

OSPF OVER MULTI-HOP AD HOC WIRELESS COMMUNICATIONS

Emmanuel Baccelli¹, Juan Antonio Cordero² and Philippe Jacquet³

Équipe HIPERCOM, INRIA Saclay, France

Emmanuel.Baccelli@inria.fr

cordero@lix.polytechnique.fr

Philippe.Jacquet@inria.fr

ABSTRACT

Efficient OSPF (Open Shortest Path First) operation on multi-hop ad hoc wireless networks has become desirable, as wireless community mesh networks and vehicular networks emerge using OLSR (Optimized Link State Routing), a link state MANET routing protocol similar to OSPF in many aspects. OSPF is already extensively deployed and well known in wired IP networks, and could provide simple, seamless unification of wired and wireless IP networking routing-wise, if extended to operate efficiently on ad hoc networks. The IETF has thus proposed three different MANET extensions to the OSPF protocol, allowing heterogeneous networks encompassing both wired and wireless routers, which may self-organize as multi-hop wireless subnetworks, and be mobile. Two of these extensions are based on techniques derived from multi-point relaying (MPR). In the following, we compare and analyze these two extensions and we propose a unique, merged approach which out-performs the existing extensions.

KEYWORDS

Network Protocols, Wireless Network, Mobile Ad Hoc Network, Multi Point Relays, Open Shortest Path First, Routing Protocols

1. INTRODUCTION

Specific protocols have been developed for multi-hop ad hoc wireless networks in the IP realm, over the last decade. This new type of networks is characterized by rather harsh constraints such as higher topology change rates [20], lower bandwidth, lower transmission quality, more security threats, more scalability issues (as well as novel energy and memory constraints aboard some mobile network elements).

Several different categories of multi-hop ad hoc wireless networks are currently emerging, such as for instance wireless community mesh networks, vehicular networks, or sensor networks. In this paper, we focus on the first two categories, and scenarios without significant energy and memory constraints, where network nodes are fixed or moderately mobile relatively to one another.

For this category of scenarios, OLSR (Optimized Link State Routing [3]) is currently being deployed and used in numerous fast growing multi-hop wireless ad hoc networks in Europe and in North America, such as [26] [27] [28] [29] [30] [31], as well as in a variety of vehicular network deployments. OLSR is based on a proactive link state approach, which, incidentally, makes it very similar to OSPF. One question then immediately comes to mind: if OSPF and OLSR are so similar, why is OSPF not also used on multi-hop wireless networks? Operating OSPF on this new type of network is indeed a seducing idea for at least two reasons (i) legacy: OSPF is extremely well deployed, known, and renowned, thus facilitating greatly the integration of multi-hop wireless networking in the existing infrastructure, and (ii) seamless unification of wired and wireless IP networking under a single routing solution: an interesting perspective in terms of flexibility, maintenance, and costs.

There are in fact multiple issues with the use of OSPF in ad-hoc networks [4] [6]. The main problem is the amount of overhead necessary for OSPF to function, which is too substantial for

the low bandwidth available so far on multi-hop wireless networks. However, OSPF has a modular design, using different modules called interface types, each tailored for specific technologies, such as Ethernet (Broadcast interface type), or Frame Relay (Point-to-Multipoint interface type).

An extension of OSPF, namely a new OSPF interface type for multi-hop wireless networks, would thus be desirable. The goal is an extension that adapts well to the characteristics of multi-hop wireless networks, while letting OSPF run unaltered on usual networks and existing interfaces; a must, for obvious reasons including legacy and backward compatibility with networks currently running standard OSPF. The devices targeted by such an extension are assumed to have reasonable CPU, memory, battery and moderate mobility characteristics. In other words: targeted devices are rather Cisco mobile routers aboard vehicles moving at low or medium speeds, and/or fixed mesh network nodes, rather than sensors and MANET nodes moving at high speed.

1.1. Related Work

Prior art concerning OSPF operation on ad hoc networks includes initial studies such as [4] and [6] which lead to experimental OSPFv2 extensions [5] [10] based on Multi-Point Relays [3], a technique developed and used in various ad hoc networking environments over the past decade. Further work focused on OSPFv3, such as the analysis provided in [11] and [13], which compares an approach based on MANET Designated Router (MDR), with another approach based multi-point relays and Smart Peering (OR/SP). On the other hand, the study presented in [15] compared OR/SP with another approach, MPR-OSPF, also based on multi-point relaying. This ground work lead to IETF standardization, including [21], [23] and [22]. The report provided in [16] compares these three IETF standards in details. Other extensions, exploiting alternative techniques, have also been proposed. For instance [18] proposes an OSPF extension based on Routable Neighbor Graph (RNG). In this paper, we review, analyze and compare the different MPR-based techniques that were so far proposed for efficient OSPF operation on MANETs.

1.2. The Multi-Point Relaying (MPR) Technique

A significant number of network protocols, including OLSR and OSPF, rely on flooding mechanisms, *i.e.* schemes that disseminate the same piece of information to all routers in the network. A naive flooding mechanism can be as simple as: when a packet must be flooded, each node in the network repeats this packet the first time it receives it. This way, starting from the source of the packet, each node in the component connected to the source will receive the packet at least once (but typically multiple times, as shown left of Fig. 1).

Several existing techniques optimize a flooding process by reducing the number of repeaters but still ensuring that each node in the network receives a flooded packet at least once, thus saving valuable bandwidth. Multi-Point Relay (MPR) is one of the most popular such optimization, having each node select a minimal set of “relay nodes” (called MPRs), responsible for relaying flooded packets. As shown right in Fig. 1, from the local point of view of a node flooding a packet – *i.e.* the center node in the figure – this corresponds to only a small number of “necessary” neighbors (the black nodes) relaying the broadcast (instead of all the neighbors, with the naive flooding mechanism).




Figure 1. Multi-Point Relays (MPR) flooding vs. pure broadcast flooding.

In addition of ensuring that the number of repeaters is drastically reduced, while flooding still covers each node in the network, MPRs have another interesting property in the context of link state routing. Sole knowledge of the links from each node to its neighbors for which it is "necessary" (in the above-described sense) is sufficient in order to compute the shortest paths network-wide, as if the knowledge of every link in the network was available. This property thus enables a drastic reduction in the amount of link state that needs to be signalled, while still ensuring optimal connectivity.

1.3. OSPF on MANETs

As a proactive link-state routing protocol, OSPF [2] [19] employs periodic exchanges of control messages to accomplish topology discovery and maintenance: packets called Hellos are exchanged locally between neighbors to establish bidirectional links, while other packets called LSAs reporting the current state of these links are flooded throughout the entire network. This signalling results in a topology map, the link state database (LSDB), being present in each node in the network, from which a routing table can be constructed. An additional mechanism, particular to OSPF, provides explicit pair-wise synchronization of the LSDB between some neighbors, via additional control signalling (database description messages and acknowledgements). Such neighbor pairs are then called *adjacent* neighbors, while other bidirectional neighbors are called TWO-WAY.

In a wireless ad hoc environment, limited bandwidth and interferences between neighbors call for a significant reduction of OSPF control traffic [6]. At the same time, router mobility requires Hello and LSA periods to be drastically shortened in order to be able to track topology changes, implying heavier control traffic, without even more efficient control traffic reduction techniques.

The standard OSPF mechanism providing control traffic reduction is the *Designated Router* mechanism [2]. However, in a wireless ad hoc environment, this mechanism is not functional, due to the fact that wireless neighbors generally do not have the same set of wireless neighbors [20].

OSPF extensions for MANET thus use alternative mechanisms. Aside of miscellaneous tweaks and tricks such as implicit acknowledgements or control traffic multicasting (instead of unicast), these alternative mechanisms can be classified in the following categories:

- **Flooding Optimization and Backup.** Instead of the usual, naive flooding scheme, use more sophisticated techniques that reduce redundant retransmissions.

- **Adjacency Selection.** Instead of attempting to become adjacent with all its neighbors, a router becomes adjacent with only some selected neighbors.
- **Topology Reduction.** Report only partial topology information in LSAs, instead of full topology information.
- **Hello Redundancy Reduction.** In some Hello messages, report only changes in neighborhood information instead of full neighborhood information.

1.4. A Note on the Quality of User Data Paths

One element that is often neglected in discussions about adapting OSPF to multi-hop wireless networking is the fate of user data. So far, reports on OSPF extensions for ad hoc networks usually focus exclusively on control data and do not really take into account the consequences of algorithm alteration on user data. However, as shown in this paper, using longer paths can have drastic consequences in terms of the overhead that the network has to bear. Standard OSPF [2] [19] has the following principles:

- *Principle 1.* User data is always forwarded over the shortest paths.
- *Principle 2.* User data is only forwarded over links between routers with explicitly synchronized link state data-base.

In wired networks, the first principle aims at reducing delays and overhead endured by data traffic. The second principle aims at reducing risks of routing loops occurrences.

In multi-hop wireless networks, these principles are in question, as shown by the extensions proposed so far [21] [22] [23]. Concerning *Principle 1*, this paper shows that an approach that does not provide optimal paths w.r.t. the chosen metric should be discarded, if for one reason, because OSPF usually operates on networks that carry substantial data traffic. Thus, *Principle 1* should indeed be kept.

Note that the question of which metric to use on wireless links is an open, but orthogonal issue. Experiments presented in this paper use the hop-count metric because it is still, for better and for worse, the most common metric used to date on multi-hop wireless networks (though paths minimizing the number of hops are for example not always the best paths in terms of bandwidth, which is crucial in a wireless context). However, the results presented in this paper are applicable to any additive metric, and the focus is put on how to provide optimal routes assuming that the separable metric question has already been answered.

Principle 2 is on the other hand more debatable. So far, a clear difference could not be identified between (i) using paths made only of synchronized links, and (ii) using paths made both with synchronized and unsynchronized links in MANETs. This could be explained by the short lifetime of links, compared to wired links: if links are too short-lived, it could be wasteful to use bandwidth to try to synchronize link state databases; there may not even be enough time to finish synchronization before the link breaks.

1.5. Outline

This paper analyses how similar MPR concepts are used differently in each specific OSPF extension. In Section 2, a coherent set of configurable parameters is identified so as to encompass both OSPF extensions within the same framework, before discussing and evaluating the respective merits of each configuration within this framework, via simulations. For details

on the simulation environment, refer to the Appendix. Other additional parameters are presented and discussed in Section 3. Based on this analysis and on the defined framework, Section 4 proposes a recommended configuration for MPR-based OSPF operation on MANETs and provides a first evaluation on its performance. Finally, Section 5 concludes the paper.

2. PARAMETERS FOR MPR-BASED OSPF

The OSPF extensions considered in this paper, [21] and [23], essentially propose different configurations of similar concepts based on MPR. Table 1 overviews the modules of the different MPR-based configurations considered in this paper: configurations 1.x correspond to [23], while 2.1 corresponds to [21]. Configuration 2.2 is another possible configuration that is also considered in this paper.

Table 1. Considered MANET-OSPF configurations.

	Configurations 1.x		Configurations 2.x	
	1.1	1.2	2.1	2.2
Flooding Optimization	MPR Flooding		MPR Flooding	
Flooding Backup	Overlapping Relays Backup		Adjacency Backup	MPR Backup
Adjacency Selection	Smart Peering Selection		MPR Adjacency Reduction	SLO-T Selection
Topology Reduction	Unsynchronized Adjacencies	Smart Peering Reduction	MPR Topology Reduction	

Such configurations are overviewed, analyzed and evaluated through simulations in the following subsections. Section 2.1 elaborates on the main elements of flooding, including flooding relay selection and backup procedures. Section 2.2 describes the adjacency-forming decision rules applied by each configuration. Section 2.3 defines the criteria used for advertising relevant topology information in Router LSAs. Finally, Section 2.4 presents and evaluates two Hello optimization techniques that have been explored in the context of OSPF MANET extensions. Note that these techniques are not necessarily tied to a particular configuration (and thus are not mentioned in Table 1).

2.1. Flooding Optimization

In all considered configurations, MPR (see Figure 1) is used to determine flooding relays and reduce the number of forwarders of a given disseminated packet, while still ensuring that this packet is sent to each router in the network. However, there are significant differences concerning two important aspects: (i) the status of neighbors among which MPRs are selected; and (ii) the acknowledgement procedure that rules when flooding topology information (LSAs) over MPRs.

2.1.1. MPR Selection

Given a node x , its set of 1-hop neighbors $N(x)$, and its set of 2-hop neighbors $N2(x)$, the MPR selection algorithm extracts from $N(x)$ a subset of nodes $MPR(x)$ such that x is connected through $MPR(x)$ to every node of $N2(x)$, as shown in Fig. 1 for instance. The considered configurations assume different 1-hop and 2-hop neighbors set, which leads to different MPR selections, as described below:

- **MPRs selected among bidirectional neighbors**

In configurations of type 2 (2.1 and 2.2), $N(x)$ contains all the bidirectional 1-hop neighbors of the computing node x , and $N2(x)$ contains all the bidirectional 2-hop neighbors of x , that is, the nodes with are bidirectionally reachable from $N(x)$ but are not 1-hop neighbors of x . This means that the MPRs selected by x are able to reach every node within 2 (bidirectional) hops from x .

- **MPRs selected among adjacent neighbors**

Configurations of type 1 (1.1 and 1.2) enforce a more restrictive rule, which makes a router x compute the flooding relays (MPRs) only among adjacent neighbors to cover, in turn, only their own adjacent neighbors.

MPR selection among adjacent neighbors is equivalent to running the MPR algorithm over a reduced topology in which nodes are only connected by adjacencies. For any non-trivial adjacency rule (see Section 2.2), this is a much sparser network than the actual network. It is seducing to perform MPR selection over such a sub-topology because it limits the number of flooding relays (see Figure 2, which displays the average size of the flooding relay selector set), which is the approach of configurations 1.1 and 1.2.




Figure 2. Average relay (selector) set size (constant density, 5 m/s).

However, this approach is wasteful from another point of view. In sparse networks, more or less every router is chosen as MPR. Indeed, the probability of relaying an MPR flood is close to $\frac{M_r}{M}$ (with M_r being the average number of relays per node and M the average number of neighbors per node), and in sparse networks we basically get $M_r = M$. Thus, the sparser the network is, the more wasteful it is to allocate CPU resources for MPR computation. And by selecting relays for the adjacency subgraph, which by definition is sparser, configurations 1.1 and 1.2 tend to select every router within this subgraph as MPR, which tends to be wasteful. Further consequences of this choice are discussed in Section 2.2.

2.1.2. Flooding Backup

Flooded LSAs are required to reach all nodes in the network. In order to guarantee the reliability of the process, receivers are expected to acknowledge flooded LSAs, either implicitly or explicitly (see Section 3.2). In the absence of acknowledgement, different backup retransmissions strategies are employed, depending on the configuration in use:

- **Backup per adjacency**

A router receiving an LSA from an adjacent neighbor must acknowledge its reception to the neighbor. Absent this acknowledgement, the neighbor must retransmit the LSA. This process is the standard OSPF policy. This is also the behavior of configuration 2.1. This approach is called Adjacency Backup.

- **Backup per neighborhood**

While an MPR relay ensures primary transmission of an LSA, neighbors which overhear the transmission ensure backup retransmissions in case they notice some router(s) in their neighborhood have not acknowledged this LSA. This is the behavior of configurations 1.1 and 1.2. This approach is called Overlapping Relays (OR).

- **Backup per MPR selector and per adjacency**

A router receiving an LSA from an MPR selector or from an adjacent neighbor must acknowledge its reception to the sender. Absent this acknowledgement, the neighbor must retransmit the LSA. This is the behavior of configuration 2.2. This approach is called MPR Backup.

Note that the MPR Backup approach is equivalent to the Adjacency Backup strategy (and to standard OSPF backup) only in case where adjacency is tied to MPR selection. If MPR selection is not necessarily related to adjacency selection (as it is for configuration 2.2, see Section 2.2), MPR Backup and Adjacency Backup policies lead to different behaviours.

The Overlapping Relays approach differs further from standard OSPF backup, and is more complex than the other approaches, in terms of synchronization and buffer management. Simulations show that Overlapping Relays also yield significantly more retransmitted LSAs (see Fig. 3), and thus more control traffic overhead. It does not, however, substantially improve routing quality in terms of delivery ratio, or path length, as observed later in this paper (see Section 4).

Figure 3 compares LSA retransmission ratios among configurations 1.1, 1.2, 2.1 and 2.2, in a moderate mobility scenario, for different link quality scenarios. Wireless link quality is modelled by the non-linear parameter α ($\alpha \in [0,1]$, with $\alpha = 1$ standing for ideal, error-free wireless channel), which is rigorously defined in [11]). A noticeable difference can be observed between the amount of retransmissions required with configurations 1.1 or 1.2 (using Overlapping Relays), compared to the amount of retransmissions required with configurations 2.1 or 2.2. Moreover, configurations 1.1 and 1.2 (using Overlapping Relays) are also quite dependent on link quality changes, while other configurations are more stable with respect to this parameter.

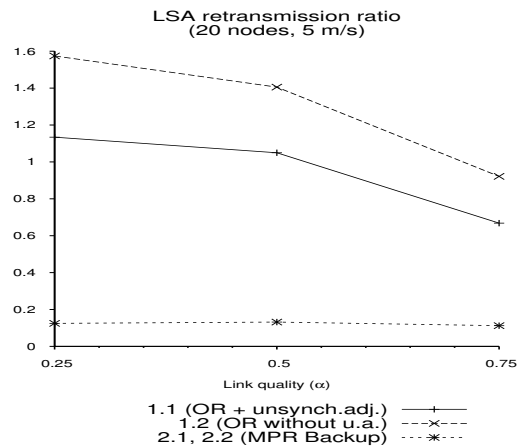


Figure 3. Number of LSA backup retransmissions over number of primary LSA transmissions (LSA retransmission ratio) for configurations 1.1, 1.2, 2.1 and 2.2 (speed: 5 m/s). Details about the simulation environment are in the Appendix.

2.2. Adjacency Selection

The decision whether or not to become adjacent with a neighbor can be taken using different criteria, depending on the configuration in use:

- **MPR selection**

A router brings up an adjacency with a bidirectional neighbor if (i) it has selected this neighbor as MPR, or (ii) it is selected as MPR by this router. These adjacencies are persistent, *i.e.*, they are maintained as long as possible. This is the behavior of configuration 2.1 and is called MPR Adjacency Selection.

- **Smart Peering selection**

The Smart Peering rule allows a router to become adjacent of a bidirectional neighbor if and only if that neighbor is not already reachable through a route formed by adjacent (Smart Peering selected) links [23].




Figure 4. The Smart Peering rule.

As far as each router maintains all the Smart Peering links selected in the network as part of the unique Link State Database, the rule brings up adjacencies to every neighbor not already present in the current LSDB. This is the behavior of configurations 1.1 and 1.2.

- **Relative Neighbor Graph selection**

Given a network links graph, the Relative Neighbor Graph [1] is an embedded subgraph that contains (but is not limited to) the shortest links. Synchronized Link Overlay approach (SLO-T) [8] [14] is such a scheme, which allows a router to bring up an adjacency if it is not eliminated by a rule breaking triangular connections (A-B-C-A), which prunes the edge with the highest cost within this triangle [18]. In a context of unit-cost links (hop-count metric), the pruning operation removes the edge with highest ID, defined as the minimum of the IDs of its vertices. An example of such triangular elimination is shown in Fig. 5, where the edge with highest ID is between node (42) and node (37), which is thus pruned, as shown on the right of the figure. Configuration 2.2 implements the unit-cost version of SLO-T as adjacency selection rule.




Figure 5. Relative Neighbor Graph (RNG) triangular elimination.

Smart Peering Selection reduces the number of adjacencies (as shown in Figure 6.a) while providing a connected set of adjacencies, but on the other hand does not generally provide a set of adjacencies that includes the shortest paths network-wide (which is an issue if adjacency selection is tied to advertised topology, as seen later in Section 2.3). SLO-T Selection produces an even smaller set of connected adjacencies. Nevertheless, it can be observed in Figure 6.b how Smart Peering tends to identify and choose more stable links.

The different properties of Smart Peering and SLO-T adjacencies (stability on one hand and minimal size on the other hand) can be explained by the following. By conditioning a new adjacency with a neighbor to its absence in the current Shortest Path Tree (SPT), the Smart Peering rule prevents a node which moves after having synchronized its LSDB from synchronizing its Link State Database again, until the formerly adjacent nodes realizes the adjacency is broken, and floods updated LSAs over the network. In particular, the moving node will not become adjacent until these LSAs are received and installed by its current potential adjacent neighbors, and vice versa. This allows a nodes to join the topology of adjacencies when it enters the network, but onwards, discourages repeated adjacency-forming processes with this node, thus punishing highly mobile nodes and giving priority the stable links rather than short-lived ones.

Both Smart Peering and SLO-T rules lead to an asymptotically connected set of adjacencies. Nonetheless, they differ in the way they handle the connectivity of the adjacency set during the convergence. Smart Peering rejects a new adjacency based on current reachability through adjacent links. On the other hand, the SLO-T algorithm may reject an adjacency candidate regardless of current adjacency topology information. For example, in Figure 5, node (42) would refuse an adjacency with (37) even if the adjacencies between (42) and (13) or between (13) and (37) have not yet been established. This behavior explains the more drastic adjacency set reduction in SLO-T compared to Smart Peering, as shown in Fig. 6.a.

MPR Adjacency selection offers a less drastic reduction in the number of adjacencies, but the provided set of adjacencies are assured to contain the shortest paths, network-wide, due to the

fact that each node becomes adjacent to those neighbors (Path MPRs) providing shortest paths from the 2-hop neighborhood [3]. In some pathological cases however, the provided set of adjacencies may not be connected network-wide [16]. In order to fix this, the adjacency set may be completed with a *synch router*, which becomes adjacent to all its neighbors and thus trivially connects the adjacency set [16], at the expense of slightly more control overhead [21].

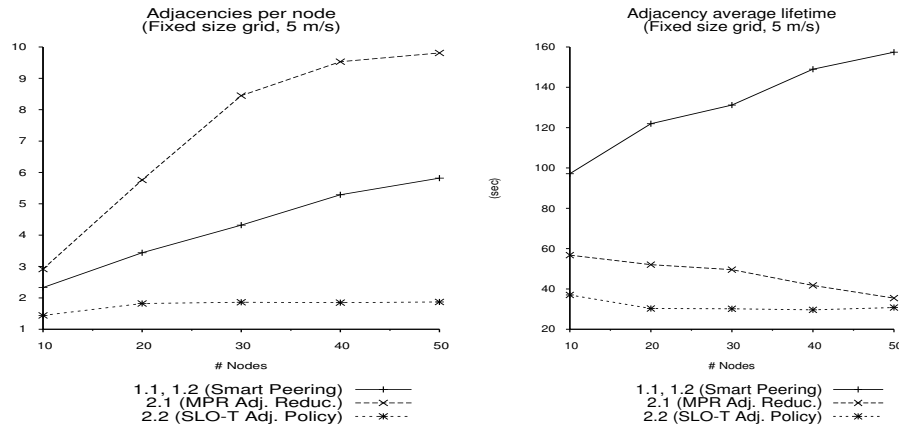


Figure 6. (a) Average number of adjacencies per node and (b) adjacency lifetime in configurations 1.1 and 2 (5 m/s).

Note that while adjacency selection and flooding relay determination are narrowly related mechanisms, this relationship differs depending on the configuration, as it was described in Section 2.1: with configurations 2.1 and 2.2, a router becomes adjacent to neighbors because it has been chosen as flooding relay to cover bidirectional 2-hop neighbors; whereas with configurations 1.1 and 1.2, flooding relays are chosen among adjacent neighbors only, to cover adjacent 2-hop neighbors, according to the Smart Peering rule.

The restriction of the set of nodes that are expected to be covered through selected MPRs leads to a reduction of the number of relays itself, as observed in Figure 2. However, this reduction is at the expense of weakening the actual flooding coverage. Indeed, configurations 1.1 and 1.2 trigger a significantly higher amount of LSA backup retransmissions, since the MPR coverage criterion only applies within the adjacency subgraph. This is shown in Fig. 7, which compares the impact of link quality degradation on control traffic composition in terms of number of packets and Kbps, when the channel quality varies from $\alpha = 0.75$ to $\alpha = 0.25$, in a small network with moderate mobility. Configurations of type 1 (1.1, 1.2) suffer from significant control traffic increase, particularly from that related to flooding operation (LSUpdate packets). This can be explained by the fact that more routers are not reached by primary transmissions, which means longer paths followed by LSAs, more backup retransmissions and more acknowledgements (which, due to more lost packets, leads in turn to even more backup retransmissions, and acknowledgements).




Figure 7. Variation of control traffic due to link quality degradation (20 nodes, 5 m/s) in terms of (a) number of packets and (b) Kbps.

2.3. Topology Reduction

In OSPF MANET configurations, Router LSAs (often referred simply as LSAs in this paper) carry all the relevant topology information that a router needs to report to the rest of the network. Router-LSAs describe different types of links depending on the configuration in use, but the contents are always closely related to the notion of adjacency:

- **All adjacencies**

The LSAs originated by a router list all adjacencies (*i.e.* links with adjacent neighbors, see Section 1.2) set up by this router. This process is the standard OSPF policy, and this is also the behavior of configuration 1.2.

- **Some selected adjacencies**

The LSAs originated by a router list a subset of the adjacencies set up by this router. This is the behavior of configuration 2.1, called MPR topology: the only links that are advertised are links to adjacent Path MPRs neighbors, *i.e.* the neighbors through which the shortest paths go, from each 2-hop neighbor towards the router [21].

- **Adjacencies and some other (bidirectional) links**

The LSAs originated by a router list some adjacencies and some TWO-WAY links, *i.e.* links with TWO-WAY neighbors (see Section 1.2), also called *unsynchronized adjacencies*. This is the behavior of configurations 1.1 and 2.2.

Unless an adjacency selection scheme is employed, listing all the adjacencies in LSAs may yield substantial control overhead. Configuration 1.2 thus uses Smart Peering to reduce the number of adjacencies, and thus the size of LSAs, which in this case report only on adjacencies. However, the impact of less link information on data traffic must be evaluated. If the subset of information is sufficient to compute the shortest paths (such as the subset provided by MPR topology in configuration 2.1), there is no impact on data traffic. If on the other hand the subset is not sufficient to compute the shortest paths, the impact on data traffic may be substantial as paths may be longer than needed. This is the case with configuration 1.2, for instance. Note that paths longer than necessary mean more radio transmissions for the network to bear with the

same goodput, while the goal is on the contrary to minimize the traffic the network has to carry, both in terms of size and number of transmissions.

Figure 8 shows the average path length provided by each configuration. It can be noticed how Smart Peering in configuration 1.2 provides substantially longer paths. Note that this result was also observed in other scenarios, with different speeds.

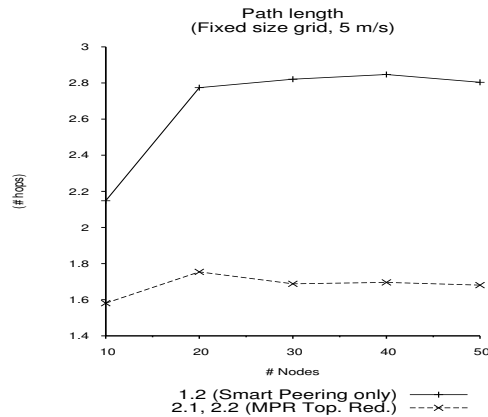


Figure 8. Average path length for data traffic (5 m/s).

If the adjacency selection scheme in use provides an adjacency set that yields longer paths, a modified scheme can complete the reported adjacency set with enough *unsynchronized adjacencies*, i.e. links with TWO-WAY neighbors (see Section 1.2), so that shortest paths can be derived from the LSDB. This is the approach of configurations 1.1 and 2.2, at the expense of more LSA overhead (with respect to configuration 1.2 for instance). This approach yields however a slightly higher risk of routing loops, since links between neighbors, that have not explicitly synchronized their LSDB, will be used for data forwarding.

Figure 9.a shows the impact of longer paths on data traffic. With configuration 1.2, which does not provide enough information to derive the shortest paths, data traffic network-wide is much higher for the same goodput, than with the other configurations, which on the other hand provide shortest paths. This gap can only be expected to grow wider with more user data input (results in Figure 9.a report up to 2Mbps).




Figure 9. (a) Data traffic and (b) total traffic (data + control) in the network (20 nodes, 5 m/s).

Note that the same gap is observed taking into account total traffic network-wide (*i.e.* both data traffic and control traffic), as shown in Figure 9.b. It shows that, in case of substantial user data input, using the shortest paths is paramount if one is to minimize the traffic overhead. Namely, inconsiderate saving on control overhead may reveal to be costly in the end, as seen with configuration 1.2. On the other hand, as explained above, configurations 2.1, 2.2, and 1.1 provide the shortest paths.

Finally, while tying adjacency selection and topology reduction is the standard OSPF approach [2] [19], it is however a seducing idea to undo this tie in a mobile ad hoc context. Further discussion on this particular subject is proposed in Section 4.

2.4. Hello Redundancy Reduction

Although the Hello traffic is a relatively small source of control traffic in mobile networks [12], some optimization techniques for information carried by Hello packets may be explored as well. Since OSPF Hello packets typically advertise all the noticed 1-hop neighbors of the originating node, a natural optimization would consist on avoid redundant notifications by only reporting changes in the neighborhood occurred since the last Hello transmission. In this case, however, single transmission failures may cause loss of Hello synchronism and take away the ability to track neighborhood changes from the Hello receivers. Thus, these optimization techniques need to provide synchronism detection and recovery mechanisms in order to restore neighbor knowledge of the Hello's originating node.

In this extent, two approaches have been explored in the framework of the OSPF MANET extension efforts. Both provide sequence number in Hello packets in order to detect synchronism gaps, but they differ in their synchronism recovery alternatives. Even when they have been implemented in specific configurations, they are conceptually autonomous and can be deployed and analyzed independently from the configurations' core.

- **Proactive synchronism recovery: *Differential Hellos*.** This approach, implemented in [22], allows routers to report in Hello slots only changes in the neighborhood, via differential (shorter then full) Hello packets. Once every n Hello transmissions (configurable), the router transmits a full Hello packet instead of a differential one. In case that any differential packet is lost, these periodical full transmissions permit every neighbor to recover Hello synchronism. The number n of differential slots per full Hello transmission reflects the trade-off between the optimization amount and the average time that a receiver would require in order to restore synchronism in case of Hello transmission failure.
- **Reactive synchronism recovery: *Incremental Hellos*.** This approach is implemented as an additional feature in [23]. Unlike the differential mechanism, which assumes a passive role from the Hello receiver, the incremental approach makes it responsible for synchronism management. In case that a node enters the network or notices a Hello transmission failure (by realizing a gap between two consecutive received Hellos), it would request the corresponding Hello originating node(s) for a full transmission.

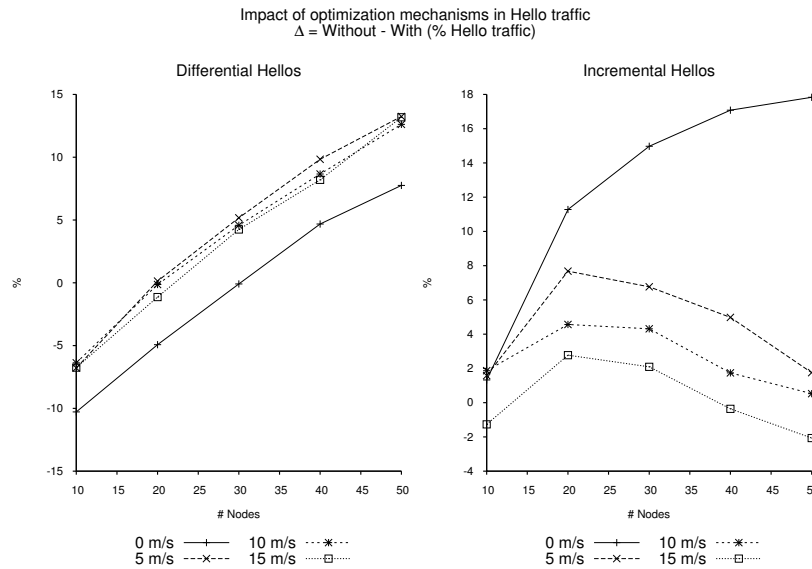


Figure 10. Impact of optimization mechanisms in Hello traffic (%).

Figure 10 shows the impact of these two optimization techniques, in terms of relative Hello traffic reduction. It can be observed that the benefits of such techniques remain in general strongly limited (less than a 18% reduction of Hello traffic is achieved at best, which represents less than 2% reduction of the total control traffic). In some cases these optimizations might even be counterproductive, generally due to additional overhead required to signal neighbor changes.

In particular, the incremental approach seems unable to significantly reduce Hello traffic in mobile and dense scenarios, in which Hello transmissions are more likely to fail and thus cause additional requests and full Hello transmissions in reply. For a fair comparison of these two techniques, however, it must be taken into account that the better overhead reduction of the differential technique w.r.t. the incremental approach is at the price of tolerating potentially longer periods of synchronism loss after a Hello failure: under the differential mechanism, a receiver cannot do anything but wait until the next full Hello transmission from the source.

3. ADDITIONAL PARAMETERS

Various additional parameters may be set differently, independently of the chosen configuration (among those considered in this paper). Most of them correspond to aspects in which the standard OSPF behavior is clearly not adapted for MANET operation. The following subsections briefly discuss the most prominent ones.

3.1. Information Determining Relays

MPR computation can be based on information contained in (i) Hellos originated by neighbor routers, or (ii) LSAs originated by neighbor routers. Both methods can be applied to any configuration discussed in this paper. However, the relay selection and update speed varies depending on this choice, as LSAs are usually generated less frequently than Hellos. Therefore, basing MPR computation on information contained in LSAs slows relays adjustments to topology changes compared to basing MPR computation on information contained in Hellos. The same reactivity could theoretically be achieved if LSA intervals were shortened to the value of *HelloInterval*, but such increase in LSA frequency would yield drastically more control overhead network-wide.

3.2. Implicit Acknowledgements

Contrary to standard OSPF policy, a flooded packet may be forwarded over the same MANET interface it was received on. This forwarded packet can thus be used as implicit acknowledgement, and eliminate the need for explicit acknowledging. The use of implicit acknowledgement can reduce the number of transmissions due to control traffic. This can be applied to any configuration discussed in this paper.

3.3. Multicasting of Control Traffic

Instead of unicast (this is standard OSPF policy) protocol packets can be multicast. The use of multicast can reduce the number of transmissions due to control traffic. This can be applied to any configuration discussed in this paper.

4. DISCUSSION

In Section 1.3, two fundamental principles of OSPF routing were mentioned, and it was discussed how these principles are applicable when operating OSPF on MANETs. The analysis and results detailed in Section 2 indicate that selecting suboptimal (non-shortest) paths for data routing has serious implications in terms of routing quality and traffic overload (see Figures 7 and 8), thus confirming clearly the pertinence of *Principle 1*. They are however less conclusive in what concerns *Principle 2*, that is, the necessity of restricting user data traffic paths to using only synchronized (adjacent) links. Indeed, no major drawbacks could be identified concerning the performance of configuration 2.2, which provides shortest paths over potentially non-adjacent links.

If, for any reason that was not explored in this paper, *Principle 2* must be kept in addition to *Principle 1*, configuration 2.1 (MPR flooding, MPR adjacency selection and MPR topology reduction, see Table 1) is the only satisfactory solution known to date, according to our result and our knowledge. If on the other hand *Principle 2* is not considered mandatory in the MANET context, we can explore other possible configurations, such as the following, which, according to the results presented in this paper, offers a better performance.

4.1. Shortest Paths with OSPF on MANETs beyond Adjacencies

Based on the analysis and simulations of the mechanisms presented in this paper, we recommend a hybrid configuration for MPR-based OSPF operation on MANETs. The main elements of this proposal are displayed in Table 2, and detailed below.

Table 2. Recommended configuration.

	Recommended Configuration
Flooding Optimization	MPR Flooding
Flooding Backup	MPR Backup
Adjacency Selection	Smart Peering
Topology Reduction	MPR Topology Reduction & Smart Peering links
Hello Redundancy Reduction	None

Flooding operation should on Multi-Point Relays (MPR). In order to ensure maximum primary flooding coverage and to decrease the overhead required for a flooding operation (see Figures 3 and 7), MPRs should be computed among bidirectional neighbors to cover every 2-hop bidirectional neighbors.

Smart Peering should be chosen as adjacency-forming strategy. As shown in Section 2.2 (see Figures 6.a and 6.b), this strategy provides a reduced adjacency backbone mainly containing the most stable links, which decreases the control traffic due to link-state databases synchronization processes.

Note that Smart Peering (see Section 2.2) normally requires links selected as adjacencies to be known by all nodes in the network, *i.e.* these links are supposed to be advertised in LSAs, and participate in the LSDB. Therefore, LSAs should advertise two types of links: (i) adjacent links selected by Smart Peering, and (ii) Path MPRs of the computing node, which are not necessarily adjacent but provide, as mentioned in Section 1.1, shortest paths.

Hello optimization techniques are generally complex and perform poorly as described in Section 2.4. Thus, normal OSPF procedure for Hello exchange should be used, enhanced only with MPR selection information such as in configurations 2.1 and 2.2.

Finally, the miscellaneous additional mechanisms described in Sections 3.2 and 3.3, should be used as described in these sections, including the use of implicit acknowledgements and of multicast transmissions for control traffic.

4.2. Recommended Configuration Evaluation

The configuration recommended in Section 4.1 offers a good bargain in terms of performance vs algorithm and implementation complexity. As shown in Figure 11, superior performance is achieved in terms of delivery ratio and delay. Using the best of both worlds produces similar route quality with less overhead, as observed in Fig. 12, which depicts the decrease in total traffic.




Figure 11. (a) Delivery ratio and (b) end-to-end delay with the recommended configuration.

Compatibility with *Principle 1* is provided using MPR topology, but *Principle 2* is left behind. The backbone of adjacencies is setup using the most stable links (using Smart Peering), where it makes more sense to synchronize databases. By doing this, a significant part of useless control traffic due to incomplete database synchronization attempts is avoided. This effect is displayed in Figure 12.a, where we can observe substantial decrease in control overhead.




Figure 12. (a) Control and (b) total traffic (control + data) with the recommended configuration.

5. CONCLUSIONS

As wireless Internet is becoming a reality, we studied in this paper a piece of tomorrow's IP protocol suite: OSPF on multi-hop wireless networks. Extending OSPF to work in such environments will allow new heterogeneous networks to exist, encompassing both wired parts and multi-hop wireless parts in the same routing domain. In this paper, we have overviewed the key challenge with routing on multi-hop wireless networks with OSPF: drastic control signalling reduction while keeping track of a topology that changes much more often compared to "usual" OSPF topology. A distinct category of solutions to this problem was identified as being different configurations of the same concept, derived from multi-point relay (MPR) techniques. A framework encompassing these configurations was identified and various possible configurations within this framework were then overviewed and evaluated via simulations. The paper concludes by recommending a specific configuration for MPR-based OSPF, which outperforms existing OSPF extensions for MANETs.

REFERENCES

- [1] G. Toussaint: "The relative neighborhood graph of a finite planar set", *Pattern Recognition*, vol. 12, pp. 261-268, 1980.
- [2] J. Moy: RFC 2328, *OSPF Version 2*. IETF. April 1998.
- [3] P. Jacquet, P. Mühlethaler, T. H. Clausen, A. Laouiti, A. Qayyum, L. Viennot: *Optimized Link State Routing for Ad Hoc Networks*. HIPERCOM Project / INRIA Rocquencourt. Proceedings of the IEEE International Multitopic Conference (INMIC). 2001.

- [4] E. Baccelli, F. Baker, M. Chandra, T. Henderson, J. Macker, R. White: *Problem Statement for OSPF Extensions for Mobile Ad Hoc Routing*. IETF Internet-Draft, draft-baker-manet-ospf-problem-statement-00 (work in progress), 2003.
- [5] T. Henderson, P. Spagnolo, J. H. Kim : *A Wireless Interface Type for OSPF*. Proceedings of the Military Communications Conference (MILCOM'03). 2003.
- [6] C. Adjih, E. Baccelli, P. Jacquet: *Link State Routing in Ad Hoc Wireless Networks*. Proceedings of the Military Communications Conference (MILCOM'03). 2003.
- [7] C. Adjih, E. Baccelli, T. H. Clausen, P. Jacquet, G. Rodolakis: Fish Eye OLSR Scaling Properties. *Journal of Communications and Networks*. 2003.
- [8] J. Cartigny, F. Ingelrest, D. Simplot: RNG Relay Subset Flooding Protocols in Mobile Ad-Hoc Networks, *International Journal of Foundations of Computer Science*, volume 14, 2003.
- [9] G. F. Riley: *The Georgia Tech Network Simulator*. Proceedings of the ACM SIGCOMM 2003 Workshops. August 2003.
- [10] J. Ahrenholz, P. Spagnolo, T. Henderson, E. Baccelli, P. Jacquet, T. Clausen: *OSPFv2 Wireless Interface Type*. IETF Internet-Draft, draft-spagnolo-manet-ospf-wireless-interface (work in progress), 2004.
- [11] T. Henderson, P. Spagnolo, G. Pei: *Evaluation of OSPF MANET Extensions*. Boeing Technical Report D950-10897-1. July 2005.
- [12] E. Baccelli: *Routing and Mobility in Large Heterogeneous Packet Networks* (PhD Thesis). École Polytechnique. Paris, 2006.
- [13] P. Spagnolo, T. Henderson: *Comparison of Proposed OSPF MANET Extensions*. Proceedings of the Military Communications Conference (MILCOM'06). 2006.
- [14] P. Jacquet: *Optimization of Point-to-point Database Synchronization via Link Overlay RNG in Mobile Ad Hoc Networks* (Rapport de Recherche 6148). INRIA Rocquencourt. February 2007.
- [15] E. Baccelli, T. H. Clausen, P. Jacquet, D. Nguyen: *Integrating VANET in the Internet Core with OSPF: the MPR-OSPF Approach*. Proceedings of the IEEE International Conference on ITS Telecommunications (ITST). June 2007.
- [16] J. A. Cordero: *Evaluation of OSPF Extensions in MANET Routing* (Master Thesis). Équipe HIPERCOM. Laboratoire d'Informatique (LIX) – École Polytechnique. September 2007.
- [18] E. Baccelli, J. A. Cordero, P. Jacquet: *Using RNG for Reliable Database Synchronization in MANETs*. Proc. of the 5th IEEE Workshop on Wireless Mesh Networks. June 2010.
- [19] R. Coltun, D. Ferguson, J. Moy, A. Lindem: RFC 5340, *OSPF for IPv6*. IETF. July 2008.
- [20] E. Baccelli, C. Perkins: *Multi-hop Ad Hoc Wireless Communication*. IETF Internet-Draft, draft-baccelli-multi-hop-wireless-communication-02 (work in progress), 2009.
- [21] E. Baccelli, T. Clausen, P. Jacquet, D. Nguyen: RFC 5449, *OSPF Multipoint Relay (MPR) Extension for Ad Hoc Networks*. IETF. February 2009.
- [22] R. Ogier, P. Spagnolo: RFC 5614, *Mobile Ad Hoc Network (MANET) Extension of OSPF Using Connected Dominating Set (CDS) Flooding*. IETF. August 2009.
- [23] A. Roy, M. Chandra: RFC 5820, *Extensions to OSPF to Support Mobile Ad Hoc Networking*. IETF. March 2010.
- [24] GNU Zebra, www.zebra.com.
- [25] INRIA OSPF Extension for MANET Code, www.emmanuelbaccelli.org/ospf.
- [26] Freifunk, Germany Wireless Community Networks, www.freifunk.net.
- [27] FunkFeuer Free Net, Austria Wireless Community Networks, www.funkfeuer.at.
- [28] AWMN, Athens (Greece) Wireless Metropolitan Community Networks, awmn.net.

- [29] Güifi.net, Catalonia (Spain) Wireless Community Networks, www.guifi.net.
 [30] Ninux.org, Rome (Italy) Wireless Community Networks, www.ninux.org.
 [31] Open Air Boston (United States) Wireless Community Networks, openairboston.net.

APPENDIX

Simulation results shown in this paper were obtained based on the Zebra OSPF implementation [24], and simulations with the GTNetS [7] simulator. Implementation for configurations 1.1 and 1.2, detailed in [9] and [13], follow specification in [22]. Implementation for configuration 2.1 follows the specification in [21]. Implementation for configuration 2.2 follows the algorithms detailed in [13]. The code for each configuration is available [25]. The following tables describe the simulation environment parameters. Table 3 shows the default value of the main parameters (when not explicitly mentioned in the figures). The values in brackets are specific for the experiment evaluating the Hello Redundancy Reduction mechanisms, when they are different from the ones used in general. Table 4 shows the parameters specific to the configurations considered in this paper.

Table 3. General simulation parameters.

Name	Value
<i>Statistic Parameters of the Experiments</i>	
Seed	0
Samples per experiment	20 (5)
<i>Traffic Pattern</i>	
Type of traffic	CBR over UDP
Packet size	1472 bytes (40 bytes)
Packet rate	85 pkts/sec (10 pkts/sec)
Traffic rate	1 Mbps (3,2 Kbps)
<i>Scenario</i>	
Mobility	Random waypoint model
Speed	$v \sim U[0, v_{max}]$ $v_{max} = 0, 5, 10, 15$ m/s
Grid shape and size	Square, 400m \times 400m
Radio range	150m
Wireless α	0.5
Pause time	40 sec
MAC protocol	IEEE 802.11b
<i>OSPF General Configuration</i>	
HelloInterval	2 sec
DeadInterval	6 sec
RxmtInterval	5 sec
MinLSInterval	5 sec
MinLSArrival	1 sec

Table 4. Specific parameters of the configurations.

Name	Value
Configurations 1.1 and 1.2	
AckInterval	1800 msec
PushbackInterval	2000 msec
Optimized Flooding?	Yes
Smart Peering?	Yes
Unsynchr. Adjacencies?	Yes
Surrogate Hellos?	Yes
Incremental Hellos?	No
Configurations 2.1 and 2.2	
AckInterval	1800 msec
Flooding MPR?	Yes
Topology Reduction	MPR Topology Reduction
Adjacency Selection	MPR Adjacency Reduction SLO-T Adjacency Policy

Authors

Emmanuel Baccelli received his Master of Science degree from EURECOM and from INT in France, in 2000. After working in industry (AT&T Labs, Metro-Optix, Hitachi Ltd.) as engineer, he obtained a PhD in 2006 from École Polytechnique, Paris, France. Since 2007, he is a scientific researcher at INRIA, France. His main research interests involve mobility, ad hoc networks, standardization, design and analysis of network protocols and algorithms.



Juan Antonio Cordero Fuertes graduated from the Facultad de Matemáticas y Estadística (Statistics & Mathematics Faculty) in 2006 and from the E. T. S. de Ingenieros de Telecomunicaciones (TelecomBCN Engineering School) in 2007, both at the Technical University of Catalonia (UPC), Spain. He is since 2007 PhD candidate at École Polytechnique, Paris, France. His research interests include routing protocols for mobile ad hoc networks and next generation Internet, as well as theory of probability and network mathematical modeling.



Philippe Jacquet graduated from École Polytechnique, Paris, France in 1981 and from École Nationale des Mines, Paris, France in 1984. He received his PhD degree from Paris Sud University, Paris, France in 1989. Since 1998, he is the scientific leader of the Hipercom Project at INRIA, France. His research interests involve information theory, probability theory, quantum telecommunication, evaluation of performance, algorithm and protocol design for mobile, wireless and ad hoc networks.

