



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Compound Wired/Wireless Internetworking with OSPF

Juan Antonio Cordero, Matthias Philipp, Emmanuel Baccelli.

N° 7642

June 2011



*R*apport
de recherche

Compound Wired/Wireless Internetworking with OSPF

Juan Antonio Cordero*, Matthias Philipp†, Emmanuel Baccelli‡.

Thème : COM – Systèmes communicants
Équipe-Projet Hipercom

Rapport de recherche n° 7642 — June 2011 — 20 pages

Abstract: As wireless ad hoc networks become more deployed, there is a growing interest for compound internetworks, that is, internetworks that contain both fixed and ad hoc networks. Routing is one of the main challenges that arise in such compound internetworks. Although specialized routing protocols exist for wired and for ad-hoc networks, and several such specialized protocols could be used together in a compound internetwork, it has been shown that the use of a single routing solution in the whole internetwork brings several advantages. The IETF has standardized extensions of the Open Shortest Path First (OSPF) protocol for ad hoc operation.

While previous performance evaluations of these extensions have focused on the wireless part of the internetwork and have been mostly performed by way of simulation tools, this paper studies practical issues of the use of a single protocol, extended OSPF, providing paths through a compound internetwork. In first term, it examines the behavior of OSPF in a real networking testbed. This testbed consists of an internetwork composed of 6 computers that form a static topology, *i.e.*, computers do not move during network lifetime. In second term, the overall behavior of extended OSPF, both considering standard OSPF and its MANET extension, is examined. Despite the limitations of the testbed, these experiments provide both a proof-of-concept and complementary results compared to prior work in the domain, which was mostly based on simulations, and focused on wireless ad hoc network scenarios only.

Key-words: OSPF, MANET, MPR, Routing, Internetwork, Compound, Testbed, Experiment

* INRIA Saclay – École Polytechnique, cordero@lix.polytechnique.fr

† INRIA Saclay – École Polytechnique, mphi@lix.polytechnique.fr

‡ INRIA Saclay – École Polytechnique, Emmanuel.Baccelli@inria.fr

Compound Wired/Wireless Internetworking with OSPF

Résumé : À mesure que les réseaux ad hoc sans fil deviennent de plus en plus déployés, il y a un intérêt croissant pour des *internetworks* (réseaux des réseaux) hybrides, c'est-à-dire, *internetworks* qui contiennent la fois des réseaux ad hoc et des réseaux fixes. En ce domain-là, le routage devient l'un des principaux défis qui se posent. Bien qu'il existe des protocoles de routage spécifiques pour réseaux filaires et des réseaux ad hoc, et plusieurs de ces protocoles pourraient être utilisés ensemble dans un *internetwork* hybride, il a été montré que l'utilisation d'une seule solution de routage dans un *internetwork* hybride a plusieurs avantages. L'IETF a standardisé trois extensions du protocole *Open Shortest Path First* (OSPF) ayant pour but le routage dans des réseaux ad hoc et à mobilité (MANETs).

Les évaluations du rendement de ces extensions développées jusqu'à présent se sont concentrées sur la partie sans fil (ad hoc) de l'*internetwork* et ont été principalement effectuées à travers de simulations. Ce rapport étudie des questions pratiques liées à l'usage d'un seul protocole de routage, en l'occurrence OSPF, sur un *internetwork* hybride. D'abord, la performance de OSPF est analysée avec des expériences sur un banc d'essai de réseaux (*testbed*). Ce *testbed* consiste en un *internetwork* hybride de 6 ordinateurs qui forment une topologie statique, c.-à.-d. où les ordinateurs ne bougent pas durant la vie du réseau. Deuxièmement, le comportement global du protocole OSPF étendu, à la fois sa version standard et son extension pour MANETs, est examiné. Malgré les limites du *testbed*, ces expériences fournissent à la fois une preuve de concept et des résultats qui confirment et complètent des travaux antérieurs dans le domaine, basés sur l'analyse du protocole sur MANETs à travers des simulations.

Mots-clés : OSPF, MANET, MPR, Routage, Internetwork, Hybride, Banc d'essai, Expérience

Contents

1	Introduction	4
1.1	Related Work	4
1.2	Paper Outline	4
2	The OSPF Protocol and its MANET Extension	5
2.1	Packet Types	6
2.2	Interface Types for Wired Links	6
2.3	MANET Interface Type	6
3	Testbed Description	7
3.1	Interfaces Configuration and Network Topology	7
3.1.1	Physical Topology	7
3.1.2	Logical Internetwork Topology	8
3.2	OSPF Routing Configuration	8
3.2.1	OSPF Adjacencies and MPRs	8
3.2.2	OSPF Flooding	9
4	Experiments and Results	9
4.1	Wireless Multi-hop Communication	10
4.2	OSPF Control Traffic Pattern	11
4.2.1	Hello Packets	12
4.2.2	LSDB Synchronization	13
4.2.3	Link State Updates, Requests and Acknowledgements	14
5	Conclusion	15
A	APPENDIX	19
A.1	Testbed Hardware	19
A.2	Testbed Software	19
A.3	Setup for PDR and RTT Measures	19
A.4	Setup for Control Traffic Measures	19
A.5	OSPF Parameters	19

1 Introduction

Since the Internet Engineering Task Force (IETF) defined in 1997 the concept of Mobile Ad hoc Networks (MANETs), several research efforts have focused on enabling communication in wireless multi-hop networks in which topology is spontaneous and dynamic. As a result of such efforts, different routing protocols have been designed for MANET operation, the most prominent to date being the OLSR [12] and AODV [14].

This paper studies the case of Autonomous Systems (AS) containing both ad hoc and fixed networks – such an AS is hereafter denominated *compound AS*. An approach for routing in a compound AS consists in splitting the AS into several routing domains, and then using a routing protocol for each domain: MANET-specific protocols in ad hoc networks, other existing protocols for fixed networks. An alternative approach consists of using a single routing protocol for the whole AS – ad hoc and fixed networks considered together. The former approach requires the presence in the AS of routers with specific hardware and software capabilities, denominated *gateways*, in order to provide connectivity between different routing domains. The latter approach, explored in this paper, reduces the cost of network maintenance and operation as no such gateways are needed, as explained in [1]. This is, however, at the expense of increasing the complexity of the employed routing protocol, which in this case needs to handle the diverse characteristics of MANETs and fixed networks with the same core mechanisms.

To date [16], two major protocols are used in the Internet for IGP routing: the Open Shortest Path First protocol (OSPF, specified in RFCs 2328 [17] and 5340 [9]) and the Intermediate-System-to-Intermediate-System protocol (IS-IS, RFC 1142 [18]). This paper explores the use of OSPF for routing in compound ASes, by way of evaluating its performance in a testbed consisting in a small compound internetwork.

1.1 Related Work

Several extensions of OSPF for MANET operation have recently been proposed and analyzed. Early studies and proposals such as [13] [20] paved the way to IETF standardizing several OSPF protocol extensions for operation over ad hoc networks [8] [5] [7]. Various studies have evaluated and compared the performance of these extensions in mobile ad hoc scenarios [6, 11, 21], mostly via simulations. Further improvements of these extensions have been proposed in [2, 3], still comparatively evaluated via simulations in MANET only scenarios. In contrast, this paper evaluates OSPF on a real testbed, consisting in both ad hoc networks and fixed wired networks.

1.2 Paper Outline

In the experiments described thereafter, all router interfaces run OSPF – interfaces attached to the wired network use the OSPF point-to-point interface specified in [9, 17], while wireless interfaces use the MANET interface type of the OSPF protocol extension specified in [8]. Analysis of the behavior of extended OSPF on such a testbed confirms results obtained in simulation-based studies. In particular, this paper concentrates on the effect of wireless links on

data path quality, on managing their coexistence with wired links in the same internetwork, as well as the impact of wireless links on overall OSPF control traffic.

Section 2 introduces the basics of OSPF and its extension for MANET operation, both used in the testbed. Section 3 describes the main characteristics of the testbed, the internetwork topology and the configuration of the participating network interfaces. Section 4 presents the performed experiments and the most significant results. Finally, section 5 concludes the paper.

2 The OSPF Protocol and its MANET Extension

OSPF [9, 17] is a link-state routing protocol for IP networks. This implies that each router maintains a local instance of the *Link-State Database* (LSDB), representing the full network topology – with the objective of the protocol being that each router should have the same information in its local instance of 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 [19]. OSPF supports network partitioning in several *areas*, in a way such that topology information maintained by different routers (the LSDB from which they maintain local instances) is the same if they belong to the same area. Throughout this document, however, a single area is configured in the internetwork.

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 are denoted by $N(x)$, whereas the set of symmetric 2-hop neighbor are denoted by $N_2(x)$.

Each router also explicitly synchronizes its local instance of 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 originating 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 network graph described in the set of received LSAs it, by way of Dijkstra’s algorithm.

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.

2.1 Packet Types

Routers in OSPF use five types of messages and packets to exchange topology information over the networks, some of them have been already mentioned in this section: *Hello* packets are used for neighbor sensing, *Database Description* (DBDesc) packets are exchanged for LSDB synchronization and *Link State Advertisements* (LSAs) are used for topology reliable flooding and update. After the exchange of DBDesc packets, a router in process of LSDB synchronization may request to its synchronizing neighbor the retransmission of particular LSAs – these requests are sent by way of *Link State Request* (LSReq) packets. Several LSAs may be sent in a single *Link State Update* packet (LSU). Several LSA acknowledgements may also be grouped in a single *Link State Acknowledgment* (LSAck) packet.

2.2 Interface Types for Wired Links

Rules for flooding and adjacency handling vary for the different *interface types* supported by OSPF. In *broadcast* and *non-broadcast multiple access* (NBMA) interfaces, 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. In *point-to-point* and *point-to-multipoint* interfaces, all links are synchronized and all interfaces participate in LSA flooding.

2.3 MANET Interface Type

The MANET interface type is defined in the extension of OSPF for operation over MANETs. Three different extensions have been standardized by the IETF [5, 7, 8], each of which specifies mechanisms to optimize topology description, flooding and LSDB synchronization in wireless ad hoc environments.

The experiments carried out used RFC 5449 [8]. Wireless interfaces following this specification select a set of Multi-Point Relays (MPRs) among its bidirectional neighbors [12, 15]. The set of MPRs selected by the wireless interface of a router must ensure that every packet received from the router reach all 2-hop neighbors of the selecting interface in 2 hops (*MPR coverage criterion*). A link between a router and one of its MPRs is denominated *MPR link*.

LSA flooding is then performed through MPRs, meaning that an LSA transmitted (originated or forwarded) by an interface is retransmitted by the MPRs of such interface. Links between interfaces and their MPRs are synchronized and thus become *adjacent*. Moreover, each interface describes in its LSAs the set of MPRs and MPR selectors. As the set of adjacencies based on MPR selection may not provide a connected subgraph, links from one additional router in the network (denominated *synch* router) are also declared adjacent to ensure adjacency set connection [4]. The Shortest Path Tree is then constructed over the set of adjacencies.

3 Testbed Description

This section describes the characteristics of the employed networking testbed. Section 3.1 presents the distribution of computers in the testbed and the network topology that they form. Section 3.2 details the implications of such topology in OSPF routing.

3.1 Interfaces Configuration and Network Topology

The testbed is composed of 6 fixed computers (routers/hosts) attached to two interconnected networks: a wired network and a wireless network. Table 1 indicates the network interfaces of each computer. For more details about computers' hardware, see Appendix A.1.

Computer	Abbr.	Wired ifs.	Wireless ifs.
server	S	eth0, eth1	–
hybrid1	h_1	eth0	wlan0
hybrid2	h_2	eth0	wlan0
wless1	w_1	–	wlan0
wless2	w_2	–	wlan0
wless3	w_3	–	wlan0

Table 1: Network interfaces of testbed computers.

3.1.1 Physical Topology

The internetwork connecting these computers was deployed in the Computer Science Lab (*Laboratoire d'Informatique*, LIX) of École Polytechnique, in Paris (France). Three scenarios –I, II and III– were configured over the resulting internetwork. These scenarios permit to test the communication between computers `wless3` and `server`, for different situations. The physical distribution of computers at LIX is displayed in Figure 1.




Figure 1: Computers position over the plan of LIX.

Positions of computers do not change, except for the case of `wless3`, which has a different position for scenario I and for scenarios II and III, as shown in Figure 1.

3.1.2 Logical Internetwork Topology

Each scenario corresponds to a specific internetwork topology. Figure 2 indicates the internetwork topology graphs for scenarios I, II and III. In the wired network, computers communicate through the IEEE 802.3 (Ethernet) standard protocol, **server** is connected with **hybrid1** by way of interface `eth0` and with **hybrid2** by way of interface `eth1`, as shown in Figure 2. In the wireless network, interfaces communicate through the IEEE 802.11b WLAN standard protocol, and all wireless routers (**hybrid1**, **hybrid2**, **wless1**, **wless2** and **wless3**) use their wireless interface `wlan0`. The topology that results from wireless reachability among computers **hybrid1**, **hybrid2**, **wless1**, **wless2** and **wless3** is modified by means of MAC filtering in order to disable links $h_1 \longleftrightarrow h_2$, $w_{1,3} \longleftrightarrow w_2$, $w_{1,3} \longleftrightarrow h_2$ and $w_2 \longleftrightarrow h_1$, as well as $h_1 \longleftrightarrow w_1$ for scenario III.

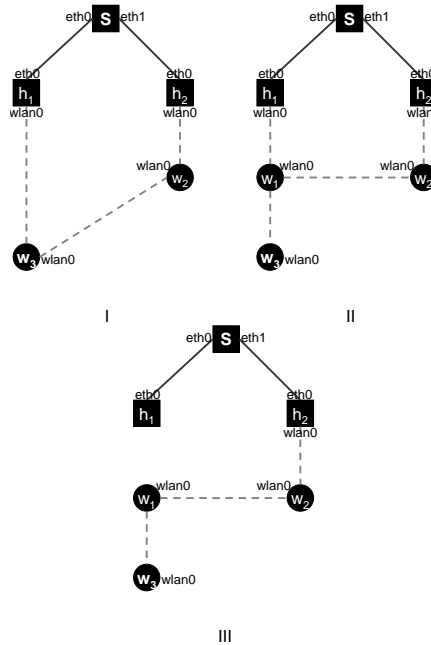


Figure 2: Considered topologies for scenarios I, II and III.

3.2 OSPF Routing Configuration

All interfaces use the extended OSPFv3 routing protocol, wired and wireless interfaces using different interface types. Wired interfaces are configured as *point-to-point interfaces*, as they are specified in RFCs 2328 [17] and 5340 [9]. Wireless interfaces are configured as *MANET interfaces*, as specified in the MPR-OSPF MANET extension for OSPF (RFC 5449 [8]).

3.2.1 OSPF Adjacencies and MPRs

According to the specification of OSPF and MPR-OSPF extension, all links in any of the considered topologies for scenarios I, II and III are adjacent. Within the wired network, every point-to-point link is an adjacency. In the wireless

network, wireless links are adjacent if they are MPR links. The list of MPRs of every wireless interface, for each scenario, is displayed in Table 2.

Interface	I	II	III
hybrid1:wlan0	w_1	w_1	–
hybrid2:wlan0	w_2	w_2	w_2
wless1:wlan0	–	w_2	w_2
wless2:wlan0	w_3	w_1	w_1
wless3:wlan0	w_2	w_1	w_1

Table 2: MPRs selected by each wireless interface, for each scenario.

It can be observed that all links are MPR links, and therefore all are adjacent. In this topology, the presence of a *synch* router (see section 2.3) is thus redundant.

3.2.2 OSPF Flooding

Flooding in the wired network is performed through adjacent links – that means, $S \longleftrightarrow h_1$ and $S \longleftrightarrow h_2$. In the wireless network, flooding is performed:

- through the MPR links (from a wireless router towards its MPR), and
- through all links connecting an interface to a hybrid router (hybrid1 and hybrid2).

4 Experiments and Results

For each scenario (I, II and III), communication between `wless3` and `server` is tested by way of two experiments. Displayed results show the averaged measures over tens of samples (see Appendices A.3 and A.4 for further details on configuration of the experiments):

- Transmission of ICMPv6¹ requests (*pings*) from `wless3` to `server`. The measure of time between the transmission of an ICMP request and its reply corresponds to the Round Trip Time (RTT) of the *ping* through the evaluated path.
- Transmission of a constant bit rate data UDP flow from `wless3` to `server`. Comparison between packets sent and packets received permits to test the quality of the traversed paths and the wireless links that compose them in each scenario. Characteristics of these UDP flows are summarized in Table 3.

Nominal sender bit rate	100 pkts/s
Packet payload	1024 bytes
CBR traffic rate	300 kbps
Flow duration	5 min/flow

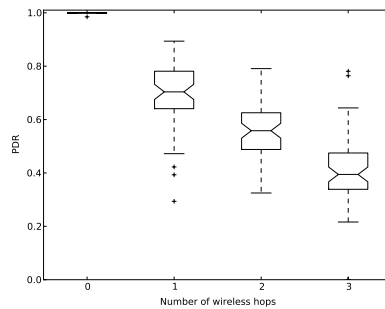
Table 3: Characteristics of transmitted UDP flows.

¹Internet Control Messaging Protocol for IPv6, RFC 4443 [10].

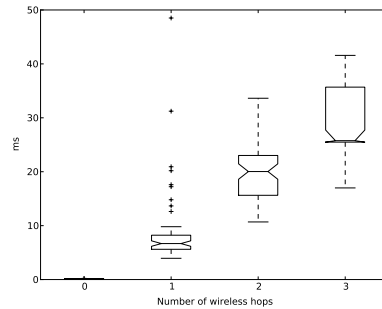
The three considered scenarios are complemented by another scenario in which information is transmitted and measured through the wired link $h_1 \longleftrightarrow S$. Results on this scenario are added for completeness and reference. Section 4.1 presents the results obtained in both experiments, for each scenario, in terms of quality of wireless links. Section 4.2 examines the amount and structure of control traffic used in OSPF for enabling routing of packets within the internetwork.

4.1 Wireless Multi-hop Communication

Figures 3.a and 3.b display the results of the performed experiments, in particular the delay for ICMP requests (*pings*) and the packet delivery ratio of CBR UDP data flows.



(a) Packet Delivery Ratio (PDR)



(b) Round Trip Time (RTT)

Figure 3: **(a)** Packet delivery ratio (PDR) of UDP flows, and **(b)** Round Trip Time (RTT) of ICMP requests, both depending on the number of wireless hops.


Both Figures 3.a and 3.b indicate the degradation of the quality of communication between routers `wless3` and `server` as the number of wireless links between them increases. As expected, the wired link $h_1 \longleftrightarrow S$ has an almost-ideal behavior: 100% PDR and no significant delay. The negative impact of wireless links in the path from source to destination is close-to-linear with the number of traversed wireless links, as shown in Figure 3.a: more than 30% of transmitted packets are lost in the first wireless link, and such percentage increases about a 15% per additional wireless link included in the path. Figure 3.b shows that such degradation is also evident in terms of round trip time (RTT).

Replies to ICMP requests are immediately delivered through a wired link, but the average and the variation of delays grow with the number of wireless links involved – is in the order of tens of miliseconds for 2 and 3 wireless links.


While the impact on communication due to the use of wireless links depends on the specific topology and the network technology that is used, two conclusions can be drawn from these experimental results. As each additional wireless hop in the route of data packets in the network implies a significant degradation of the quality of communication, routing in wireless networks should preserve the principle of shortest (wireless) paths, meaning that the number of wireless links traversed by data packets should be minimized. This confirms the conclusions of simulation-based studies such as [21] which highlights the importance of not sacrificing path optimality for less control traffic.

Moreover, in the context of compound internetworks with both wired and wireless links, it is obvious that 'optimal' does not necessarily mean 'the least number of hops', as implicitly used in previous work such as [6, 11, 2, 3]. Indeed, in compound internetworks, it is better for a path to use wired links than wireless links, whenever possible, even if it means more hops in the end. This observation confirms that metrics used should indeed be able to track link quality, and that OSPF specifications such as [8] [5] [7] should be completed with a standard way to do that in MANETs.

4.2 OSPF Control Traffic Pattern



(a) Bytes




(b) # Packets


Figure 4: Control traffic overhead at `server:eth1`.

Figures 4, 5, 6 and 7 display the evolution of OSPF control traffic transmitted by wireless interfaces `wless3:wlan0` and `hybrid1:wlan0`, on one side, and wired interfaces `hybrid1:eth0` and `server:eth0`, on the other. The five packet formats used in OSPF (Hello, LSUpdate, LSRequest, LSAck and DBDesc, see section 2.1) can be distinguished in these figures. Measures were taken with the topology of scenario I, each point corresponding to the number of packets or bytes sent within an interval of 5 seconds. The traffic load of the internetwork was composed of a CBR UDP data traffic flow from `wless3` towards `server` (see Table 3 for details), and OSPF control traffic. The figures show the structure of such control traffic, both in terms of number of packets and number of bytes, during the first 335 seconds of network operation, *i.e.*, after routers' startup. All interfaces are configured with the same OSPF parameters, in order to facilitate the comparison between control traffic patterns of each of them. See Appendix A.4 for further details.

4.2.1 Hello Packets



(a) Bytes




(b) # Packets

Figure 5: Control traffic overhead at `wless3:wlan0`.


The amount of Hello packets sent by each interface is kept constant along the monitored time. As $HelloInterval = 2sec$, interfaces transmit 2.5 Hello packets per interval of 5 seconds. The length of Hello packets is significantly longer in wireless interfaces (Figures 7.a and 5.a) than in wired interfaces (Figures 4.a and 6.a). For the same number of neighbors, Hellos from `hybrid1:eth0` have 40 bytes while those from `hybrid1:wlan0` have 75.34 bytes. This is due to the

fact that Hello packet format in RFC 5449 [8] includes additional information about link costs, adjacencies and MPR selection, which is added to the format specified in OSPF [17] and OSPFv3 [9].

4.2.2 LSDB Synchronization



(a) Bytes




(b) # Packets

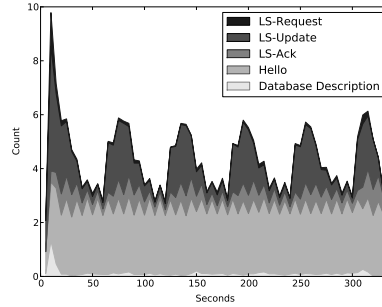
Figure 6: Control traffic overhead at `hybrid1:eth1`.

The existence of ongoing LSDB synchronization processes during the monitored time interval can be noticed in the OSPF control traffic structure by way of the presence of Database Description (DBDesc) packets. The fact that such packets are only present, for wired interfaces, in the first part of the monitored interval (from $t = 0\text{sec}$ to $t = 10\text{sec}$, as shown in Figures 4 and 6) indicates that links become synchronized only when the routers are switched on. In contrast, DBDesc are transmitted in the whole monitored interval for wireless interfaces. This is consistent with the fact that wired links are mostly stable and therefore there is no need to repeat synchronization process after the first LSDB exchange. Wireless links, in contrast, are more prone to packet losses and link failures, and need thus to be synchronized several times during the network lifetime, even in the absence of router mobility. The same phenomenon can be observed with LSRequest packets, which can only be sent during the last phase of the LSDB synchronization process, when the synchronizing neighbors have completed the exchange of DBDesc packets. These observations are also consistent with simulation-based studies such as [21] which have studied the impact of link quality on OSPF control traffic.

4.2.3 Link State Updates, Requests and Acknowledgements



(a) Bytes



(b) # Packets

Figure 7: Control traffic overhead at `hybrid1:wlan0`.

LSUpdate packets contain one or more Link State Advertisements (LSAs). Such LSAs can be either originated by the sending interface, either originated by another interface and flooded (forwarded) by the sending interface. Transmission of LSUpdate packets follows a common pattern in all the interfaces in the internetwork. Thus pattern consists of periodic peaks followed by time intervals (*valleys*) in which the number and size of LSUpdate transmissions is lower and roughly constant.

The time interval between two consecutive peaks corresponds to the value of parameter *LSRefresh*, set to *60sec* for all interfaces. This is the time interval at which an interface floods its topology description (periodically) if there are no topology changes in the meanwhile.

Despite the common pattern in the LSUpdate traffic, several differences can be observed between wired and wireless interfaces. This section concentrates on three particular aspects: peak width, height of valleys between consecutive peaks and transient state (after routers are switched on).

- **Transient state.** Immediately after switching on, wireless interfaces transmit a high number of packets – mostly, LSUpdate packets sent in response to LSRequest packets received during the first LSDB synchronization processes in all wireless links (Figures 7.a and 5.a, between $t = 0sec$ and $t = 50sec$). This amount of transmissions involves traffic rates above

220Bps (1.1kB per interval of 5sec), then decreases and stabilizes in a slightly lower level (maximum peak of 130Bps). The opposite behavior is found in wired interfaces (Figures 4.a and 6.a), in which the initial transient period of low LSUpdate traffic rate (about 26Bps = $\frac{130B}{5sec}$ for `hybrid1:eth0`) is followed by a steady period in which the minimum LSUpdate rate is slightly higher (about 30Bps = $\frac{150B}{5sec}$ for `hybrid1:eth0`). These different behaviors can be explained by the different roles that flooding has over wired and wireless links. Due to their stability, packets sent over wired links are mostly forwarded packets – that is, they come from other interfaces than those involved in the links. In the first instants in which there is no flooding over the network because adjacencies have not been formed in the network and flooding links have not yet been identified, the overall traffic traversing such wired links is temporarily low. The opposite is observed in wireless links.

- **Peak width.** Peaks are narrower in wired interfaces ($\sim 10sec$ for `server:eth1`, $\sim 15sec$ for `hybrid1:eth0`) than in wireless interfaces ($\sim 25sec$ for `hybrid1:wlan0`, $\sim 30sec$ for `wless3:wlan0`). For interfaces attached to wireless links, there is a high probability that a topology change causes a new topology update before the *LSRefresh* interval – therefore, intervals between consecutive transmission of interfaces’ topology descriptions are shorter than *LSRefresh* and the width of the peak increases. In stable wired links, in contrast, intervals between consecutive transmissions are closer to the *LSRefresh* parameter and, therefore, LSUpdate transmission events are less spread in time.
- **Height of valleys.** Besides the peaks caused by transmission of its own topology description, either periodic or as a reaction to a topology change, two other events may lead an interface to transmit Link State Advertisements (LSAs): (i) forwarding of LSAs originated by other interfaces in the internetwork, and (ii) retransmission of LSAs not acknowledged by their intended destinations. Both additional events explain the presence of valleys with significant traffic rate, *i.e.*, a non-zero minimum level of LSUpdate transmissions in the monitored interfaces. In wired (reliable) links such as `server:eth1` and `hybrid1:eth0`, such transmissions are caused by flooding, and involve about 25Bps (127B per interval of 5sec). Wireless interfaces such as `wless3:wlan0` have a minimum LSUpdate transmission rate of about 16Bps (80B per interval of 5sec) caused by LSA retransmissions and flooding.

5 Conclusion

Results from the experiments carried out on the testbed confirm the effect of the presence of wireless links in an OSPF network, confirming prior simulation-based studies concerning control traffic composition. Analysis of OSPF control traffic over the wireless network reveals that even with a very small number of neighbors per wireless interface and static routers, link synchronization processes may involve a continuous and substantial amount of traffic. Reduction of the number of synchronized links is, therefore, essential when using OSPF over ad hoc networks.

Degradation of data communication due to wireless links implies that suboptimal data paths should be avoided in routing on ad hoc networks, as the quality of resulting communication decreases substantially with each additional wireless hop. For internetworks combining wired and wireless networks, wired links should be preferred to wireless links when possible. A standard way to track link quality in compound internetworks is thus necessary in order to complete current OSPF specifications for operation on MANETs. The presented results also point out the importance of leveraging every possible wired connection in the compound AS, which advocates for the use of a single routing protocol in the AS. By using a single protocol, indeed, wired links would be natively included in path computation, without requiring the use of mandatory gateways between different routing domains, which may lead to suboptimal paths (and additional hardware, software and maintenance costs for these gateways).

References

- [1] Cordero, J. A.; Baccelli, E.; Clausen, T.; Jacquet, P. (2011). Wired/Wireless Compound Networking. In: Wang, X. (Ed.) (2011). *Mobile Ad-Hoc Networks: Applications*, InTech Publishers, ISBN 978-953-307-416-0.
- [2] Cordero, J. A.; Clausen, T.; Baccelli, E. (2011). MPR+SP – Towards a Unified MPR-based MANET Extension for OSPF. *Proceedings of the 44th Hawaii International Conference on System Sciences (HICSS)*, Hawaii, HI (United States), January 2011.
- [3] Baccelli, E.; Cordero, J. A.; Jacquet, P. (2010). Optimization of Critical Data Synchronization via Link Overlay RNG in Mobile Ad Hoc Networks. *Proceedings of the 7th IEEE International Conference on Mobile Ad-hoc and Sensor Networks (MASS)*, San Francisco, CA (United States), November 2010.
- [4] Cordero, J. A. (2010). MPR-based Pruning Techniques for Shortest Path Tree Computation. *Proceedings of the 18th IEEE International Conference on Software Telecommunications and Computer Networks (SoftCOM)*, Split (Croatia).
- [5] Roy, A.; Chandra, M. (2010). RFC 5820, *Extensions to OSPF to Support Mobile Ad Hoc Networking*, IETF, March 2010.
- [6] Baccelli, E.; Cordero, J. A.; Jacquet, P. (2009). Multi-Point Relaying Techniques with OSPF on Ad Hoc Networks. *Proceedings of the 4th IEEE International Conference on Sensor Networks and Communications (IC-SNC)*, Porto (Portugal).
- [7] Ogier, R.; Spagnolo, P. (2009). RFC 5614, *Mobile Ad Hoc Network (MANET) Extension of OSPF Using Connected Dominating Set (CDS) Flooding*, IETF, August 2009.
- [8] Baccelli, E.; Jacquet, P.; Nguyen, D.; Clausen, T. (2009). RFC 5449, *OSPF Multipoint Relay (MPR) Extension for Ad Hoc Networks*, IETF, February 2009.
- [9] Coltun, R.; Ferguson, D.; Moy, J. (2008). RFC 5340, *OSPF for IPv6*, IETF, July 2008.
- [10] Conta, A.; Deering, S.; Gupta, M. (2006). RFC 4443, *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*, IETF, March 2006. (Updated by RFC 4884)
- [11] Henderson, T.; Spagnolo, P.; Pei, G. (2005). *Evaluation of OSPF MANET Extensions*, Technical Report D950-10897-1, The Boeing Company, July 2005.
- [12] Clausen, T.; Jacquet, P. (2003). RFC 3626, *Optimized Link State Routing Protocol (OLSR)*, IETF, October 2003.

- [13] Henderson, T. *et al.* (2003). A Wireless Interface Type for OSPF, *Proceedings of the IEEE Military Communications Conference (MILCOM)*, pp. 137-145, IEEE ComSoc, Boston, MA (United States), October 2003.
- [14] Perkins, C.; Belding-Royer, E.; Das, S. (2003). RFC 3561, *Ad hoc On-Demand Distance Vector (AODV) Routing*, IETF, July 2003.
- [15] Jacquet, P.; Laouiti, A.; Minet, P.; Viennot, L. (2002). Performance of Multipoint Relaying in Ad Hoc Mobile Routing Protocols, in E. Gregori *et al.*: *Networking 2002, LNCS*, Vol. 2345.
- [16] Halabi, S.; McPherson, D. (2000). *Internet Routing Architectures*, 2nd Edition, Cisco Press, ISBN 1-57870-233-X.
- [17] Moy, J. (1998). RFC 2328, *OSPF Version 2*, IETF, April 1998. (*Updated by RFC 5709*)
- [18] Oran, D. (1990). RFC 1142, *OSI IS-IS Intra-domain Routing Protocol*, IETF, February 1990.
- [19] Dijkstra, E. W. (1959). A Note on Two Problems in Connection with Graphs, In: *Numerische Mathematik*, No. 1, pp. 269-271.
- [20] F. Baker, M. Chandra, R. White, J. Macker, T. Henderson, E. Baccelli *Problem Statement for OSPF Extensions for Mobile Ad Hoc Routing*, IETF Internet Draft, Spet. 2003.
- [21] E. Baccelli, J. A. Cordero, P. Jacquet, *OSPF over Multi-Hop Ad Hoc Wireless Communications*, International Journal of Computer Networks & Communications (IJCNC) Vol.2, No.5, September 2010.

A APPENDIX

A.1 Testbed Hardware

Networking interface drivers were the following:

- Wired interfaces: Digital Equipment Corporation DECchip 21140.
- Wireless interfaces: Broadcom BCM4306 WLAN.

A.2 Testbed Software

Software used in all computers was as follows:

- Operating System: Ubuntu v.10.04 with kernel 2.6.32.
- Routing Protocol Implementation: `ospf6d` daemon of Quagga/Zebra routing suite v.0.99.15.
 - Wired interfaces: Point-to-point.
 - Wireless interfaces: MANET, as specified in RFC 5449 [8].

A.3 Setup for PDR and RTT Measures

- Routers were switched on between $t = 0sec$ and $t = 2sec$.
- PDR results averaged over 84 iterations.
- `ospf6d` daemon NOT restarted in each iteration.
- UDP flows, started 60sec after `ospf6d` daemon switch-on:

Nominal sender bit rate	100 pkts/s
Packet payload	1024 bytes
CBR real traffic rate	<i>sim</i> 300 kbps
Flow duration	5 min/flow

Table 4: Characteristics of transmitted UDP flows.

- RTT results averages over 60 iterations (ICMPv6 requests).
- ICMPv6 request did not overlap with UDP flows.

A.4 Setup for Control Traffic Measures

- Routers were switched on between $t = 0sec$ and $t = 2sec$.
- Results averaged over 84 iterations.
- `ospf6d` daemon restarted in each iteration.
- UDP flows: see Table 4.

A.5 OSPF Parameters

Table 5: General Simulation Parameters.

Name	Value
<i>OSPF General Configuration</i>	
HelloInterval	2 sec
DeadInterval	10 sec
RxmtInterval	5 sec
AckInterval	2 sec
Jitter	100 msec
MinLSInterval	5 sec
MinLSArrival	1 sec
LSRefreshInterval	60 sec



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

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

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