

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 Wireless Sensor Networks. The current effort of this working group is development of a uni-cast 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 possibility for providing (efficient) network-wide broadcast mechanisms in WSNs, using the logical structure already provided by RPL. The aim hereof is to not impose any additional state requirements on WSN routers already running RPL. This paper presents two 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, or even never as would be the case with factory set thresholds.

A. 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 being used as the primary justification for “in which order to optimize” algorithm development, with unicast routing of the multipoint-to-point kind having been the primary design target for RPL.

B. 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 (capacity) 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 of these 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 on control traffic. A second trade-off 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.

C. 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

¹Note that while this may be commonly assumed, this is not an universal distribution of traffic patterns in WSNs – there are scenarios in which sensor-WSN router to sensor-WSN router traffic is assumed a more common occurrence, such as [5].

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 is 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.

D. IPv6 routing for WSNs: RPL

RPL – the “Routing Protocol for Low Power and Lossy Networks” (RPL) [4] – 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”, 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 their respective *Rank* in the DODAG (*i.e.* their 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.

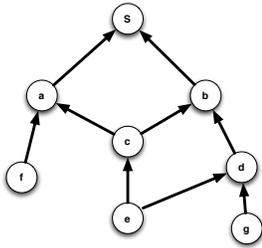


Figure 1. RPL Basic Construct: DODAGs

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 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.

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* [12].

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, add 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 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.

E. Problem Statement

RPL, as *currently* specified in [4], discusses and supports only unicast traffic, albeit two different variations hereof: sensor-to-controller traffic is supported by regular hop-by-hop routing, whereas controller-to-sensor unicast traffic is supported by way of source routing. Sensor-to-sensor routing is supported by way of combining these two: hop-by-hop routing to the “DODAG root”, and 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 I-A. 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.

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, 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, a third reason applies: in order for the DODAG root to be able to reach all sensors by way of an unicast, 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.

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

F. Paper Outline

The remainder of this paper is organized as follows: section II suggests two different mechanisms for, by way of using the data structures and topologies already maintained by RPL, providing also support for broadcast traffic in a WSN, and also briefly presents MPR Flooding – an optimized broadcast mechanism, widely employed in wireless ad hoc networks – as a point of reference and comparison. 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 two 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 transited 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 [4], wherein traffic from the source sensor transits to the DODAG root, for relaying to the destination sensor.

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 (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 causes increased battery drain, both when transmitting and receiving (and discarding) the redundant packets, and increases contention on the wireless media, increasing the probability of data loss due to collisions.

B. 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 [11]) 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 [7].

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 2 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”).

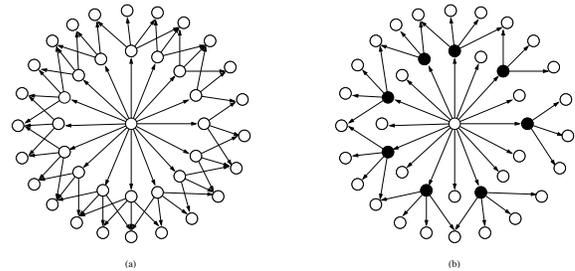


Figure 2. (a) Classic flooding and (b) MPR Flooding

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 [2], [3] and [1] – as well as for network-wide data broadcast [13].

C. Parent Flooding (PF)

Observing that a broadcast is always launched “at the 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.

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 to ensure that no router will forward a broadcast packet more than once: as each RPL router is required to select exactly one Preferred Parent, restricting retransmissions of broadcast packets to only those received from the RPL router’s preferred parent accomplishes this.

III. RPL BROADCAST PERFORMANCE STUDY

In order to explore the performance of RPL-enabled broadcast, simulations of MPRF, 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 [4]. The specific scenario settings are detailed in table I; for each datapoint in the results presented in this section, 10 different scenarios – each conforming to the same abstract scenario description – have been simulated, with the results presented being the average from these ten runs.

Parameter	Value
Ns2 version	2.34
Mobility scenarios	No mobility, random 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
Frequency	2.4 GHz

Table I
NS2 PARAMETERS

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). For MPRF, the neighborhood discovery and MPR selection part of the Java based OLSRv2 implementation [9] has been used (TC message generation disabled).

In the following comparison, PF is analyzed in two versions, *with* and *without* duplicate packet detection (PF+DD). PPF never applies duplicate packet detection, and MPR always uses it.

A. Basic Results

Figure 3 depicts the maximum and the average rank of 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 depicts the average number of parents of each WSN router in the DODAG. Keeping the density of the

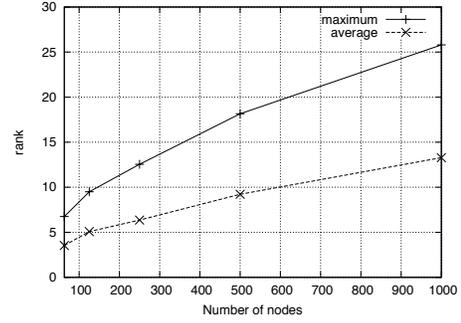


Figure 3. Maximum and average rank of the DODAG

network constant with increasing number of WSN routers, the average number of parents grows logarithmically.

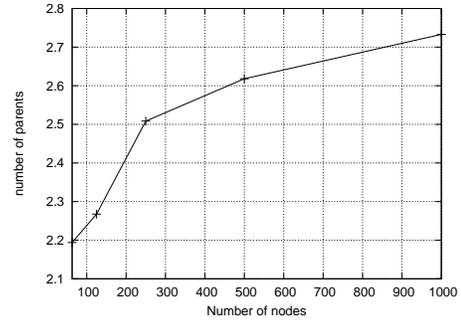


Figure 4. Average number of parents per WSN router in a DODAG

Figure 5 depicts the collision ratio of frames on the MAC layer, for PF, PF+DD, PPF, and MPRF. MPRF yields the lowest collision rate among the four analyzed protocols, while PPF yields a lower collision ratio than PF and PF+DD.

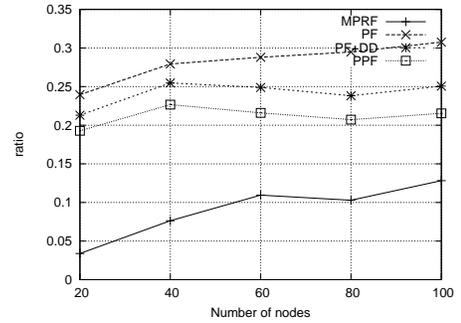


Figure 5. Collision ratio

Figure 6 depicts the delivery ratio of broadcast packets. The delivery ratio for MPRF is observed to be higher than both PF and PPF, and furthermore to not be significantly changing 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 5) of MPRF. PF+DD has a higher delivery ratio than PPF, due to the redundancy of transmissions – when a node receives

the same broadcast packet from several of its parents, chances are higher that at least one of the packets will reach the node, while if the one transmission from the preferred parent in PPF is lost due to a collision, the node will not forward the other incoming packets from its (non-preferred) parents.

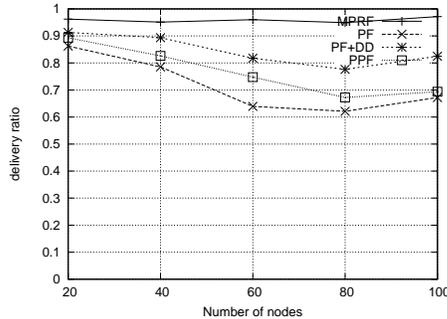


Figure 6. Delivery ratio

Figure 7 depicts the total overhead of broadcast messages; *i.e.* the cumulative number of bytes transmitted during the simulations and by each protocol. PF+DD causes a significantly lower overhead than PF, due to far fewer retransmissions. PPF has a lower overhead than PF+DD with MPRF still outperforming PPF by a large, and constant, margin.

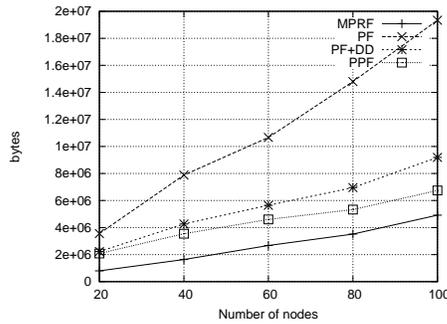


Figure 7. Total retransmission overhead

Figure 8 shows the average end-to-end delay for data traffic from the root to every WSN router in the network. MPRF incurs the lowest delay of the four protocols, while PPF incurs a much lower delay than does PF, and a slightly lower than does PF+DD.

B. PPF with Jitter

In the results presented above, data traffic has been promptly forwarded by each WSN router, without delay. As has been shown in [6], [8], 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 9 depicts the collision ratio of frames when using no jitter, and a random jitter uniformly distributed between 0 and 500 ms. With jitter, the collision ratio is much lower than it is

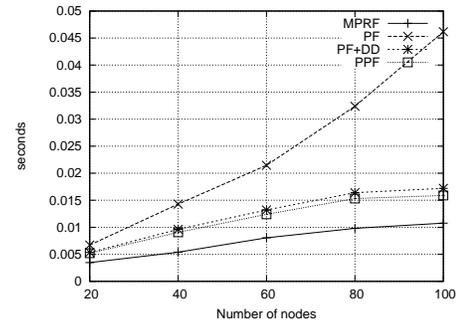


Figure 8. Average delay

without. This is due to the fewer concurrent retransmissions by adjacent WSN routers. Comparing to figure 5, PPF with jitter yields a collision ratio comparable to, or lower than, MPRF without jitter.

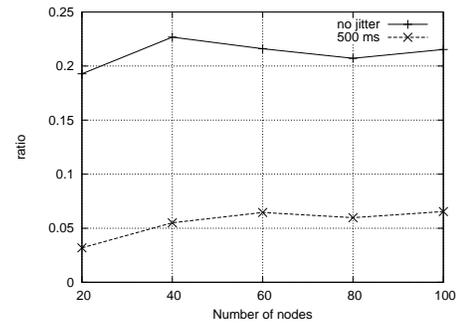


Figure 9. Collision ratio of PPF with jitter

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 10. Comparing to figure 6, the delivery ratio of PPF still remains consistently below that of MPRF.

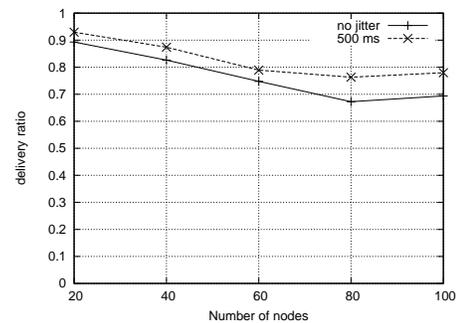


Figure 10. Delivery ratio of PPF with jitter

The drawback of using jitter is a higher end-to-end delay of packets, as depicted in figure 11. With jitter, the delay is considerably higher than it is without.

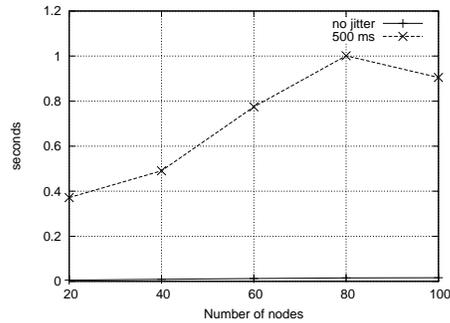


Figure 11. Average delay of PPF with jitter

IV. CONCLUSION

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

PF and PPF 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.

MPRF, however, remains much more efficient than PPF (and, obviously, PF), obtaining both a consistently lower bandwidth consumption and a consistently higher data delivery rate. 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 also necessitates a neighborhood discovery mechanism for selecting and designating MPRs – the price of a more efficient broadcast protocol.

REFERENCES

- [1] 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
- [2] T. Clausen, P. Jacquet, “Optimized Link State Routing Protocol (OLSR)”, Experimental RFC3626, <http://www.ietf.org/rfc/rfc3626.txt>
- [3] 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>
- [4] 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
- [5] 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
- [6] 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
- [7] 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
- [8] T. Clausen, C. Dearlove, B. Adamson, “Jitter Considerations in Mobile Ad Hoc Networks (MANETs)”, Informational RFC5148, <http://www.ietf.org/rfc/rfc5148.txt>
- [9] U. Herberg, “JOLSRv2 – An OLSRV2 implementation in Java”, Proceedings of the 4th OLSR Interop workshop, October 2008
- [10] U. Herberg, Ian Taylor, “Development Framework for Supporting Java NS2 Routing Protocols”, Proceedings of the 2010 International Workshop on Future Engineering, Applications and Services, May, 2010
- [11] 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
- [12] 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>
- [13] J. Macker, “Simplified Multicast Forwarding”, <http://tools.ietf.org/id/draft-ietf-manet-smf-10> (Work In Progress), March 2010