MPR+SP: Towards a Unified MPR-Based MANET Extension for OSPF

— **Source link** ↗

Juan Antonio Cordero, Thomas Clausen, Emmanuel Baccelli

Institutions: French Institute for Research in Computer Science and Automation, École Normale Supérieure

Topics: Optimized Link State Routing Protocol, Ad hoc wireless distribution service, Wireless Routing Protocol, Wireless ad hoc network and Link-state routing protocol

Related papers:

- A New Routing Protocol for Heterogeneous Mobile Ad Hoc Networks

- Energy Based Analysis of Traditional Routing in Mobile Ad-Hoc Networks

- Routing Protocols in Mobile Ad-Hoc Network

- The Cambridge Ad-Hoc Mobile Routing Protocol

- Optimization of critical data synchronization via link overlay RNG in mobile ad hoc networks

MPR+SP: Towards a Unified MPR-based MANET Extension for OSPF

Juan Antonio Cordero Fuertes, Thomas Heide Clausen, Emmanuel Baccelli

MPR+SP: Towards a Unified MPR-based MANET Extension for OSPF

Juan Antonio Cordero
cordero@lix.polytechnique.fr
INRIA Saclay
France

Thomas Clausen
thomas@thomasclausen.org
École Polytechnique
France

Emmanuel Baccelli
emmanuel.baccelli@inria.fr
INRIA Saclay
France

Abstract— **Heterogeneous networks combining both wired and wireless components – fixed routers as well as mobile routers – emerge as wireless mesh networks are being deployed. Such heterogeneity is bound to become more and more present in the near future as mobile ad hoc networking becomes a reality. While it is possible to cope with heterogeneity by employing different routing protocols for the fixed / wired part and for the wireless / ad hoc part of the network, this may lead to sub-optimal performance, *e.g.* by way of longer routing paths due to these routing protocols sharing prefixes and "connecting" the network only at distinct gateways between the two routing domains. Thus, the establishment of a single unified routing domain, and the use of a single routing protocol, for such heterogeneous networks is desired. OSPF is a natural candidate for this task, due to its wide deployment, its modularity and its similarity with the popular ad hoc routing protocol OLSR. Multiple OSPF extensions for MANETs have therefore been specified by the IETF. This paper introduces a novel OSPF extension for operation on ad hoc networks, MPR+SP, and compares it with the existing OSPF extensions via simulations, which show that MPR+SP outperforms prior art.**

I. INTRODUCTION

A Mobile Ad hoc Network (MANET) is an "autonomous system of mobile routers (and associated hosts) connected by wireless links, the union of which form an arbitrary graph", and in which "routers are free to move randomly and organize themselves arbitrarily". In such a network, routers "form a dynamic topology which may change unpredictably and rapidly", and are connected via wireless "links" – presenting characteristics uncommon to IP networks [1]. Such networks present several challenges and differences with respect to usual IP networking, yielding extensive academic research in the domain, and standardized MANET routing protocols such as OLSR [2] or AODV [3].

These routing protocols were specifically optimized for ad hoc scenarios, without particular attention to operation of heterogeneous networks, *i.e.*, networks combining both wired and wireless components, as well as both fixed and mobile routers. Networks with such heterogeneity emerge, with the deployment of wireless mesh networks becoming more common.

One solution for managing such heterogeneity is to deploy different routing protocols in the wired and in the ad hoc part of the network, *i.e.* OSPF [4] and OLSR [2]. However, using different protocols is suboptimal in several ways: it may lead to suboptimal paths between the two parts of the network, through a single gateway – and this even in cases where more diverse connectivity might be leveraged, and the network may benefit from traffic engineering. Moreover, familiarity

with a single protocol is an advantage – training engineers to operate and maintain an additional routing protocol is quite costly. For these reasons, the use of a single routing protocol is desired.

OSPF is one of the most widely deployed protocol for Internet routing inside Autonomous Systems (AS) [5]; it has been in continuous use since the 1990s and is therefore well known and understood. A proactive link-state routing protocol, OSPF is powered by the same core algorithms as OLSR – the predominant MANET routing protocol. While there are aspects of OSPF which *as-is* are incompatible with operation of a MANET, the modular architecture of OSPF enables development of extensions – in particular, extensions specifically designed for MANET operation. Development of such extensions enables handling of heterogeneous networks, with both ad hoc and wired parts, and where the particularities of each such part is managed by appropriate mechanisms – all within the same routing protocol instance.

The first issue that needs to be addressed while designing an OSPF extension for MANET, is the hierarchical 2-level routing structure used by OSPF to split the Autonomous System (AS) into different areas connected via a central backbone area as shown in figure 1. Automatic maintenance of such a structure in face of node mobility is hard – and, for this paper , considered out of scope. Rather, the paper addresses issues that pertain to OSPF operation over a single area, comprising both wired and ad hoc routers.

Multiple OSPF extensions for MANET operation in a single area have been standardized by the IETF[1], including [6], [7] and [8]. This paper proposes a combination of some of the techniques, developed in these different existing OSPF extensions, in order to present a novel OSPF extension for MANET operation – obviously, with the goal of providing better performance when compared to these existing extensions.

A. Paper Outline

The remainder of this paper is organized as follows: Section II further details relevant MANET characteristics and basic OSPF concepts, and provides an overview of the different existing extensions for enabling OSPF operation of MANETs. Of these extensions, the MPR [6] and the SP [8] extensions present interesting opportunities for being combined. Thus, these are described in further details in section III, which also analytically explores some asymptotic

---

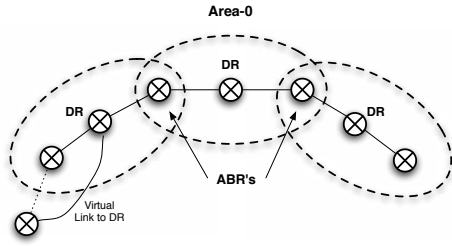[1]The Internet Engineering Task Force, http://www.ietf.org/.

Fig. 1. An OSPF autonomous system split in different areas (Area-0 is the backbone area).

properties of these. Section IV presents the architecture of MPR+SP and details how the different operations related to a link-state routing protocol are performed. Section V provides a comparison with existing OSPF extensions for MANET, based on simulation, before we conclude in section VI.

II. BACKGROUND

This section briefly describes relevant MANET characteristics, as well as basic OSPF operation. The purpose is to identify key algorithmic challenges for enabling OSPF operation of MANETs. Furthermore, this section outlines the three OSPF extensions, standardized by the IETF for MANETs.

A. Mobile Ad hoc NETworks – MANETs

MANETs present a set of properties which challenge not only OSPF, but IP-networking in general. Wireless network interfaces as well as router mobility generally leads to relatively short-lived, low bandwidth connections between routers. Moreover, during its lifetime, the quality of such a connection may vary a lot, due to interferences, obstacles, the weather etc. The term "connection" is employed in place of "link" since, as described in [9] and [10], the very notion of IP link in a MANET environment is often difficult to grasp: issues such as semi-broadcast and non-transitivity (figure 2) poses challenges to protocols running between routers in a MANET. The self-organized nature of MANETs means that, typically, no central authority is available to help alleviate such issues, which thus have to be solved by distributed algorithms.
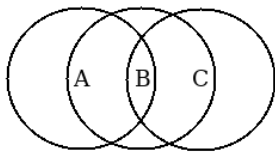

Fig. 2. Non-transitivity issues: the hidden node problem. *B can hear both A and C on the same interface, but A and C do not hear each other on this link.*

Implications of non-transitivity issues in wireless communication are illustrated in figure 2, in which three routers, each with a single interface, are displayed. If the interface of each of these three routers were connected to a *classic IP*

link[2], the fact that node A is able to communicate directly (*i.e.* no IP forwarding) with B, and B with C, implies that (i) B can communicate directly with A (connectivity is symmetric or, at least, bi-directional) and (ii) A also can communicate with C directly. In a wireless ad hoc network, this can not be assumed to be true. In figure 2, the disks represent the radio range of each of the wireless network interfaces of A, B and C. In this illustration A and C cannot communicate directly, whereas B can communicate directly with both nodes A and C. Thus $\overline{ABC}$ may appear to B to form a classic IP link – whereas from the point of view of neither A nor C does this appear to be the case. This simple example illustrates that multi-hop ad hoc wireless communication cannot be adequately represented in terms of classic IP links, due to the fact that the wireless nodes may not agree on which interfaces would constitute or be part of a "link".

B. Open Shortest Path First – OSPF

OSPF [4] [11] is a link-state routing protocol for IP networks. Each router maintains a local *link state database* (LSDB), representing the full network topology – with the objective of the protocol being that each router should have the same LSDB and, thus, the exact same view of the network topology. Paths to every possible destination are derived from the *Shortest Path Tree* (SPT) that every router computes, by way of Dijkstra's algorithm [12].

Routers acquire local topology information and advertise their own presence by periodically exchanging *Hello* messages with all their 1-hop neighbors (*i.e.* neighbor sensing). With such signaling, each router becomes aware of its immediate network topology, *i.e.* its 2-hop neighborhood. This also allows verification of bidirectional connectivity with 1-hop neighbors (then called *bidirectional* neighbors). The set of symmetric 1-hop neighbors of a router x will be hereafter denoted by $N(x)$, whereas the set of symmetric 2-hop neighbor will be denoted by $N_2(x)$.

Each router also explicitly synchronizes its LSDB with a subset of its bidirectional neighbors. Links between a router and its synchronized neighbors are called *adjacencies*, and are required to form a network-wide connected backbone, connecting all routers in the network, in order to ensure paths can be computed correctly.

Finally, routers also acquire remote topology information by way of receiving *Link State Advertisements* (LSA). Each such LSA lists mainly the current adjacencies of the router which generated the LSA. LSAs are disseminated through the entire network in reliable fashion (explicit acknowledgements and retransmissions) via the backbone formed by adjacencies; this operation is called *LSA Flooding*. Thus, any router which has formed adjacencies must advertise this periodically by way of constructing an LSA and performing LSA flooding.

Remote topology information is then used for the construction of the Shortest Path Tree: each router computes the shortest paths over the set of LSAs it has received.

[2]The example par-excellence of a classic IP link is an Ethernet.

According to this structure, OSPF distinguishes several types of links: a subset of bidirectional links become adjacent, among which a new subset is selected to be part of the SPT. While data traffic is routed on the SPT, control traffic is sent over adjacent links.
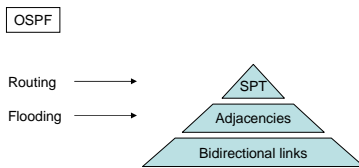

Fig. 3.   Link characterization in OSPF.

Rules for flooding and adjacency handling vary for the different *interface types* supported by OSPF. In a *non-broadcast multiple access* interface (NBMA), the existing OSPF interface type with closest characteristics to those of a MANET, the flooding procedure is mainly managed by *Designated Routers* (DRs). A Designated Router is elected from among routers whose interfaces are connected to the same *link*. Such a DR forms adjacencies with all the routers connected to the same link, and it becomes responsible for flooding of LSAs, originated by routers on that link.

As discussed in section II-A, MANET routers may not agree on which routers are connected to a given link. Thus, in a MANET, DR election may be inconsistent between different routers, causing flooding to disfunction and, possibly even preventing the protocol from converging. Handling flooding and adjacency rules in the context of wireless communication is therefore an essential aspect of OSPF operation on a MANET.

C. OSPF and MANETs

As indicated above, OSPF supports different link layer technologies by way of defining interface types, and specifying appropriate protocol behaviors according to these. MANET characteristics do not fit any existing OSPF interface type. This is in particular due to the non-transitive nature of the connectivity between routers – or, more directly, due to the fact that routers may not necessarily agree on which are or are not on the same "link", DR selection becomes problematic and thus convergence difficult.

The IETF has therefore specified a new OSPF interface type tailored for ad hoc networks, and published the three OSPF extensions for MANETs [6], [7] and [8]. Each extension provides a specific approach to OSPF operation on MANETs, *i.e.* essentially, different mechanisms for LSA flooding, adjacency setup and SPT construction:

- **MPR-OSPF** [6] performs these three operations by relying on the Multi-Point Relays (MPR [13] and section III-A). Nodes select MPRs from among their bidirectional neighbors in order to provide 2-hop coverage, and use this to disseminate their LSAs. A router becomes adjacent to both those neighbors which it has selected as multi-point relays (MPRs) and those

neighbors who have selected the router as their multi-point relay (MPR selectors). Each router advertises in its LSAs its own MPRs and MPR selectors; consequently, the Shortest Path Tree is constructed over the set of adjacencies.
- **Overlapping Relays / Smart Peering (OR/SP)** [8] disseminates LSAs via MPR flooding as in MPR-OSPF, where the multi-point relays selected among the adjacent (synchronized) neighbors of the electing router. Adjacencies are selected following the Smart Peering (SP) rule, in which a neighbor becomes adjacent if it is not already reachable through the computing router's current Shortest Path Tree (see subsection III-B for further details). LSAs list adjacent neighbors, and may also list additional bidirectional neighbors (so-called *un-synchronized adjacencies*). The SPT is thus constructed over adjacencies and a subset of bidirectional neighbors.
- **OSPF-MDR** [7] relies on two connected dominating sets (CDS) called MANET Designated Router (MDR) and Backup MDR (BMDR), which aim at extending the NBMA philosophy of "designated routers" and "backup designated routers" to MANETs. This implies that routers behave differently depending on their role: MDRs are the only nodes allowed to flood LSAs. Every non-MDR router becomes adjacent at least to the closest MDR, and MDRs must become adjacent to other MDRs. LSAs list a configurable subset of links of the originator, which must at least include the adjacent neighbors. The SPT is thus constructed over adjacencies and a subset of bidirectional neighbors.

These three extensions present two different philosophies. MPR-OSPF and OR/SP are based on multipoint relaying, and an essentially entirely distributed approach in which all routers follow the same rules – OSPF-MDR provides a centralized approach in which a router's behavior depends on its role in the network.

III. MPR AND SP – THE DETAILS

This paper proposes to unite ideas from MPR-OSPF and OR/SP in a single OSPF extension for MANETs. This section will, therefore, detail the main algorithms in MPR-OSPF and OR/SP, and discuss some of the asymptotic properties of these algorithms.

A. The Multi-Point Relaying Technique

Multi-Point Relaying (MPR) [13] is an algorithm, through which a node selects a subset of its 1-hop neighbors (*multi-point relays*) such that each 2-hop neighbor is reachable through (at least) one of the selected 1-hop neighbors (*MPR coverage criterion*). MPR selection requires that the selecting node knows the 2-hop neighbors that will be covered by its MPRs

Limiting retransmission of a packet to a subset of the neighbors of the source (see figure 4) significantly reduces the overhead for a network-wide broadcast transmission [14]. Hence, the main interest of the MPR technique in OSPF is the pruning of the number of relays for LSA flooding.
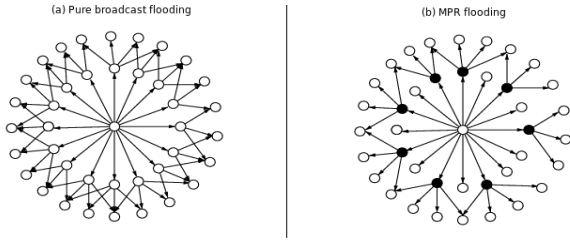
Fig. 4. Flooding based on Multi-Point Relays *vs.* pure broadcast.

The performance of the MPR technique has been throughly analyzed in [13], [14] and [15]. From the definition, it is clear that the subgraph generated by the MPRs elected by every node forms a dominating set [15]. From [14], the average size of the MPR set (that is, the average number of relays selected by a node), $|MPR(x)|$, in an infinite, 2-dim. network is upper-bounded by the expression

$$|MPR(x)| < \sqrt[3]{(3\pi)^2 M} < M, \tag{1}$$

where M represents the average number of links per node (analysis in an infinite square with uniformly distributed nodes). This bound on equation (1), which will be shown in Section V to be still far from the empiric results, illustrates the benefits of MPR with respect to pure broadcast, in terms of number of 1-hop retransmissions (M with pure broadcast).

MPRs can be used for other goals, besides reducing the flooding overhead, specifically as part of an algorithm for pruning the set of neighbors that must be advertised through LSAs, as is utilized by *Path MPRs* [6]. Path MPRs produce a reduced subset of neighbors that the computing node has to report to the rest of the network (through Router LSAs) in order to facilitate the computation of shortest paths from every possible source towards the computing node.

Lemma 1 characterizes the overlay that a node s needs to know in order to compute the shortest paths from s to all possible destinations within the network, assuming that the MPR-based link pruning algorithm keeps the shortest paths from the 2-hop neighborhood of the source towards the source:

Lemma 1: Let $G = (V, E)$ be a network graph (with V being the set of vertices and E the set of edges), an edge metrics function $cost(e \in E) > 0$, a node $s \in V$ and a subgraph $G'_s = (V, E'_s)$ including:

1) the edges connecting s to its 1-hop neighbors, and
2) for every node x of the network, the edges from x to those 1-hop neighbors of x providing shortest paths from every 2-hop neighbor of x to x.

Then, the Dijkstra algorithm computed on a source node s over G'_s selects the shortest paths in G from the source to every possible destination.

Proof: Since the Dijkstra algorithm selects the shortest paths of the graph (w.r.t. a given metrics $cost$) over which it is computed, we need to prove that the shortest paths from s in G are contained in G'_s, *i.e.*, $SPT_s(G) \subset G'_s \subset G$. Let z be an arbitrary node $z \in V$, $\overline{sz}_{sh-p}$ be the shortest path (w.r.t. $cost$) between s and z, and let $d(x, y)$ be the distance in hops between x and y.

- If $d(s, z) = 1$, $\overline{sz}_{sh-p} \in G'_s$ by condition 1 of the hypothesis.

- If $d(s, z) = 2$, let m be the intermediate node between s and z in the shortest path $\overline{sz}_{sh-p} = \{\overline{sm}, \overline{mz}\}$. The edge $\overline{sm}$ belongs to G'_s by definition (condition 1). Since $s \in N_2(z)$, the edge $\overline{mz}$ belongs to G'_s as m provides the shortest path from s to z (condition 2).

- For $d(s, z) = n > 2$, let $\{m_i\}_{0 \leq i \leq n}$ be the nodes involved in $\overline{sz}_{sh-p}$ ($m_0 \equiv s$, $m_n \equiv z$, $d(s, m_i) = i$). The edge $\overline{sm_1}$ belongs to G'_s by definition of G'_s (condition 1). The edges $\overline{m_i m_{i+1}}$ ($i \leq 1$) are included in G'_s because m_i provides shortest path from m_{i-1} (2-hop neighbor of m_{i+1}) to m_{i+1} (condition 2 of the hypothesis about G'_s). Repeating the argument along $\overline{sz}_{sh-p}$ for $\{m_j\}_{1 \leq j < n}$, we conclude that all segments $\overline{sm_1}, ..., \overline{m_i m_{i+1}}, ..., \overline{m_{n-1} z}$ belong to G'_s and thus $\overline{sz}_{sh-p}$ belongs too.

Any MPR heuristic is permitted, as long as it satisfies the coverage criterion. We assume the heuristics specified in [15] (see figure 5) for the MPR flooding algorithm and the Appendix B of [6] for the Path MPR topology pruning mechanism.

$$\begin{cases} MPR = \{\emptyset\} \\ MPR \longleftarrow \{\text{relays providing exclusive coverage to 2-hop neigh.}\} \\ while(\exists \text{ uncovered 2-hop neigh.}), \\ \quad MPR \longleftarrow \text{relay : covers the max. \# of uncovered 2-hop neigh.} \end{cases}$$

Fig. 5. Summary of the MPR heuristics.

B. Smart Peering

The Smart Peering (SP) principle provides a rule for adjacency-formation. As specified [8], a node x shall become adjacent to a bidirectional neighbor $y \in N(x)$ in case that at least one of the following two conditions is satisfied:

- There are not *enough* available paths from x to y within the overlay of (Smart Peering) adjacent links maintained by x.
- The new candidate link would provide a *significantly* cheaper path from x to y.

Depending on the definition for *enough* (threshold of available paths to discard a new adjacency) and *significantly* (with respect to the metrics), different variations of the rule might be implemented. The simplest version is presented in figure 6, and allows an adjacency with a neighbor if and only if that neighbor cannot be reached through an (Smart Peering) adjacent path from the computing node.
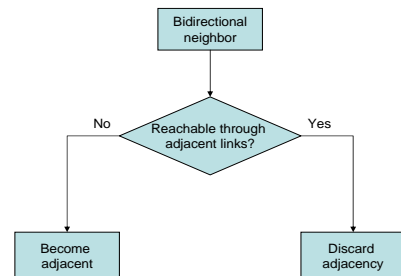

Fig. 6. The Smart Peering decision.

Note that this rule, and in particular its simplified version, presents three properties of adjacency-forming decisions:

- **Overlay density and connection in static conditions.** By construction, every node is expected to join the Smart Peering overlay, so it is trivially dense. *Lemma 2* shows that the Smart Peering overlay is also connected. In terms of a link-state routing protocol synchronization, this implies that all nodes belonging to the same Smart Peering overlay share the same link-state database (LSDB).

  Lemma 2: Using the Smart Peering, every pair of nodes (A, B) of a connected network are connected through at least one SP-adjacent path.

  Proof: Let d be the minimum distance in bidirectional hops from A to B ($d < \infty$). Then, by induction over d,
  - Case $d = 1$: if A and B are not already connected via an SP-path, the two nodes will become adjacent; this is by definition of Smart Peering.
  - $d \Rightarrow d+1$. Assume that every pair of nodes at distance d are SP-connected; let us prove the property for two nodes A and B at distance $d + 1$. Let us consider the set of bidirectional neighbors of A, $N(A)$. There exists at least one $x \in N(A)$ for which $d(x, B) = (d + 1) - 1 = d$, and is thus SP-connected to B (induction hypothesis). Calling $\overline{Ax}$ the SP-route between A and x (which exists as proved for the case $d = 1$), $\overline{xB}$ the SP-route between x and B, it is clear that the route $\overline{Ax} \cup \overline{xB}$ is an SP-route between A and B, and that concludes the proof.

- **Short-lived links filtering.** Once the first adjacency of a node has been formed and advertised to the whole network, no other node will accept a new adjacency with such node as long as the trace of the first one remains. Highly mobile nodes will therefore have difficulties to form new adjacencies after the completion of their first adjacency, while nodes presenting a lower relative speed to their neighbors will stay synchronized by means of the initial adjacencies formed.

- **Conservative minimization of the number of links.** In an ideal, static network with instantaneous flooding and a completely ordered sequence of adjacency-forming processes[3], every node would create a single adjacency when entering into the network, in order to join the adjacent overlay, and possibly an additional adjacency to a new neighbor, in order to incorporate it to the adjacent overlay. This leads to an asymptotic number of adjacencies per node between 1 and 2. In case of real mobile networks, though, the Smart Peering rule preserves the unity of the adjacent set even at the expense of redundant adjacencies (conservative minimization): a link is rejected as adjacency only if there is already a valid path in the locally stored adjacent overlay.

IV. MPR+SP ARCHITECTURE

The proposed MPR-based routing extension of OSPF, MPR+SP, combines the techniques described in section III from the two approaches already presented [6] [8]. The MPR algorithm is used for control traffic flooding and for the selection of links taking part in the Shortest Path Tree

[3]*I.e.*, a sequence in which no new adjacency is considered until the previous one has been completed.

(SPT) computation. In contrast, link-state database synchronizations (adjacencies) are minimized through Smart Peering, due to the fact that point-to-point synchronization becomes expensive and ineffective in a mobile scenario, as will be argued in the following.

Sections IV-A and IV-B describe how MPR+SP neighbors relate to each other and how do they update, diffuse and maintain the topology information across the network. Finally, section IV-C outlines the impact of this architecture in the link model.

A. Neighbor Sensing

Nodes learn their close topology and report their presence to their neighbors by exchanging Hello packets. As mentioned in section III-A, these Hellos need to contain the list of 1-hop neighbors of the originating node. By doing so, the receiving nodes can learn their 2-hop neighborhood and thus elect their MPRs.

In MPR+SP, MPRs are elected among their bidirectional 1-hop neighbors and are expected to cover all bidirectional 2-hop neighbors. Nodes selected as MPRs by a router are marked as MPRs in Hello packets from the selector.

B. Topology Information Diffusion

A link-state routing protocol is defined by the way in which the network topology information flows across the network and reaches every router within. In MPR+SP, as for any other OSPF MANET extension, this information is carried through Router LSAs that are disseminated by way of two mechanisms:

- Selective retransmission (reliable flooding over a selected subset of neighbors), and
- Link-state database synchronization (adjacency-forming processes and adjacency maintenance).

Selective retransmission follows the MPR principle: a router only forwards (and acknowledges) Router-LSAs if they have been received from one of the router's MPR selectors. Adjacencies are elected according to the Smart Peering rule and expected to exchange their respective link-state databases. Router-LSAs received during adjacency-forming processes may be flooded as well by the receiver if the LSA contains newer topology information than the one locally stored on the receiver.

The topology information collected by Router-LSAs and Hello packets is used for computing the Shortest Path Tree (SPT). In MPR+SP, routers reconstruct a network subgraph that contains the following components:

1) Path MPRs of every router in the network, listed in the corresponding Router-LSAs.
2) Adjacencies maintained by every router in the network, reported in Router-LSAs.
3) 1-hop and 2-hop neighbors of the router that performs the computation, reported via Hello packets.

From *Lemma 1*, the subgraph formed by components 1) and 3) contains the shortest path of the computing router to every other router in the network (vertex in the network graph). Adjacencies are however required for the Smart

Peering adjacency selection. This is due to the fact that adjacency candidates' acceptance or rejection depends on whether there is an existing adjacent path between the source and the candidate neighbor (see section III-B).

Figure 7 depicts a simple static network example and illustrates the three components of the subgraph that node *1* would reconstruct. Figure 7.a displays the complete network graph, and Figures 7.a, b, c and d indicates (thick lines) the subgraphs corresponding to the Path MPRs overlay, node 1's 1-hop and 2-hop neighborhood and the Smart Peering adjacent overlay, respectively. Note that the SP overlay in a static network cannot be unambiguously deduced from the network graph. For the example at figure 7.d, it has been assumed that (i) the order of appearance of the nodes correspond to their id (that is, node i will appear in the network before node j if $i < j$), (ii) adjacency-forming nodes are not concurrent, and (iii) older nodes have priority to form an adjacency to a new neighbor. It can be observed that the three components may overlap, since some links may fall into several of such categories.
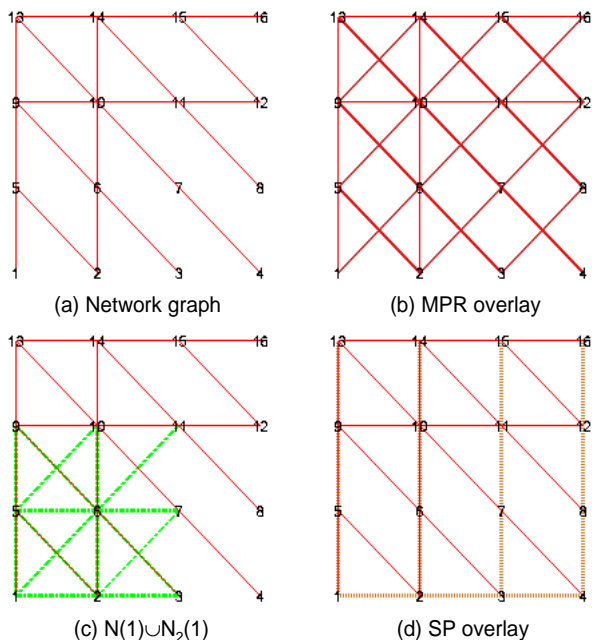

Fig. 7. Example of static network and the components of the topology subgraph reconstructed by node (1): **(a)** Network graph, **(b)** Path MPR overlay, **(c)** 1-hop and 2-hop neighborhood of (1), and **(d)** (a possible) Smart Peering overlay.

Inclusion of Path MPR links and the Smart Peering overlay in the LSDB leads to a dual network topology representation: the complete graph is used for computation of optimal routes and thus for data traffic routing, whereas the restricted subgraph containing SP links is only used for adjacency selection purposes.

C. Link Hierarchy

MPR+SP's architecture has a non-negligible impact on the link hierarchy typically supported by OSPF (see figure 3) and some of its MANET extensions (*e.g.*, RFC 5449). Figure 8 indicates the changes that MPR+SP implies in this ambit.
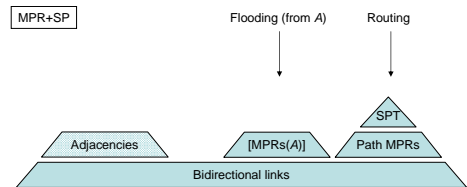

Fig. 8. Link characterization in MPR+SP.

For each node x from the network, MPR+SP generates two subgraphs based on the graph of bidirectional links within the network: the MPR subset, formed by the MPRs of x, the MPRs of these MPRs and so on; and the Path MPR subgraph containing Path MPRs of every node in the network. These two subgraphs are used in MPR+SP for control traffic flooding and data traffic routing, respectively: flooding of Router LSAs is performed over the MPR subset, while the Shortest Path Tree of x is mostly extracted from the Path MPR subset. Contrary to OSPF and its existing MANET extensions, neither of these subgraphs is necessarily contained in the subgraph of adjacencies. Such subgraph is only used for point-to-point synchronization purposes.

V. EVALUATION

This section presents a performance evaluation of MPR+SP, and compares it with the performance of the other existing MPR-based OSPF extensions. The goal is to understand to which degree the combination of different MPR-based techniques significantly improves the performance of these same techniques separately. The analysis is done by simulating the considered configurations in different mobile network scenarios, in which all nodes have the same properties and mobility pattern (see Appendix). Two experiments are performed to test the behavior of the protocols with respect to network density, on one side, and link quality, on the other. For simplicity, only the mean values of the different parameters are presented in this section, for nodes moving at a moderate speed (max. $5\frac{m}{s}$). The link quality is modeled by the non-linear parameter $\alpha \in [0, 1]$ (where $\alpha = 1$ represents an ideal channel). For a detailed description of α, see [16]. Further details on the simulation parameters are shown in tables I, II, III and IV of the Appendix.

To summarize, the results show that the hybrid configuration MPR+SP achieves similar (or slightly better) levels of routing quality (that is, delivery ratio, data path optimality and data traffic delay) to existing MPR-based extensions MPR-OSPF and Overlapping Relays – however does so by imposing a significantly lower control traffic overhead on the network.

A. Routing Quality

In general terms, MPR+SP achieves similar performance to MPR-OSPF and Overlapping Relays. Figure 9.a shows that it has a slightly higher delivery ratio than OR/SP and it copes better than MPR-OSPF with high density scenarios: its delivery ratio remains stable around 75% while MPR-OSPF delivery drops as density increases.
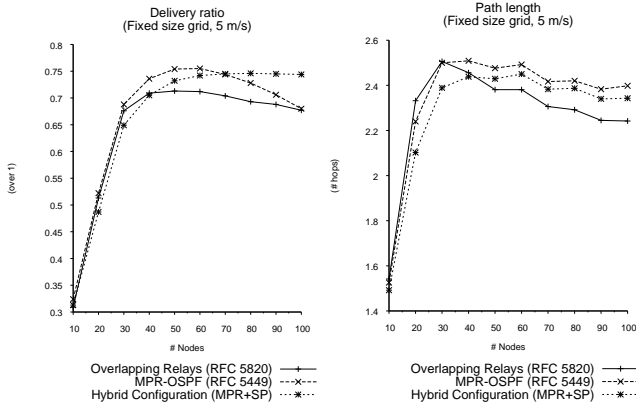
Fig. 9. (a) Delivery ratio and (b) path length (fixed size grid, 5 m/s).

Figure 9.b shows that MPR+SP provides degree of route optimality, similar to that of MPR-OSPF. This is not surprising, as both extract their Shortest Path Tree from among Path MPR-selected links. Is can also be observed that path lengths in MPR+SP are slightly lower than MPR-OSPF's. This is due to the fact that the hybrid configuration includes, together with Path MPR links (which in static conditions roughly contain shortest paths, see section III-A), the Smart Peering overlay in the computation. MPR+SP achieves a denser and thus more accurate topology map of the network.

B. Relays

Figure 10 shows two different patterns for the relay size and lifetime. MPR+SP mostly follows the pattern of MPR-OSPF: both configurations elect MPRs from among bidirectional neighbors to cover every 2-hop bidirectional neighbor, whereas Overlapping Relays computes MPRs only from among adjacent neighbors to cover only adjacent 2-hop neighbors. This is a significant difference between MPR-OSPF and MPR+SP, on one side, and Overlapping Relays, on the other, and it explains the higher relay stability and smaller relay set size achieved by OR. Such relay stability and set size reduction may seem convenient, but restricting MPR operation to the synchronized overlay distorts the role of the multi-point relays and might be harmful for the flooding efficiency, as shown in [17].
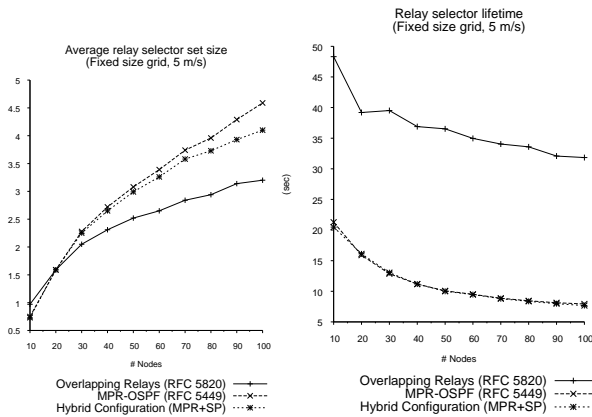

Fig. 10. (a) Average number of relays per node and (b) average relay lifetime (fixed size grid, 5 m/s).

It can also be seen from figure 10.a that the average relay set size from MPR-OSPF slighly diverges from the average size from MPR+SP for dense networks. This gap might be due to the increase of traffic density in these networks (both control and data traffic, see Figs. 15 and 16), which might prevent nodes to properly select their relays.

C. Adjacencies

Figure 11 shows two different adjacency rules: Smart Peering, used in Overlapping Relays and MPR+SP configurations, and the adjacency based on MPR, used by MPR-OSPF. figure 11.a confirms that Smart Peering reduces significantly the number of adjacencies with respect to MPR-OSPF. This latter reaches its maximum in the displayed scenario (fixed size grid, 5 m/s) with $9.34\frac{adj}{node}$, before decreasing due to network saturation. It has to be noted, however, that this count does not take into account that MPR-OSPF forms *persistent* adjacencies that are not torn down when they no longer correspond to MPR links. Such persistent adjacencies take part in the flooding operation, but are not expensive in terms of LSDB exchange.
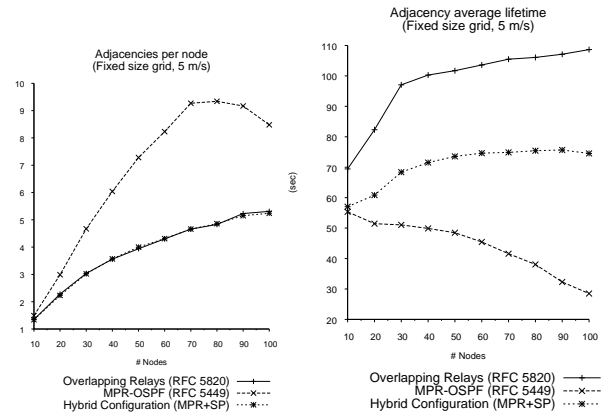

Fig. 11. (a) Average number of adjacencies per node and (b) average adjacency lifetime (fixed size grid, 5 m/s).

The adjacency lifetime in each of the configurations is shown in figure 11.b. As it was described in section III-B, adjacencies selected through the Smart Peering rule (both in MPR+SP and Overlapping Relays) are significantly more stable than those selected by MPR-OSPF. The Smart Peering capacity for discriminating the most stable links is also illustrated in figure 13, where the adjacent set of Smart Peering configurations becomes roughly stable at $\alpha \simeq 0.5$. In MPR-OSPF, in contrast, the set of adjacencies per node keeps growing as α increases.

Nonetheless, there is a non-negligible gap between the adjacency lifetime curves from MPR+SP and from Overlapping Relays. Such gap has no relation with the adjacency-forming rule (Smart Peering in both cases), but to the neighbor keep-alive mechanism. In OSPF, a node declares a neighbor *dead* if it has not received a Hello packet from it during a *DeadInterval* period. However, in a lossy channel Hello packets can be lost with a probability that increases with the length of the packet (see the lossy channel model in [16]).

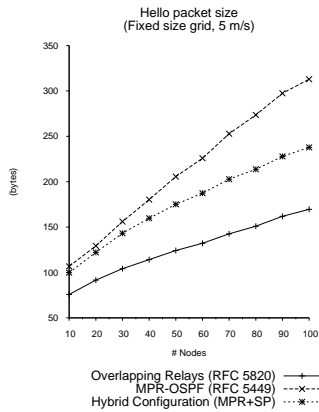Figure 12 shows the average Hello packet size for the three configurations.

Fig. 12. Average size of Hello packets (fixed grid, 5 m/s).

Aside from the fact that such keep-alive does not take into account packets other than Hellos, this policy causes that configurations with longer Hello formats (such as MPR-OSPF or MPR+SP) are more likely to declare false *dead* neighbors in lossy channels than those with shorter formats (such as Overlapping Relays). That makes the adjacency stability of configurations with longer Hellos more sensitive to link quality, as it can be seen from figure 13.b.
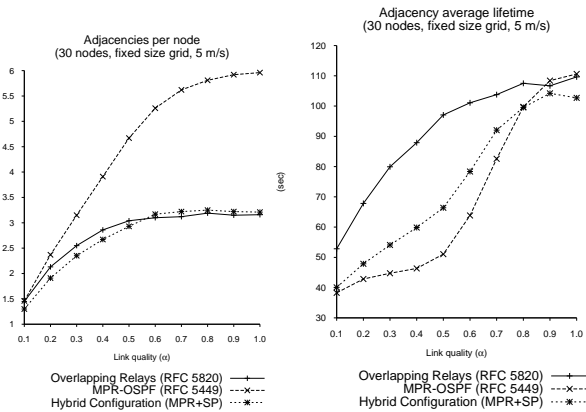

Fig. 13. (a) Average number of adjacencies per node and (b) average adjacency lifetime (30 nodes, fixed size grid, 5 m/s).

Figure 14 illustrates the effect of the keep-alive configuration in the adjacency lifetime value. It shows the adjacency lifetime achieved with MPR+SP in normal conditions (keep-alive only based on Hello reception), and the value achieved with the same configuration, when Link State Update (LSU) packets are used as keep-alives together with Hellos.
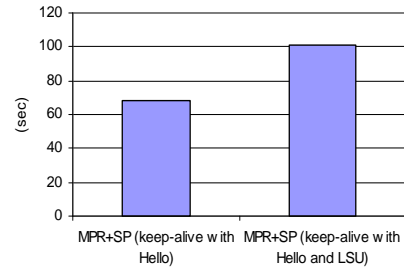

Fig. 14. OSPF keep-alive (*InactivityTimer*) impact in adjacency lifetime.

D. Control and Overall Traffic

Control traffic is one of the main arguments in favor of MPR+SP. While reaching similar levels of routing quality and data traffic optimization, the hybrid configuration manages a significantly lower control traffic overhead, both in terms of Kbps and number of packets (accesses to the channel). This can be observed in figures 15, 16 and 17. In dense networks, such overhead reduction of MPR+SP has a positive impact in the routing quality parameters, as it was already mentioned in sections V-A and V-B.
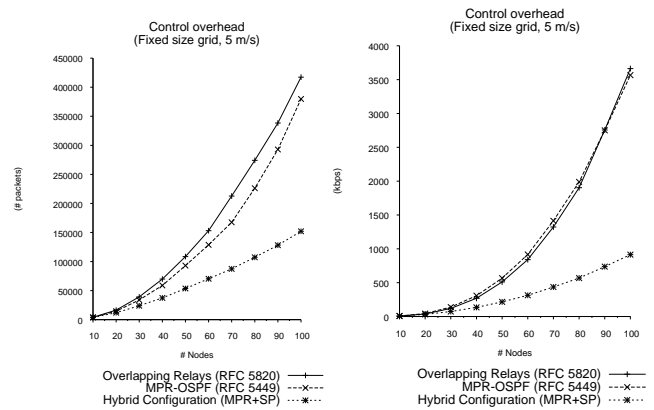

Fig. 15. Control traffic overhead in (a) number of packets and (b) Kbps (fixed size grid, 5 m/s).

MPR-OSPF provides shortest paths for data traffic (by means of the MPR technique) but it requires a significant control traffic overhead for adjacency forming and maintenance based on MPR. For dense networks, the amount of control traffic may be significant enough to have a relevant impact on the routing quality (figure 9) and internal procedures such as relay election (figure 10). This is the cost of respecting the OSPF-like notion according to which data paths should be synchronized (thus adjacent) paths. Overlapping Relays reduces the amount of control traffic dedicated to adjacency maintenance by minimizing the adjacency set (Smart Peering rule). Since this is insufficient for providing shortest paths to data traffic [17], the adjacent overlay needs to be completed with additional bidirectional links (*unsynchronized adjacencies*). At the end, this leads to similar amounts of control traffic.
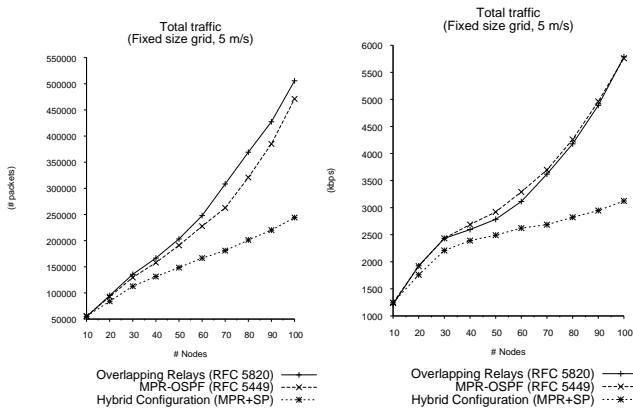
Fig. 16. Total traffic in (a) number of packets and (b) Kbps (fixed size grid, 5 m/s).

MPR+SP combines both strategies: it assures shortest paths (through MPR) for data traffic while keeping the overhead dedicated to adjacencies (through SP) extremely low. This is at the expense of breaking the relationship between synchronized (adjacent) links and SPT-selected links for data traffic. The simulations show that this relationship, which is appropriate in the context of mostly static, stable scenarios, has no longer interest for highly dynamic networks.
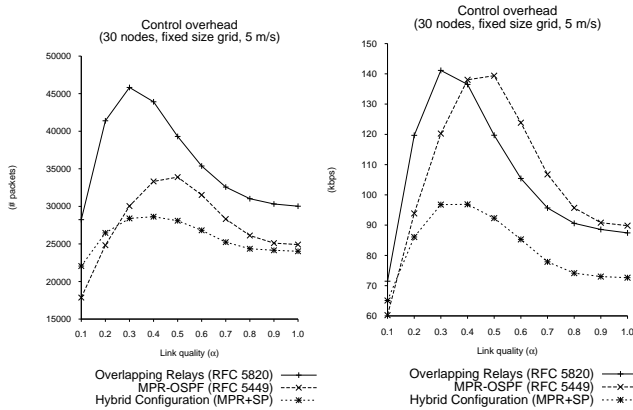

Fig. 17. Control traffic overhead (a) number of packets and (b) Kbps (30 nodes, fixed size grid, 5 m/s).

Figure 17 illustrates a different aspect of the control traffic: its evolution depending on the channel quality for each of the considered configurations. From this perspective, MPR+SP is the most robust configuration out of the three considered configurations, with respect to channel variations.

All the curves show a similar shape, with an initial region of positive slope (corresponding to very lossy channels) followed by a zone of negative slope. In terms of control traffic structure, the first region can be understood as the region in which flooding traffic (the main type of traffic increasing when the channel quality improves) is insufficient for spreading topology changes across the network. The inflection point, which varies in each configuration, corresponds to the point in which channel improvements do not longer imply increases in the control traffic (in terms of number of packets), that is, the channel is reliable

enough for performing flooding operation. In this sense, the configurations reaching faster (w.r.t. α) the inflection point are in figure 17.a those in which the flooding (LSA) traffic has more relative weight in the control traffic as a whole: OR in first term, followed by MPR+SP and then MPR-OSPF. In the latter, database exchange traffic is more significant than the whole flooding.

E. Discussion

The presented results indicate that adjacencies do not play an essential role, for neither flooding of control traffic, nor for routing of data traffic. Link synchronization is a costly process, and while it may be beneficial in case of long-living links, this does not apply in the case of MANETs, where links appear and disappear quickly: the benefit from forming adjacencies is less significant. Since adjacencies are furthermore not required in order to produce optimal routes, the size of the adjacent set can be reduced as long as the adjacent set stays connected and thus assures coherence of the LSDB in the whole network.

VI. CONCLUSION

This paper has investigated the subject of heterogeneous networks, *i.e.* networks comprising fixed wired routers and wireless mobile ad hoc routers. This environment is emerging, with mesh networks and mobility exiting the research labs and finding their place in real-world deployments. A single protocol is desired to provide routing over such networks in order to avoid suboptimality due to paths through gateways between incompatible protocols, and lack of efficient traffic engineering. OSPF is a prominent candidate to fulfill this duty, as it is both a popular routing solution for wired IP networks, and similar to OLSR, the most deployed MANET routing protocol. This paper has presented MPR+SP, a novel OSPF extension for efficient operation on ad hoc networks, and has compared MPR+OSPF with the three existing OSPF extensions for MANET as standardized by the IETF. The simulation results presented in this paper have show, that MPR+SP significantly outperforms the existing MPR-based OSPF extensions in terms of control overhead amount and robustness (w.r.t. channel variations), while keeping similar if not better data traffic delivery properties.

REFERENCES

[1] *Charter of the MANET working group*. IETF Secretariat. Proceedings of the 55th IETF Meeting in Atlanta, Georgia. 2002. *(obsolete)*
[2] T. Clausen, P. Jacquet: RFC 3626, *Optimized Link State Routing Protocol (OLSR)*. IETF. October 2003.
[3] C. Perkins, E. Belding-Royer, S. Das: RFC 3561, *Ad hoc On-Demand Distance Vector (AODV) Routing*. IETF. July 2003.
[4] J. Moy: RFC 2328, *OSPF Version 2*. Internet Society (ISOC). April 1998.
[5] S. Halabi: *Internet Routing Architectures*. Cisco Press, 2000.
[6] E. Baccelli, P. Jacquet, D. Nguyen, T. Clausen: RFC 5449, *OSPF Multipoint Relay (MPR) Extension for Ad Hoc Networks*. IETF. February 2009.
[7] R. Ogier, P. Spagnolo: RFC 5614, *Mobile Ad Hoc Network (MANET) Extension of OSPF Using Connected Dominating Set (CDS) Flooding*. IETF. August 2009.
[8] A. Roy, M. Chandra: RFC 5820, *Extensions to OSPF to Support Mobile Ad Hoc Networking*. IETF. March 2010.

[9] E. Baccelli, T. Clausen, U. Herberg, C. Perkins, *IP Links in Multihop Ad Hoc Wireless Networks?* Proceedings of the IEEE International Conference on Software Telecommunications and Computer Networks (SOFTCOM). Split, Croatia. September 2009.

[10] E. Baccelli, C. Perkins, *Multi-hop Ad Hoc Wireless Communication.* `draft-baccelli-multi-hop-wireless-communication-03.txt`. IETF Internet Draft. November 2009.

[11] R. Coltun, D. Ferguson, J. Moy: RFC 5340, *OSPF for IPv6.* IETF. July 2008.

[12] E. W. Dijkstra: "A Note on Two Problems in Connexion with Graphs", in *Numerische Mathematik* # 1 (pp. 269-271), 1959.

[13] A. Qayyum, L. Viennot, A. Laouiti: *Multipoint Relaying for Flooding Broadcast Messages in Mobile Wireless Networks.* Hawaii International Conference in System Sciences (HICSS-35), 2002.

[14] P. Jacquet, A. Laouiti, P. Minet, L. Viennot: "Performance of Multi-point Relaying in Ad Hoc Mobile Routing Protocols", in E. Gregori *et al.*: *Networking 2002, LNCS 2345*, pp. 387-398, 2002.

[15] C. Adijh, P. Jacquet, L. Viennot: RR 4597, *Computing Connected Dominated Sets with Multipoint Relays.* INRIA Research Report. October 2002.

[16] T. Henderson, P. Spagnolo, G. Pei: *Evaluation of OSPF MANET Extensions.* Boeing Technical Report D950-10897-1. July 2005.

[17] E. Baccelli, J. A. Cordero, P. Jacquet: *Multi-Point Relaying Techniques with OSPF on Ad Hoc Networks.* Proceedings of the 4th IEEE International Conference on Sensor Networks and Communications (ICSNC). Porto, Portugal. September 2009.

[18] G. F. Riley: *The Georgia Tech Network Simulator.* Proceedings of the ACM SIGCOMM 2003 Workshops. October 2003.

[19] T. Henderson, P. Spagnolo: *Comparison of Proposed OSPF MANET Extensions.* Proceedings of the Military Communications Conference (MILCOM'06). 2006.

[20] J. A. Cordero: RR 6827, *On MPR-OSPF Specification and Implementation in Quagga/GTNetS.* INRIA Research Report. February 2009.

[21] INRIA OSPF Extension for MANET Code, `www.emmanuelbaccelli.org/ospf`.

APPENDIX

The simulations have been performed with the Quagga/Zebra OSPF implementations of the considered configurations, under the GTNetS [18] environment. Implementation for OR/SP is detailed in [16], validated in [19] and follows the specification [8]. Implementation for MPR-OSPF follows the specification [6]. Code for MPR+SP is publicly available in [21].

The following tables indicate the main parameters of the simulation environment: table I shows the default value of the environment parameters (when not explicitly mentioned in the pictures) and tables II, III and IV detail the specific parameters of each analyzed configuration.

TABLE I

GENERAL SIMULATION PARAMETERS.

Name	Value
Experiment Statistic Parameters	
Number of samples	20 samples/experiment
Traffic Pattern	
Type of traffic	CBR UDP
Packet size & rate	1472 bytes, 85 pkts/sec
Traffic rate	1 Mbps
Scenario	
Mobility	Random waypoint model
Speed	$U[0, v_{max}], v_{max} = 0, 5, 10, 15 \frac{m}{s}$
Grid shape and size	Square, 600 m $\times$ 600 m
Radio range	150 m
Wireless α	0.5
Pause time	40 sec
MAC protocol	IEEE 802.11b
OSPF General Configuration	
HelloInterval	2 sec
DeadInterval	6 sec
RxmtInterval	5 sec
MinLSInterval	5 sec
MinLSArrival	1 sec

TABLE II

OR/SP SPECIFIC PARAMETERS.

Name	Value
AckInterval	1800 msec
PushbackInterval	2000 msec
Optimized Flooding?	Yes
Smart Peering?	Yes
Unsynch. adjacencies?	Yes
Incremental Hellos?	No

TABLE III

MPR-OSPF SPECIFIC PARAMETERS.

Name	Value
AckInterval	1800 msec
Flooding MPR?	Yes
Topology Reduction	MPR Topology Reduction
Adjacency Selection	MPR Adjacency Reduction

TABLE IV

MPR+SP SPECIFIC PARAMETERS.

Name	Value
AckInterval	1800 msec
Flooding MPR?	Yes
Topology Reduction	Adapted MPR Topology Reduction (with Smart Peering)
Adjacency Selection	Smart Peering