



HAL
open science

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

Thomas Heide Clausen, Ulrich Herberg

► **To cite this version:**

Thomas Heide Clausen, Ulrich Herberg. Comparative Study of RPL-Enabled Optimized Broadcast in Wireless Sensor Networks. [Research Report] RR-7296, 2010, pp.20. inria-00488030v1

HAL Id: inria-00488030

<https://inria.hal.science/inria-00488030v1>

Submitted on 31 May 2010 (v1), last revised 6 Oct 2010 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

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

Thomas Clausen, Ulrich Herberg,

N° 7296

May 2010

 ***Rapport
de recherche***

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

Thomas Clausen*, Ulrich Herberg†,

Thème : COM – Systèmes communicants
Équipe-Projet Hipercom

Rapport de recherche n° 7296 — May 2010 — 20 pages

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 memorandum investigates the possibility for providing (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 memorandum presents a selection of such broadcast mechanisms for RPL routed WSNs, and evaluates their performances. As part of this evaluation, the memorandum compares with MPR Flooding – an established efficient flooding optimization, widely used in MANETs.

Key-words: RPL, Sensor Network, Broadcast

* LIX - Ecole Polytechnique, Thomas@ThomasClausen.org

† LIX - Ecole Polytechnique, Ulrich@Herberg.name

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

Résumé : Les tendances récentes suggèrent une convergence des réseaux de capteurs sans fils (WSNs ou Wireless Sensor Networks) vers IPv6. C'est pourquoi l'IETF (Internet Engineering Task Force) a mis en place un groupe de travail, chargé de spécifier un protocole de routage permettant aux réseaux de capteurs sans fil multi-sauts de fonctionner avec IPv6. Les efforts du groupe de travail se concentrent actuellement sur le développement d'un protocole de routage unicast appelé RPL. RPL construit une architecture logique de type DAG (graphe orienté acyclique) avec un nœud racine unique sur lequel se termine la majorité des flux de trafic, et suppose des restrictions sur les dynamiques du réseau et sur les types de trafic supportés afin de satisfaire les contraintes fortes des routeurs en terme d'états et de traitement.

Cet article examine la possibilité de fournir aux réseaux de capteurs sans fil des mécanismes efficaces de broadcast (diffusion), en utilisant la structure logique déjà proposée par RPL. Le but est ici de ne pas imposer d'exigences supplémentaires aux routeurs WSN fonctionnant déjà avec RPL. De tels mécanismes de broadcast pour les réseaux WSN utilisant le routage RPL, l'article en présente plusieurs et évalue leur performance. Dans le cadre de cette évaluation, ils sont comparés aux mécanismes reconnus et efficaces d'optimisation de diffusion du protocole MPR Flooding, largement utilisé dans les réseaux MANETs.

Mots-clés : RPL, Réseaux de capteurs, Diffusion

1 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, or even never as would be the case with factory set thresholds.

1.1 WSN Traffic Patterns

The communications requirements for WSNs are in contrast to “traditional networks”, wherein communications devices (network interfaces, switches, routers) have carrying data traffic as their sole raison d’être, and in which the devices do not make any a-priori assumptions as to the characteristics of the traffic they will be carrying. WSNs assume an a-priori knowledge of the traffic patterns to optimize for – with sensor-to-controller traffic (*multipoint-to-point*) being predominant, controller-to-sensor traffic (*point-to-multipoint*) being rare and sensor-to-sensor traffic being somewhat esoteric¹.

A-priori assumptions as to traffic patterns are used as the primary justification for “in which order to optimize” algorithm development, and to determine which delivery mechanisms to support in WSN routing protocols.

1.2 WSN Trade-off’s

Low-power consumption, minute physical sizes, low price-points and ruggedness against the environment are among the industrial or commercial keywords, often associated with wireless sensors – and which entail challenging constraints (in terms of the computational power, permanent and temporary storage and in the characteristics of the wireless interfaces) for designing routing algorithms. WSN routing protocols are therefore inherently compromises: trade-offs are made in adapting to the specific constraints under which they are to operate – the first sacrifice is usually “generality”. WSN routing protocols generally and narrowly consider only the traffic characteristics of their target environment as “valid” and discard all other traffic characteristics in the name of satisfying their operational constraints; two of the most common such constraints brought forward are strict bounds on in-router state and control traffic. A second sacrifice is often in route optimality: stretched (non-optimal) routing paths are an acceptable trade-off for lower control traffic from a routing protocol, with the hypothesis that traffic flows will be such that the impact of such stretched paths will be negligible.

¹Note that while this may be commonly assumed, this is not a universal distribution of traffic patterns in WSNs – there are scenarios in which sensor to sensor traffic is assumed a more common occurrence, such as [1].

1.3 WSNs and the IETF

The ROLL working group of the IETF is currently developing support for IPv6 routing in WSNs. The objective of this working group is to support WSNs:

- which comprise up to thousands of WSN routers,
- where the majority of the WSN routers have very constrained in-router resources,
- where the network to a large degree is “managed” by a (single or few) central “super-WSN routers”, and
- where handling mobility is not an explicit design criteria.

For the latter of these, it is assumed that network dynamics is a consequence of rare, interfering events (say, a sensor getting crushed or running out of battery) necessitating a correction, rather than tracking continuously changing conditions such as due to mobility.

The ROLL working group aims at supporting multipoint-to-point, point-to-multipoint and point-to-point traffic patterns. The emphasis among these traffic patterns is to *optimize for* multipoint-to-point traffic, to *reasonably support* point-to-multipoint traffic and to *provide basic features for* point-to-point traffic, in that order.

1.4 IPv6 routing for WSNs: RPL

RPL – the “*Routing Protocol for Low Power and Lossy Networks*” [2] – is an early proposal for an IPv6 routing protocol for WSNs by the ROLL working group in the IETF. The basic construct in RPL is the DODAG — a destination oriented DAG, rooted in a “controller”, illustrated in figure 1. 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 position – distance according to some metric(s), in the simplest form hop-count – with respect to the root). 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.

As a Distance Vector protocol, RPL contains rules, restricting the ability for a router to change its rank. Specifically, a router is allowed to assume a smaller rank than what it has previously advertised (*i.e.* to logically move closer to the root) if it discovers a parent advertising a lower rank (and it must, then, disregard all previous parents with higher ranks), while the ability for a router to assume a greater rank (*i.e.* to logically move farther from the root) in case all its former parents disappear, is restricted to avoid count-to-infinity problems. The root can trigger “global recalculation” of the DODAG by way of increasing a sequence number in the DIO messages.

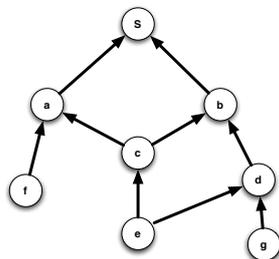


Figure 1: RPL Basic Construct: DODAGs

1.4.1 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 (notably, destination prefixes and message emission timers) 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 (up to and including all its 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* [3].

1.4.2 RPL Traffic Patterns

The DODAG constructed by RPL is intended to serve as a sort of “backbone topology” for carrying data traffic.

“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 belong to, and thus 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 [2] 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.

1.5 Problem Statement

RPL, as *currently* specified in [2], discusses and supports only unicast traffic, albeit two different variations hereof: sensor-to-controller traffic, supported by regular hop-by-hop routing, and controller-to-sensor unicast traffic, supported by way of source routing. Sensor-to-sensor routing is supported as a combination of these two: hop-by-hop routing to the “DODAG root”, then source routing from the DODAG root to the destination sensor.

Thus, RPL supports, at least functionally though not optimally, the WSN traffic patterns identified in section 1.1. 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 / trigger that all sensors in the WSN transmit their sensor information – *e.g.* to verify if an extreme/alarming condition, signaled by a single sensor, is confirmed by other sensors in the WSN, or for bulk-setting configuration parameters in all sensors.

While such a “broadcast” could be accomplished by the DODAG root performing “bulk-unicast” to all sensors in the network, this is hardly efficient. The two “classic” reasons for why this is inefficient apply: (i) when performing broadcasting via bulk-unicast, the same data is being transmitted repeatedly across the same links (especially, over links close to the source – in case of an RPL network, close to the DODAG root – thus creating potentially congested areas), and (ii) it requires that the source of the bulk-unicast knows all possible recipients for the broadcast. In the case of an RPL-routed WSN, two additional reasons apply. First, in order for the DODAG root to be able to reach all sensors by way of unicast transmissions, all WSN routers are required to generate DAO messages – and the DODAG root is required to maintain explicit paths to all sensors, for use as source-routes. Second, as root-to-sensor unicast may by default be source-routed, each such unicast packet explicitly contains the path it must take, increasing the size of each broadcast packet (and, so, the energy consumption and media overhead)

Thus, this memorandum investigates ways of providing a reasonable optimized broadcast capability for an RPL routed network.

1.6 Memorandum Outline

The remainder of this memorandum is organized as follows: section 2 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. MPR Flooding an optimized broadcast mechanism, widely employed in wireless ad hoc networks, and is introduced as a point of reference and comparison. Section 3 provides a comparative performance study of the suggested broadcast mechanisms. Section 4 concludes this paper.

2 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 broad-

casting. 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 – this is similar to the basic mechanism for sensor-to-sensor unicast in [2], wherein traffic from the source sensor transits to the DODAG root, for relaying to the destination sensor.

2.1 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 (commonly ensured by embedding a unique sequence number in each broadcast packet, emitted by a given source), (ii) that each router must maintain information (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.

2.2 MultiPoint Relay Flooding (MPRF)

A common improvement over Classic Flooding is for each router to select and designate a subset of its neighbors (MultiPoint Relays – MPRs [4]) 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 [5].

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, as illustrated in figure 2. To this end, each router must maintain, at a minimum, state describing both its neighbor routers, as well as its 2-hop neighbors (“neighbor routers of 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 [6], [7] and [8] – as well as for network-wide data broadcast [9]. Comparing RPL-specific broadcast mechanisms with MPRF is therefore, to a certain extent, a comparison with “state of the art” of broadcasting in wireless multi-hop networks.

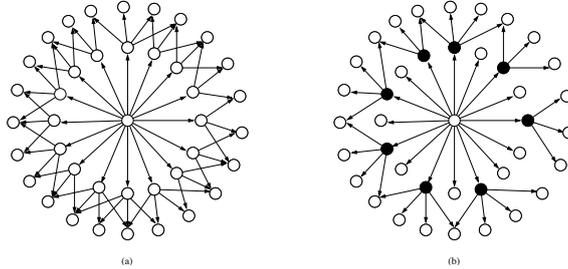


Figure 2: (a) Classic flooding and (b) MPR Flooding

2.3 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.

2.4 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 (parents or otherwise) are ignored for retransmission.

2.5 Preferred Parent MPR Flooding (PPMPRF)

PPF is fundamentally a “duplicate avoidance” mechanism, which can be combined with 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”.

Compared to “classic” MPR selection, this “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.

2.6 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.

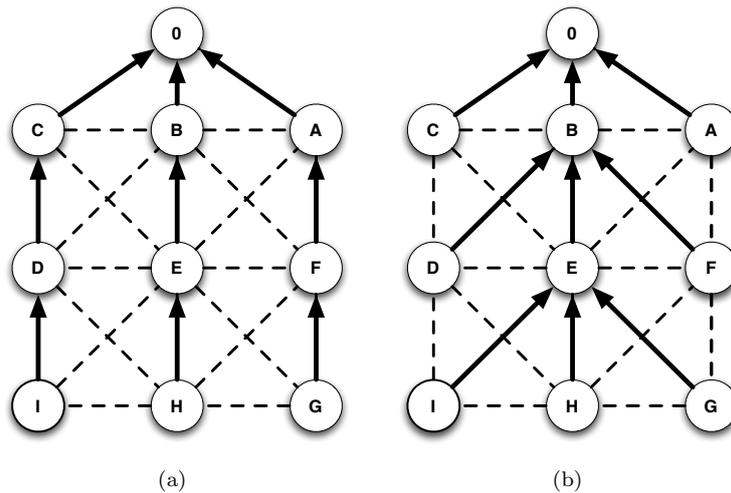


Figure 3: Uncoordinated PP selection (a) and coordinated PP selection (b) in the same network. Solid arrows indicate the selection of a Parent as Preferred Parent; dotted lines the connectivity of the network.

Figure 3(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 3(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.

3 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 [2]. For MPRF and PPMPRF, the neighborhood discovery and MPR selection part of the Java based OLSRv2 implementation [10] has been used (TC message generation disabled). The specific scenario settings are detailed in table 1; for each datapoint in the results presented in this section, ten different scenarios – each conforming to the same abstract scenario description – have been simulated, with the results presented being the average from these ten runs.

For the purpose of this study, only a single RPL instance with a single DODAG is considered, and only in a static network. 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 comparison, PF is analyzed in two versions, *with* and *without* duplicate packet detection (PF+DD). PPF and PPMPRF never apply duplicate packet detection, and MPRF always uses it.

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

Table 1: Ns2 parameters

3.1 Basic Results

Figure 4 depicts the maximum and the average rank of the routers in the DODAG, where the number represents the distance of a WSN router to the root in terms of hops (*i.e.* the maximum rank represents the diameter of the network, and the average rank represents the average over all WSN routers). The maximum and average ranks grow logarithmically with the number of WSN routers in the network.

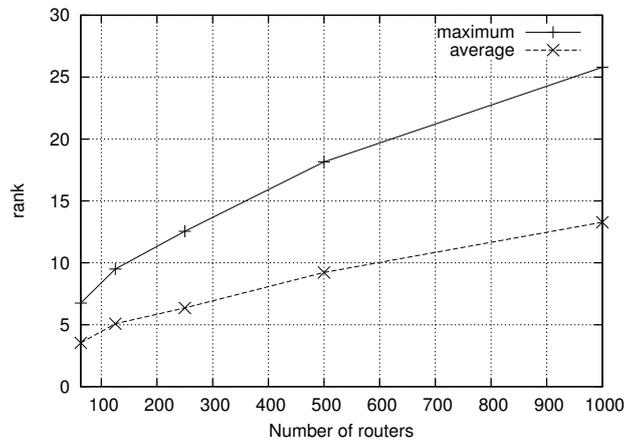


Figure 4: Maximum and average rank of the DODAG

Figure 5 depicts the average number of parents of each WSN router in the DODAG. Keeping the density of the network constant with increasing number of WSN routers, the average number of parents grows logarithmically.

Figure 6 depicts the collision ratio of frames on the MAC layer, for the different broadcast mechanisms. The MPR-based mechanisms yield the lowest collision rate among the analyzed protocols, while PPF yields a lower collision ratio than PF and PF+DD. This is expected, as in MPRF, relays are explicitly selected so as to avoid redundant retransmission by topologically close routers

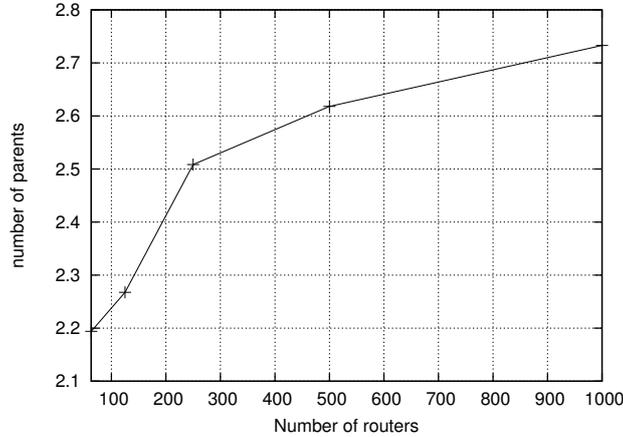


Figure 5: Average number of parents per WSN router in a DODAG

– whereas 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 7, wherein a broadcast transmission is made by router 0 and relayed as indicated by the solid arrows. In PPF, as indicated in figure 7(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 7(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 8 depicts the delivery ratio of broadcast packets. The delivery ratio for MPRF is higher than that of both PF (with and without DD) and PPF, and furthermore does not change significantly when scaling the network (note that the density of the network stays constant in the simulations). This is hardly surprising, considering the lower collision ratio (see figure 6) of MPRF. PPMPRF has a delivery ratio similar to that of PF+DD, with the optimized Preferred Parent selection (PPMPRF-opt) yielding a higher delivery ratio. 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 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 is at the expense of vastly higher media load, as depicted in figure 9: the cumulative number of bytes transmitted during the simulations are significantly higher for PF+DD. PPF incurs a lower overhead than PF+DD with MPRF still outperforming PPF by a large, and

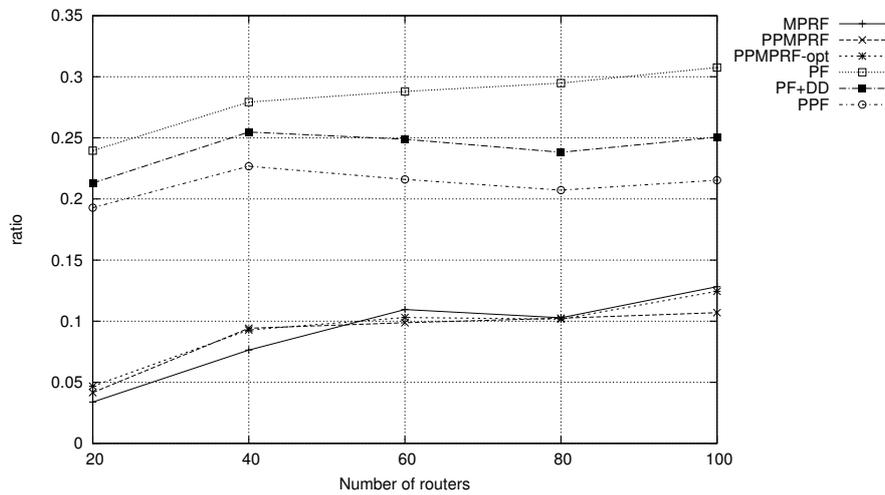


Figure 6: Collision ratio

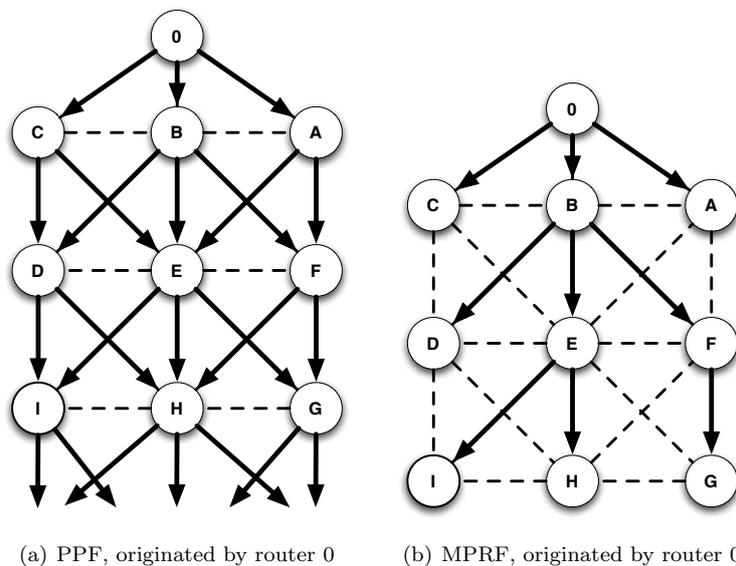


Figure 7: PPF (a) and MPRF (b) in the same network. Solid arrows indicate transmission of a packet; dotted lines the connectivity of the network.

constant, margin. The reason that PPMRF (and its optimized variant) has an even lower overhead than MPRF is that its delivery ratio is lower, which means that less routers are “reached” by the broadcast, and thus do not retransmit the traffic.

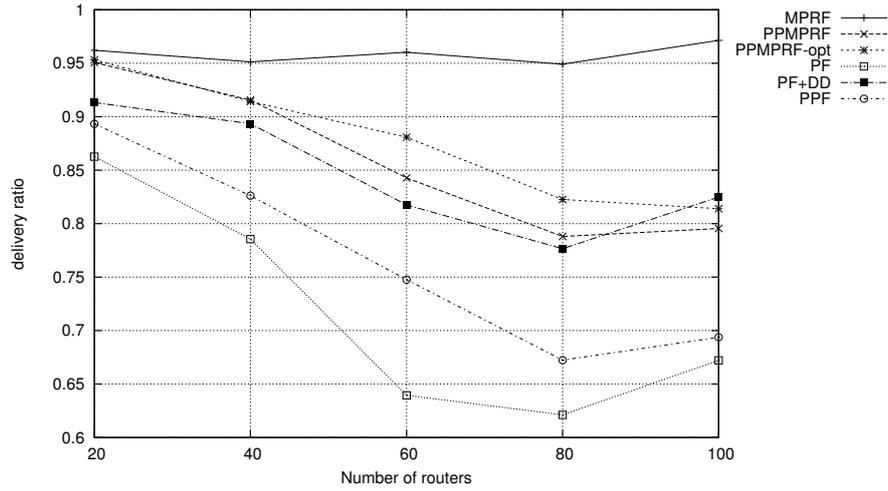


Figure 8: Delivery ratio

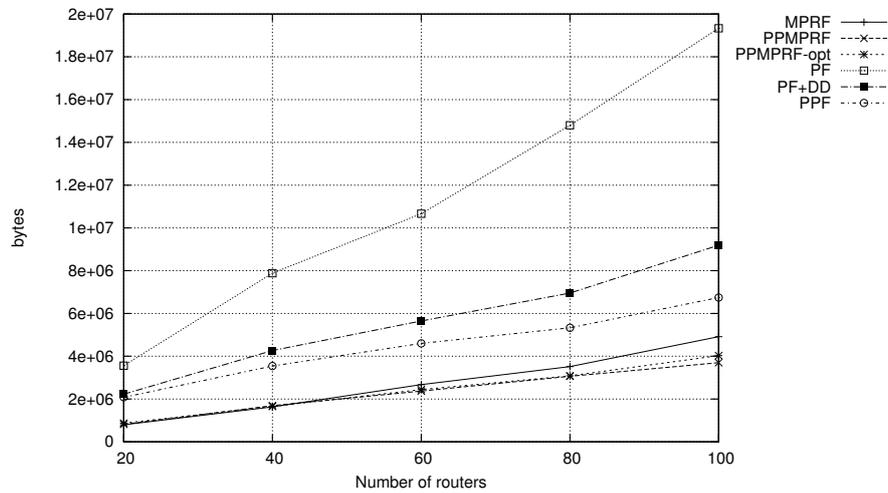


Figure 9: Total retransmission overhead

Figure 10 depicts the average end-to-end delay for data traffic from the root to every WSN router in the network, and figure 11 depicts the average path length of successfully delivered data packets. The MPR-based broadcast mechanisms incur the lowest delay of the protocols, while PPF incurs a much lower delay than does PF, and a slightly lower than does PF+DD. The, on average, longer path lengths of MPRF is due to the data delivery ratio being higher – MPRF successfully “reaches” routers farther away from the root (as

depicted in figure 12). It is worth observing that MPRF does so 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.

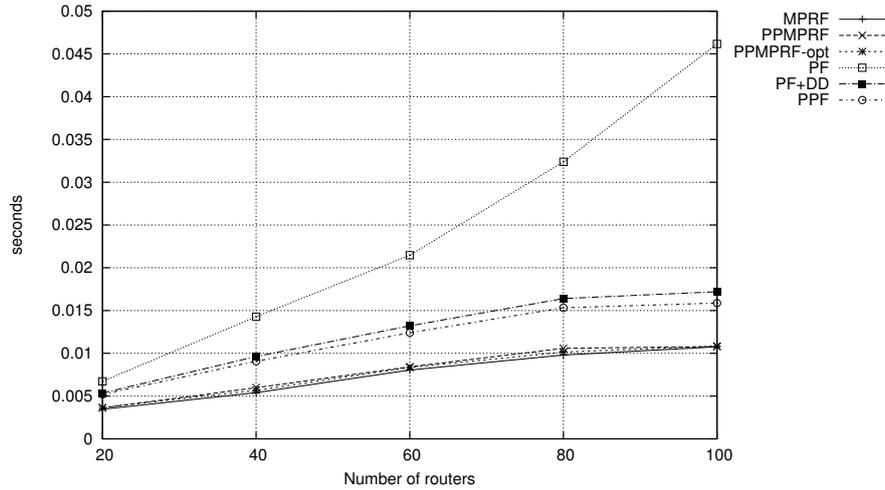


Figure 10: Average delay

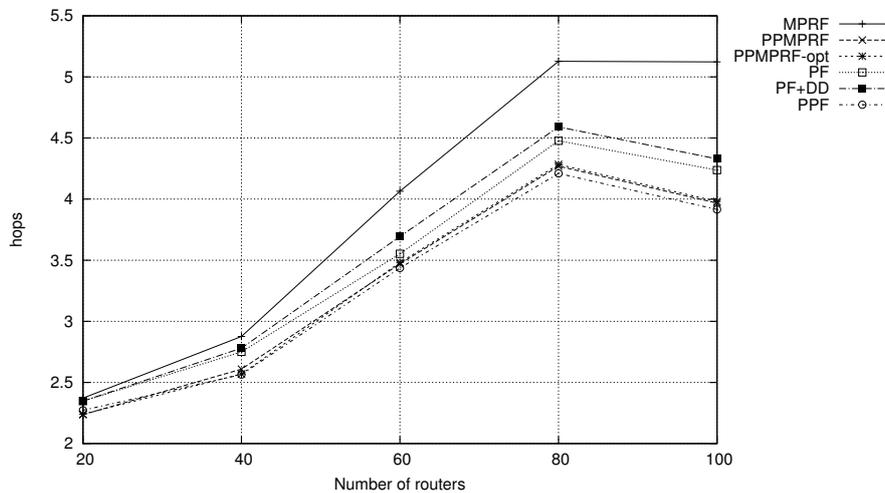


Figure 11: Average path length

3.2 PPF with Jitter

In the results presented in section 3.1, data traffic has been promptly forwarded by each WSN router, without explicit delay. As has been shown in [11, 12], adding a random jitter before retransmitting a broadcast packet can significantly reduce the number of collisions and, therefore, increase the delivery ratio

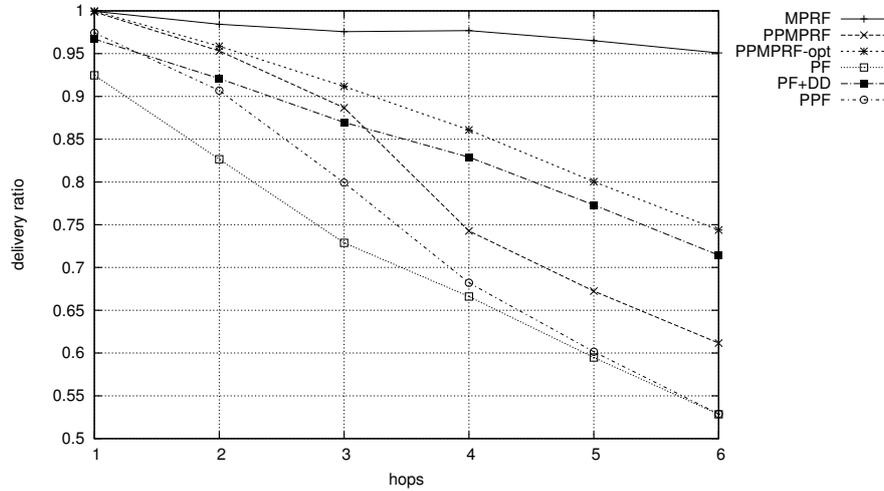


Figure 12: Traffic delivery ratio with respect to distance from the root in hops

for broadcast packets. In the following, the effect of adding jitter to PPF is investigated.

Figure 13 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 retransmissions by adjacent WSN routers. Comparing to figure 6, PPF with jitter yields a collision ratio comparable to, or lower than, MPRF without jitter.

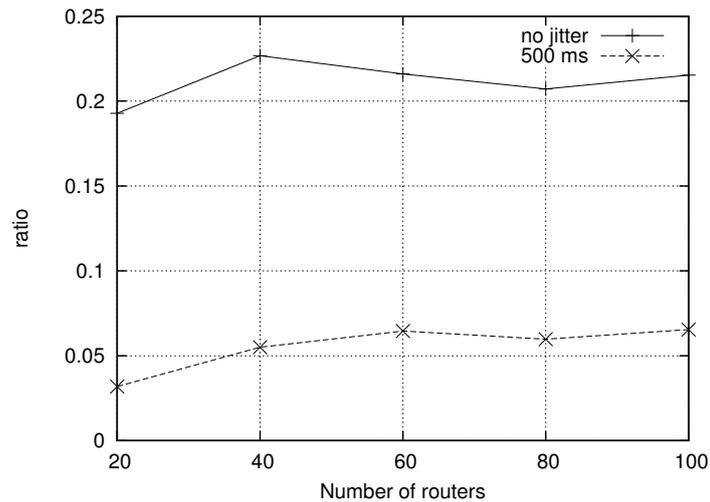


Figure 13: Collision ratio of PPF with jitter

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 14. Comparing to figure 8,

the delivery ratio of PPF still remains consistently below that of MPRF, even when PPF is used with jitter.

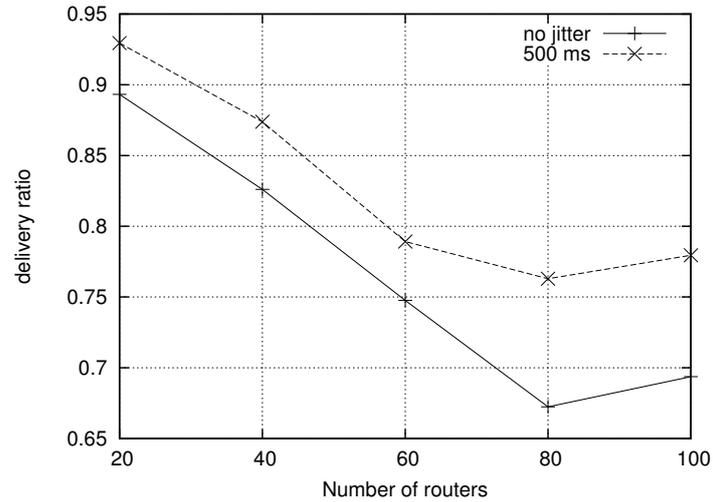


Figure 14: Delivery ratio of PPF with jitter

The drawback of using jitter is a higher end-to-end delay of packets, as depicted in figure 15. With jitter, the delay is considerably higher than it is without.

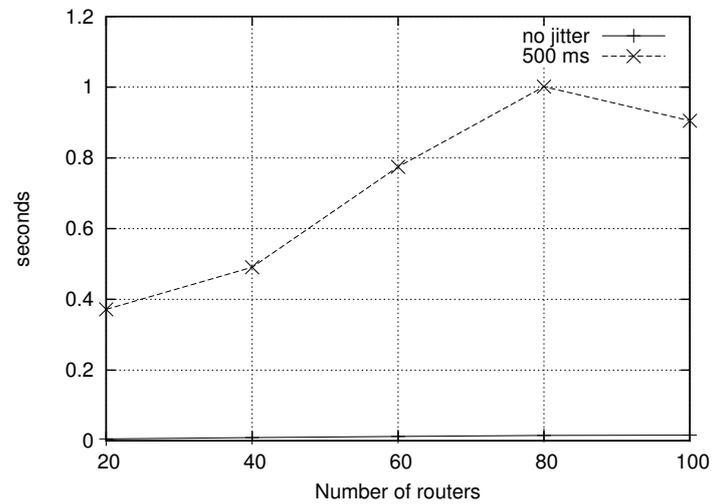


Figure 15: Average delay of PPF with jitter

4 Conclusion

This memorandum has presented 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. The simulations have shown that, not unexpectedly, PPF is more efficient than PF without duplicate packet detection (DD). When PF is used with a DD mechanism, the delivery ratio is higher than PPF due to redundant retransmissions; however, this redundancy naturally leads to a higher overhead on the radio channel. Furthermore, PPF does not require maintenance of extra state for storing packet sequence numbers for the DD mechanism.

Similar to PPF, PPMPRF does not require per-data-packet state on each router, however does require additional local topology information to be available, at least while selecting Preferred Parents. PPMPRF performs significantly better than PPF in terms of delivery ratio and traffic overhead. Using an optimized Preferred Parent selection mechanism, which coordinates the selection between adjacent routers, the performance of PPMPRF is observed to be further improved.

MPRF, however, remains much more efficient than PPMPRF, PPF, and PF, obtaining both a consistently lower bandwidth consumption, a consistently higher data delivery rate and consistently lower delays – even when reaching routers “farther away” from the originator of the broadcast transmission. On the potential downside, however MPRF requires additional state in each WSN router for duplicate packet suppression.

Compared to PPF, which requires only RPL signaling, MPRF and PPMPRF also necessitate a neighborhood discovery mechanism for selecting and designating MPRs – the price of a more efficient broadcast protocol.

References

- [1] J. Martocci et. al., "Building Automation Routing Requirements in Low Power and Lossy Networks", <http://tools.ietf.org/html/draft-ietf-roll-building-routing-reqs-09> (Work In Progress), January 2010
- [2] The ROLL Design Team, "RPL: Routing Protocol for Low Power and Lossy Networks", <http://tools.ietf.org/html/draft-ietf-roll-rpl-7> (work in progress), March 2010
- [3] 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>
- [4] 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
- [5] 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
- [6] T. Clausen, P. Jacquet, "Optimized Link State Routing Protocol (OLSR)", Experimental RFC3626, <http://www.ietf.org/rfc/rfc3626.txt>
- [7] 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>
- [8] 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
- [9] J. Macker, "Simplified Multicast Forwarding", <http://tools.ietf.org/id/draft-ietf-manet-smf-10> (Work In Progress), March 2010
- [10] U. Herberg, "JOLSRv2 – An OLSRv2 implementation in Java", Proceedings of the 4th OLSR Interop workshop, October 2008
- [11] 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, Aalborg, Denmark
- [12] T. Clausen, C. Dearlove, B. Adamson, "Jitter Considerations in Mobile Ad Hoc Networks (MANETs)", Informational RFC5148, <http://www.ietf.org/rfc/rfc5148.txt>

Contents

1	Introduction	3
1.1	WSN Traffic Patterns	3
1.2	WSN Trade-off's	3
1.3	WSNs and the IETF	4
1.4	IPv6 routing for WSNs: RPL	4
1.4.1	RPL Operational Requirements	5
1.4.2	RPL Traffic Patterns	5
1.5	Problem Statement	6
1.6	Memorandum Outline	6
2	Data Broadcasting in RPL	6
2.1	Classic Flooding (CF)	7
2.2	MultiPoint Relay Flooding (MPRF)	7
2.3	Parent Flooding (PF)	8
2.4	Preferred Parent Flooding (PPF)	8
2.5	Preferred Parent MPR Flooding (PPMPRF)	8
2.6	Optimized Preferred Parent MPR Flooding (PPMPRF-opt)	9
3	RPL Broadcast Performance Study	10
3.1	Basic Results	11
3.2	PPF with Jitter	15
4	Conclusion	18



Centre de recherche INRIA Saclay – Île-de-France
Parc Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 Orsay Cedex (France)

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Grenoble – Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Nancy – Grand Est : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex
Centre de recherche INRIA Paris – Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex
Centre de recherche INRIA Rennes – Bretagne Atlantique : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex
Centre de recherche INRIA Sophia Antipolis – Méditerranée : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399