on Wireless Personal Multimedia Communications (WPMC), October 2002

[5] T. Clausen, P. Jacquet, “Optimized Link State Routing Protocol (OLSR)”, Experimental RFC3626, http://www.ietf.org/rfc/rfc3626.txt

[6] T. Clausen, P. Jacquet, D. Nguyen, E. Baccelli, “SPF Multipoint Relay (MPR) Extension for Ad Hoc Networks”, Experimental RFC5449, http://www.ietf.org/rfc/rfc5449.txt

[7] T. Clausen, C. Dearlove, P. Jaquet, “The Optimized Link State Routing Protocol version 2 (OLSRv2)”, http://tools.ietf.org/html/draft-ietf-manet-olsrv2-11 (work in progress), April 2010

[8] J. Macker, “Simplified Multicast Forwarding”, http://tools.ietf.org /id/draft-ietf-manet-smf-10 (Work In Progress), March 2010

[9] U. Herberg, “JOLSRv2 – An OLSRv2 implementation in Java”, Proceedings of the 4th OLSR Interop workshop, October 2008

[10] T. Clausen, G. Hansen, L. Christensen, G. Behrmann, “The Optimized Link State Routing Protocol, Evaluation through Experiments and Simulation”, Proceedings of the IEEE conference on Wireless Personal Multimedia Communications (WPMC), October 2001

[11] T. Clausen, C. Dearlove, B. Adamson, “Jitter Considerations in Mobile Ad Hoc Networks (MANETs)”, Informational RFC5148, http://www.ietf.org/rfc/rfc5148.txt