

Adjacency Persistency in OSPF MANET

Juan Antonio Cordero

N° 7390

Aôut 2010

Thème COM

 ***Rapport
de recherche***

Adjacency Persistency in OSPF MANET

Juan Antonio Cordero*

Thème COM — Systèmes communicants
Équipe-Projet Projet Hipercom

Rapport de recherche n° 7390 — Août 2010 — 15 pages

Abstract: Link-state routing remains as one of the most challenging issues in ad hoc networking, due to the special conditions and requirements that hold in such networks, which cannot be handled by classical routing protocols. In the last decade, several efforts have been deployed either to design new routing solutions adapted to ad hoc conditions, either to extend existing solutions for wired networks to the domain of wireless mobile scenarios. This paper elaborates on the latter alternative, focusing on the standard OSPF MANET extension RFC 5449. It analyzes the impact and interest of the persistency principle to the main OSPF MANET operations, in particular the adjacency synchronization and the other operations that relate to it (flooding and route construction). The presented results show that such persistent approach is appropriate for managing adjacencies in the context of RFC 5449, and significant improvements might be achieved by extending the persistent principle into the topology selection mechanism.

Key-words: MANET, Mobile, Ad hoc, Network, Routing, OSPF, MPR, Persistency, Adjacency, Flooding, Synchronization, Link-State

* Juan-Antonio.Cordero@inria.fr

Adjacency Persistency in OSPF MANET

Résumé : Le routage d'état-lieu (link-state) reste comme l'une des questions les plus difficiles dans un réseau ad hoc, en raison des conditions particulières et des exigences qui tiennent à de tels réseaux, qui ne peuvent être traitées par les protocoles de routage classiques. Dans la dernière décennie, plusieurs efforts ont été déployés, soit pour concevoir de nouvelles solutions de routage adaptées aux conditions ad hoc, soit d'étendre les solutions existantes pour les réseaux filaires au domaine des scénarios wireless mobiles. Ce document détaille la dernière alternative, en se concentrant sur l'extension standard RFC 5449 OSPF MANET. Il analyse l'impact et l'intérêt du principe de "persistance" sur des opérations principales au context OSPF MANET, en particulier la synchronisation des adjacences et les autres opérations qui s'y rapportent (flooding et construction de routes). Les résultats présentés montrent que cette approche persistante est approprié pour la gestion des adjacences dans le cadre du RFC 5449, et des améliorations significatives peuvent être obtenues en étendant le principe de "persistance" au mécanisme de sélection topologique.

Mots-clés : MANET, mobilité, ad hoc, réseau, routage, OSPF, MPR, persistance, adjacence, flooding, synchronization, état-lieu

Contents

1	Introduction	3
1.1	Paper Outline	4
2	Background	4
2.1	OSPF Architecture	4
2.2	The Multi-Point Relays Extension of OSPF RFC 5449	5
3	Persistency in OSPF MANET Adjacencies	6
3.1	The Role of Adjacencies in OSPF MANET	6
3.2	Persistent Adjacencies	7
4	Evaluation	8
4.1	Considered Configurations	8
4.2	Persistent Adjacencies and Data Routing Quality	9
4.3	Control Traffic Structure	11
5	Conclusion	14
6	References	14

1 Introduction

Due to the challenging conditions in which communication is performed in ad hoc networks (short-lived links, radio channel unreliability, scarce bandwidth, etc.), link-state routing remains as one of the outstanding issues in this realm. Therefore, there have been many efforts in the last decade to provide efficient routing mechanisms in such networks, either by designing new link-state approaches specific for Mobile Ad hoc Networks (such as OLSR [6]), either by extending existing protocols in order to facilitate operation in MANET conditions. This latter approach becomes particularly attractive for networks in which mobile ad hoc components may coexist with fixed infrastructure for which efficient routing solutions are already known.

The most significant efforts in such approach have produced several extensions of the Open Shortest Path First protocol (OSPF) [1] [8], which have been standardized by the IETF in [9], [10] and [11]. This paper focus on the standard extension RFC 5449 [9] and explores some aspects of the link-state database (LSDB) synchronization in the context of ad hoc networks. While the analysis is restricted to the particular configuration of MPR-OSPF, results might be interesting for more general conditions. Link synchronization is a basic operation in OSPF (and, more in general, in every link-state routing approach) which plays an essential role in assuring a consistent knowledge of the network topology shared by all routers, such that any of them can thus select optimal (shortest) paths to every possible destination. Due to the cost of such operation (in terms of overhead required for database exchange and utility in ad hoc networks), the study of synchronization properties and the analysis of different optimization possibilities constitute a relevant aspect to focus on.

In all the mentioned OSPF MANET extensions, synchronized links (so-called adjacencies in OSPF terminology) are selected among the set of network links

according to a specific rule to form a reduced overlay in which reliable link-state synchronization is performed in a point-to-point fashion. In the case of RFC 5449, adjacent links are persistent, meaning that adjacencies may be conserved, once they have been formed, even when they would be no longer selected as adjacent.

This paper addresses the impact of such persistency policy in the adjacency management of OSPF MANET extensions, by studying the case of RFC 5449. More precisely, it discusses the advantages of using a persistent or non persistent adjacency rule and evaluates the effect of both approaches, in the context of RFC 5449, in terms of routing quality, performance of flooding operations and amount of control traffic dedicated to LSDB synchronization processes.

1.1 Paper Outline

The remaining of this paper is organized as follows. Section 2 briefly summarizes the architecture of OSPF and its standardized MANET extensions, with special attention to RFC 5449. Section 3 describes the role of adjacencies in such routing extensions and discusses the impact of forming persistent or non-persistent ones. Section 4 presents the results of simulating persistent and non-persistent approaches in a wide range of mobile scenarios. Finally, section 5 concludes the paper.

2 Background

This section provides a brief summary of OSPF architecture (section 2.1) and presents the main techniques used in the MANET extension standardized in RFC 5449 (section 2.2).

2.1 OSPF Architecture

OSPF [1] [8] is one of the most prominent protocols for link-state IP routing within an Autonomous System (AS) [3]. As a link-state approach, it relies on the dissemination of the topology information across the network. Such dissemination allows all routers to keep a copy of the same distributed link-state database (LSDB). Maintaining a consistent LSDB enables every router to compute autonomously optimal routes (shortest paths) to all possible destinations. Such routes are computed through the Dijkstras algorithm.

Such maintenance requires the flooding of the topology information in an efficient way. This flooding is mostly handled by adjacencies. A link between two routers is called adjacent if both routers have synchronized their link-state databases. This operation requires a database exchange that is performed in a point-to-point fashion. When there are modifications in the network topology, they are advertised through the adjacent links, which thus need to form an overlay connecting all routers within the OSPF domain to each other. Topology updates are carried by Link State Advertisements (LSA).

Rules for determining which links must become adjacent depend on the type of the interfaces (point-to-point, broadcast, NBMA, point-to-multipoint or virtual link [1]) involved in the links. For instance, a point-to-point link is always declared to be adjacent. For Non-Broadcast Multiple Access networks (NBMA),

interfaces elect a Designated Router (DR) to be responsible for topology information flooding. Then, every interface in the network becomes adjacent to this elected DR and every topology update is first sent to the DR and then flooded by the DR to the whole network, through its adjacent links.

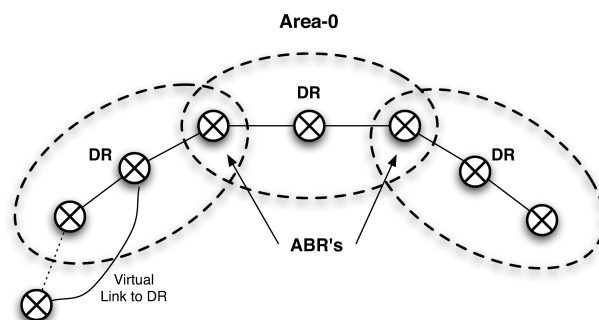


Figure 1: Area structure in an OSPF NBMA Autonomous System.

An Autonomous System ruled by OSPF can be split in routing areas. An area is a group of routers that share the same link-state database and thus share the same view of the network topology. If there are more than a single area, all OSPF areas are connected to a backbone, so-called Area Zero (see Figure 1), which handles inter-area routing. Such a partition requires a 2-level hierarchy of routers (internal routers belonging to an area, IR, and area border routers connecting two or more areas, ABR), and enables the network to contain most of the impact of a topology change (in terms of control traffic overhead to update LSDBs) in the area in which that change occurred, keeping the rest of the network relatively unaffected.

2.2 The Multi-Point Relays Extension of OSPF RFC 5449

RFC 5449 [9] specifies the standard MANET extension of OSPF based on the multi-point relays (MPR) technique [4].

The Multi-Point Relaying technique provides a mechanism for performing efficient broadcast in wireless networks. It is known that naive broadcast operation leads to undesirable saturation effects (the broadcast storm problem) [2], so MPR aims to reduce the redundancy of such naive scheme by only allowing a limited subset of neighbors of the source to retransmit broadcast packets. Figure 2 shows the difference between a broadcast packet retransmission in a pure flooding (naive broadcast) fashion and under the MPR technique, only allowing selected relays (solid balls in the figure) to retransmit the received packet. Such multi-point relays are selected by the source node and the election can be done through different heuristics, but the subset of MPRs has to be able to reach every 2-hop neighbor of the computing source (MPR coverage criterion). The selection of the most convenient relays requires naturally that the computing source is aware of its own 2-hop neighbors, since any MPR selection is performed to provide coverage to a particular set of 2-hop neighbors, using a particular set of 1-hop neighbors.

RFC 5449 relies on this MPR technique to perform the classic operations of a link-state routing protocol such as OSPF. Hence, routers synchronize their LSDBs mostly with their MPRs, topology information traffic (LSAs) are only flooded through the overlay formed by links in which any of the endpoints is MPR of the other, and only MPR links are advertised in LSAs.

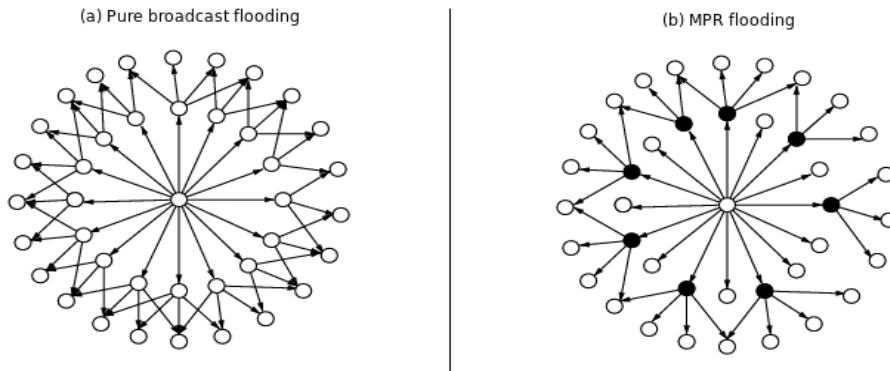


Figure 2: (a) Pure flooding vs. (b) flooding through Multi-Point Relays.

3 Persistency in OSPF MANET Adjacencies

This section explores the notion of adjacency in the context of OSPF MANET and discusses the applicability and interest of persistent approaches in such domain. Section 3.1 elaborates on the impact of adjacent links in the main routing operations (database synchronization, control traffic flooding and route selection) for the standardized OSPF MANET extensions. Section 3.2 presents the persistency concept and the implications of its implementation in the different OSPF MANET operations. The case of RFC 5449 is described specifically for both subsections.

3.1 The Role of Adjacencies in OSPF MANET

It has been already mentioned in section 2.1 that adjacencies play an essential role in regular OSPF, in particular in NBMA interfaces, those from classic OSPF which best correspond to the semibroadcast situation in MANETs. In such cases, adjacencies populate Router-LSAs and the area-wide shortest paths are thus mainly computed among them.

In the standardized OSPF MANET extensions [9] [10] [11], this role is mostly kept, and hence adjacencies take part in the following aspects:

- Database synchronization. The adjacency rule enables every node to perform link-state database synchronization with a subset of its neighbors. Such rule is based in RFC 5449 in the Multi-Point Relaying criterion, as already mentioned in section 2.2. Flooding. Control traffic (Router-LSAs) is flooded across a subset of the adjacent links in [9], [10] and [11]. In RFC 5449, nodes flood their Router-LSAs to their neighbors selected as MPR.

- Topology selection. If the adjacency rule is able to select network-wide shortest paths, as in [9], Router-LSAs advertise a subset of the adjacent links of the originating nodes (MPR selectors, for RFC 5449). Otherwise, Router-LSAs may need to list adjacent links and some additional bidirectional links, as in [10] and [11].

In all cases, changes in the adjacency rule have a direct impact in the three above-described main operations of OSPF.

3.2 Persistent Adjacencies

In the context of adjacency-related decisions, persistency is defined as an asymmetry between the condition for adding a link to the adjacent set and the condition for removing it, in which the latter is more restrictive than the former. Figure 3 illustrates an example of persistency in adjacency selection and maintenance, in which an adjacent link is only degraded in case it is no longer bidirectional.

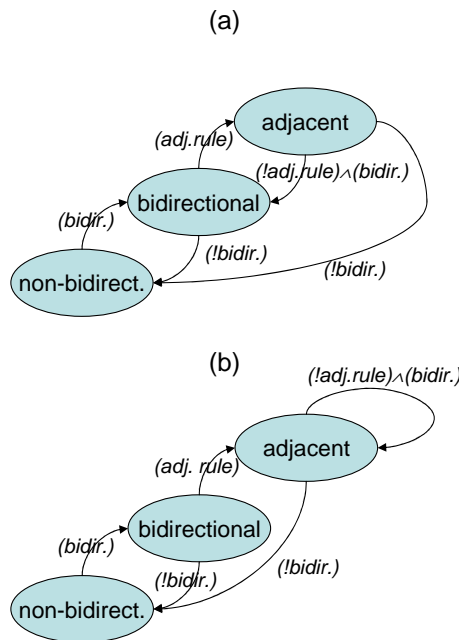


Figure 3: (a) Non-persistent and (b) persistent approaches for adjacency handling.

RFC 5449 applies such approach: a neighbor is selected by a node as adjacent if it is either an MPR or an MPR selector (neighbor that has selected the node as MPR) of the computing node. Once the adjacency has been formed (LSDB databases have been exchanged), the link is conserved as much as possible that is, as long as it is a two-way (bidirectional) link. When applied to LSDB synchronization, the main effect of this hysteresis is to reduce the number of database exchanges (in particular, for those neighbors which are alternatively added and removed from the MPR set) while increasing the overlay of reliable links (that is, those for which LSA exchange is secured by acknowledgements).

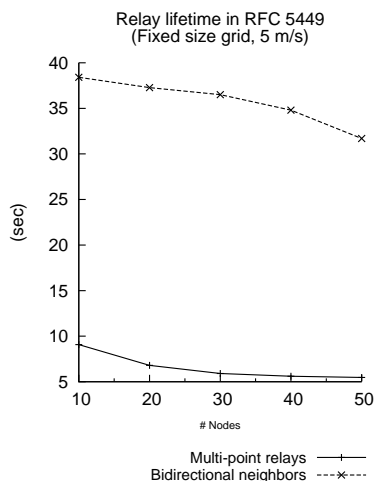


Figure 4: Relay and neighbor lifetime in RFC5449 (5 m/s).

Such approach is reasonable, given the short lifetime of multi-point relays (see Figure 4). As it was mentioned in section 2.2, MPR selection depends on the set of 1-hop and 2-hop neighbors of the computing node. Any change in this set of nodes may imply a recalculation of the whole MPR set, thus making the relay stability significantly smaller than the bidirectional neighbors.

Since adjacencies play a relevant role not only in terms of LSDB synchronization, but also for flooding and route selection operation (see section 3.1), adjacency hysteresis might be extended naturally to such aspects as well. Evaluation of these alternatives (flooding through persistent adjacencies and advertising persistent adjacent links in LSAs) is addressed in the following section.

4 Evaluation

This section analyzes the impact of adjacency persistency by evaluating the performance of four configurations based on RFC 5449, each of them exploring the implementation of the persistency principle into a different operation. Such configurations have been simulated for a wide range of density scenarios and mobility patterns. The results for moderately mobile scenarios (5 m/s), which are presented in the following subsections, are representative of the observations performed. For details of the simulations parameters, see the Appendix.

4.1 Considered Configurations

The following configurations are considered (see also table 1)¹:

- Configuration PPM (adjacency and flooding persistency). Adjacency rule is persistent, and Router LSAs are flooded through all adjacent links (including those persistent).

¹Acronyms for the considered configurations stand for [P]ersistent / non-persistent [M]PR for (i) adjacencies, (ii) flooding and (iii) topology selection.

Conf.	Adjacencies	Flooding	Topology
PPM	Persistent	Persistent	Path MPRs
PMP	Persistent	Flooding MPR selectors	Persistent
PMM (RFC 5449)	Persistent	Flooding MPR selectors	Path MPRs
MMM (Non-Persistent)	MPRs	Flooding MPR selectors	Path MPRs

Table 1: Considered configurations.

- Configuration PMP (adjacency and topology persistency). Adjacency rule is persistent, and all adjacent links (including those persistent) are advertised in LSAs.
- Configuration PMM (only adjacency persistency). Adjacency rule is persistent, but only non-persistent adjacencies (i.e., links to Path MPRs or Path MPR selectors) are advertised in the topology selection. LSAs are only flooded through Flooding MPR selectors.
- Configuration MMM (non-persistent approach). Links are no longer adjacent when none of the involved nodes is MPR of the other.

Configurations P**² are RFC5449-compliant [9], configuration PMM following literally the specification. Configuration MMM breaks with the behavior prescribed in section 5.3.4 of [9], in what concerns the treatment of adjacencies when the corresponding neighbor ceases to be part of the MPR set.

4.2 Persistent Adjacencies and Data Routing Quality

Figure 5 shows the data delivery ratio for each of the considered configurations. In general terms, it can be observed that persistent approaches (P**) present a significantly better behavior than configuration MMM, the only one which does not apply the persistency principle in any main operation of the routing extension. Moreover, the figures show that the configuration literally following the specification of RFC 5449 [9] can be still improved in terms of delivery ratio by implementing the persistency principle in OSPF operations other than LSDB synchronization, as configurations PPM and PMP do. The simulations show as well that such increase of the delivery ratio was achieved with no significant damages on the routing quality, measured in terms of data packets delay.

²That is, all configurations with persistent adjacencies, regardless of the approaches implemented for the other OSPF operations. This includes PMP, PPM and PMM.

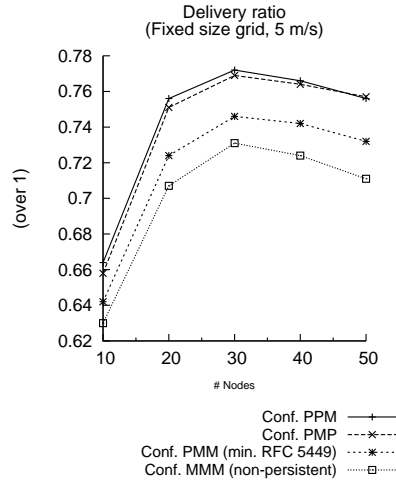


Figure 5: Delivery ratio (5 m/s).

This improvement of the routing quality is at the expense, first of all, of a relevant growth of the adjacent overlay, as it can be observed in Figure 6.a, which compares the size of the average adjacent set in persistent configurations (1, 2 and 3) with the non-persistent configuration (4). Figure 6.b, in turn, shows the gap between persistent adjacencies, with significantly longer lifetime, and non-persistent ones.

Difference in the size of adjacent overlays is more significant as the network density increases: in 50 nodes networks, with a density $\nu = \frac{50 \text{ nodes}}{400m \times 400m} = 312 \frac{\text{nodes}}{\text{km}^2}$, about 80% of the adjacencies are persistent. It has to be noted, however, that the cost of such adjacencies is extremely low in terms of control traffic exchange. They correspond indeed to links which have been already synchronized when they become persistent, and thus can be conserved with very little additional overhead. While forming an adjacency is in general a quite expensive and hazardous process in an ad hoc network, mainly due to the rigid conditions in which databases are exchanged, conservation of such adjacency has almost no cost this argument will be more developed in section 4.3.

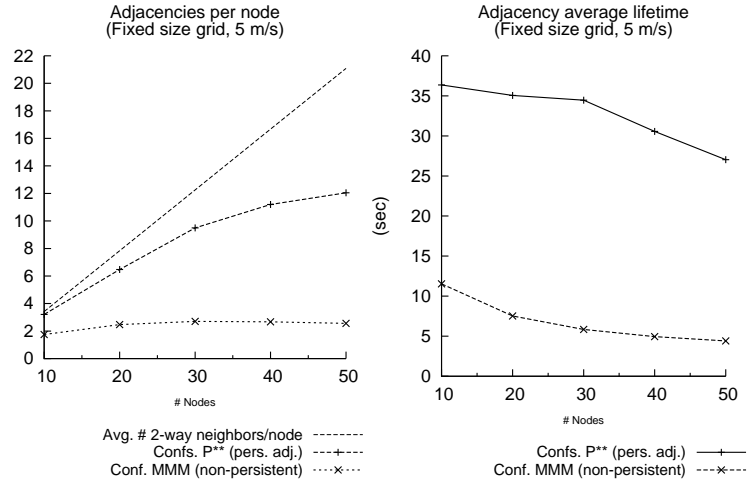


Figure 6: Adjacency characterization: (a) Adjacencies per node, (b) Adjacency lifetime (5 m/s).

4.3 Control Traffic Structure

The implementation of the persistent approach in any of the OSPF main operations link synchronization, LSA flooding, route construction necessarily implies an increment of the density (number of links) of the corresponding overlay. One of the main objectives of the persistency evaluation is to identify the cost of this overlay density increase. Control traffic structure and amount are also some of the main differences among the considered configurations.

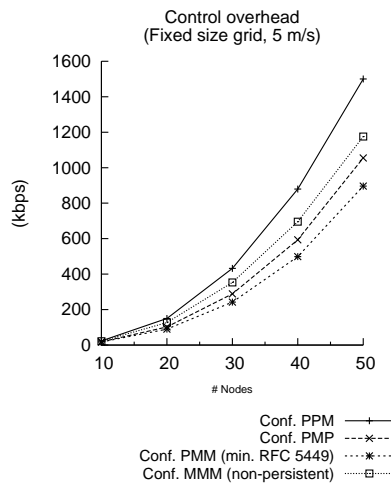


Figure 7: Total control overhead, in kbps (5 m/s).

Configuration PPM presents the highest overhead utilization, at least in terms of kbps. It is the only persistent approach that suffers from such a high overhead, being the other persistent configurations (PMP and PMM) the ones

handling lowest amounts of control traffic. While link synchronization persistency (common to configurations P**) appears to have a positive impact in the overall traffic with respect to the non-persistent configuration, the application of the persistent principle to LSA flooding overcomes that effect by making explode the amount of control traffic dedicated to flooding. This is confirmed by Fig. 8, which represents separately the number of packets dedicated to the flooding operation (mostly Link State Update packets carrying Router LSAs). This figure points out the increase of flooding traffic happening under configuration PPM, which mobilizes more than three times the number of packets in the other configurations.

The fact that PPM performs flooding across an overlay containing all persistent adjacencies has a positive effect in the quality of the flooding operation, as the evolution of the LSA retransmission ratio indicates (Fig. 9). This ratio measures the average number of times that a Router-LSA needs to be retransmitted in order to be successfully received through an adjacent link. Other than the slight improvement for configuration PPM, it can be observed that flooding over non-persistent adjacencies might be seducing from the point of view of the overlay minimization, but results in a less stable set of adjacencies (see Fig. 4) and, consequently, a more unreliable operation that requires more LSA retransmissions than any persistent configuration.

The slight improvement of the flooding quality achieved by PPM with respect to confs. PMP and PMM (Fig. 5) are at the expense of a significant increase in the control traffic. This increase is mostly caused by the explosion of the flooding overhead (see Fig. 7 and 8).

Topology information in OSPF is disseminated through LSA flooding and through point-to-point synchronization. There is a trade-off between control traffic dedicated to synchronization (adjacency-forming processes) and dedicated to LSA flooding via multicast, for similar levels of routing quality (data delivery ratio). Figure 10 shows the amount of synchronization control traffic (packets taking part in the database exchange) under each of the considered configurations. Configurations with low levels of flooding traffic are as well those with more significant amounts of synchronization control traffic, and vice versa. It can be observed (see Fig. 7) however that the implementation of persistency in the topology selection operation (configuration 3) produces a less significant impact on the overall control overhead than than in flooding (configuration 1), while achieving equivalent levels or delivery ratio (Fig. 5) and keeping reasonable flooding quality values (in terms of LSA retransmission ratio, below 10%, see Fig. 9).

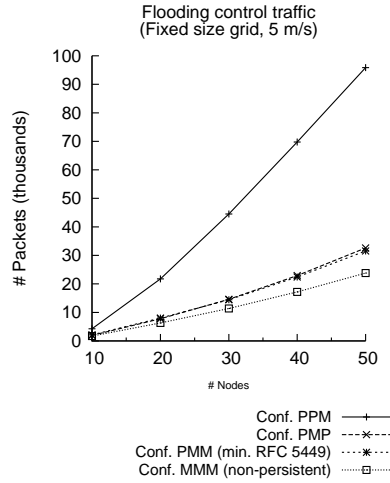


Figure 8: Flooding control traffic (LSUpdate packets via multicast), in number of packets (5 m/s).

Figure 10 shows as well the main inconvenient of the non-persistent approach, confirming what was mentioned in section 4.2. Once formed, an adjacency can be conserved with almost no additional overhead only the corresponding to LSA reliable flooding over that link. But this additional overhead is far less relevant than the overhead caused by tearing down adjacencies that might be required again in a short time, as Figure 9 indicates. Therefore, keeping a small set of adjacencies, when the adjacency rule is as unstable as MPR (see Figure 4), can only be done by tearing down cheap existing adjacencies, while increasing the number of expensive database exchanges to perform. That leads to a significant synchronization control traffic growth, as it can be appreciated in Figure 10.

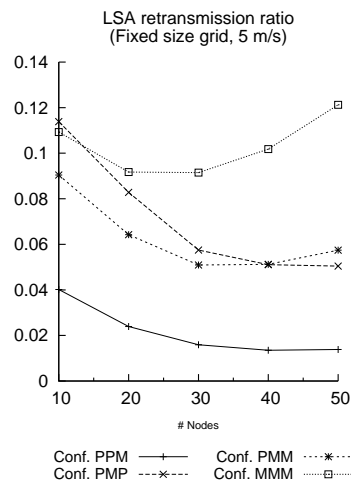


Figure 9: LSA retransmission ratio (5 m/s).

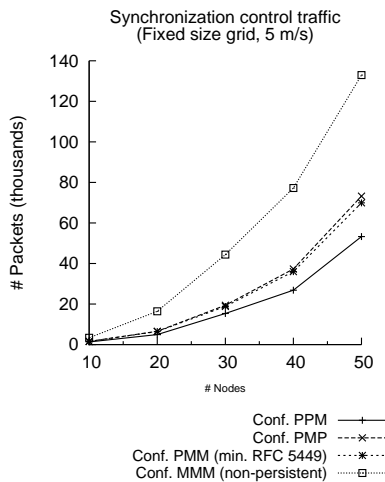


Figure 10: Flooding control traffic (LSUupdate packets via multicast), in number of packets (5 m/s).

5 Conclusion

This paper has addressed the effect and interest of the principle of persistency in the context of link-state routing synchronization. The analysis has been done by focusing on the standard extension of OSPF for MANET specified by the IETF in RFC 5449. In such extension, like in other OSPF-like approaches, link-state database synchronization is handled by the notion of adjacency.

RFC 5449 specifies already a persistent rule for adding and removing links from the adjacent set. This paper analyzes the impact of such approach by comparing it with a non-persistent configuration. The results of such evaluation, performed by means of simulating both configurations in a wide range of mobile scenarios, clearly confirm the positive impact of persistency in terms of overhead optimization and routing quality.

The paper explores as well some other applications of the persistency principle, mainly focusing on the other main OSPF operations: LSA flooding and route construction. In these configurations, adjacency persistency is thus complemented with flooding and topology persistency, respectively. The analysis and evaluation of such configurations indicate that, while in both cases there are non-negligible benefits in terms of routing quality, flooding persistency causes a significant rise of the control traffic. In contrast, advertising persistent adjacent links has a very limited impact in terms of overhead while achieving a remarkable improvement of the overall performance.

6 References

- 1 J. Moy: RFC 2328, OSPF version 2. IETF. April 1998.
- 2 S.-Y. Ni, Y.-C. Tseng, Y.-S. Chen, J.-P. Sheu: The Broadcast Storm Problem in a Mobile Ad Hoc Network. Proceedings of ACM MobiCom99. 1999.

- 3 S. Halabi: *Internet Routing Architectures*. Cisco Press, 2000.
- 4 A. Qayyum, L. Viennot, A. Laouiti: Multipoint Relaying for Flooding Broadcast Messages in Mobile Wireless Networks. *Proceedings of the 35th Hawaii International Conference in System Sciences (HICSS)*. 2002.
- 5 G. F. Riley: The Georgia Tech Network Simulator (GTNetS). *Proceedings of the ACM SIGCOMM Workshop on Models, Methods and Tools for Reproducible Network Research (MoMeTools2003)*. 2003.
- 6 T. Clausen, P. Jacquet: RFC 3626, Optimized Link State Routing Protocol (OLSR). IETF. October 2003.
- 7 T. Henderson, P. Spagnolo, G. Pei: Evaluation of OSPF MANET Extensions. Boeing Technical Report D950-10897-1. July 2005.
- 8 R. Coltun, D. Ferguson, J. Moy, A. Lindem: RFC 5340, OSPF for IPv6. IETF. July 2008.
- 9 E. Baccelli, T. Clausen, P. Jacquet, D. Nguyen: RFC 5449, OSPF Multi-point Relay (MPR) extension for ad hoc networks. IETF. February 2009.
- 10 R. Ogier, P. Spagnolo: RFC 5614, Mobile Ad hoc Network (MANET) extension of OSPF using Connected Dominating Set (CDS) flooding. IETF. August 2009.
- 11 A. Roy, M. Chandra: RFC 5820, Extensions to OSPF to support mobile ad hoc networking. IETF. March 2010.
- 12 GNU Zebra, www.zebra.org.
- 13 INRIA MANET Extension of OSPF, publicly available in www.emmanuelbaccelli.org/ospf.



Centre de recherche INRIA Paris – Rocquencourt
Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay 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 Rennes – Bretagne Atlantique : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex
Centre de recherche INRIA Saclay – Île-de-France : Parc Orsay Université - ZAC des Vignes : 4, rue Jacques Monod - 91893 Orsay 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