QoS Routing in OLSR with Several Classes of Service

Jérémie Leguay*+, Vania Conan*, Timur Friedman+
* Thales Communications
+ Université Pierre et Marie Curie, Laboratoire LiP6–CNRS

Abstract

The Optimized Link State Routing (OLSR) protocol is a proactive link-state protocol bafor ad-hoc networks. It uses an optimization called Multi-Point Relays (MPRs) to provide an efficient broadcast structure and to reduce the number of link advertisements. Various propositions have been made to enhance the quality of service (QoS) of the routes that OLSR provides, selecting routes according to metrics such as their expected loss rate or throughput. This paper analyzes the existing propositions and finds them lacking. By building upon OLSR's link announcement mechanism, these solutions permit the advertisement of links with non desirable properties and can only support one class of service. We propose ways to improve one of the propositions, QOLSR. Our QOLSR$^+$ uses MPRs only to create a reliable broadcast structure and employs modified link announcements to increase the number of advertised links.

1. Introduction

Mobile ad-hoc networks, called MANETs [5], allow the spontaneous set up of wireless communication systems. A MANET is composed of mobile nodes that share one or more wireless channels without centralized control. Many routing protocols have been defined for MANETs [10, 4]. However, no standard solution exists for quality of service (QoS) support. This paper addresses the issue of integrating QoS considerations into one of the MANET protocols, the Optimized Link State Routing (OLSR) [4] protocol, in order to provide several classes of service to applications.

QoS provisioning in ad-hoc networks is a challenge when using the IEEE 802.11 MAC layer because of the limited network resources available, the time-varying quality of the radio medium, and the mobility of nodes. Moreover, the fact that the availability of network resources is hard to predict and to guarantee makes QoS solutions based on reservations difficult to achieve. This paper presents solutions that perform *QoS routing*, or routing with QoS constraints. The goal is to find routes that are better than classical shortest hop count paths in ensuring certain path metrics. Indeed, the use of the hop count might lead to poor quality routes that follow long range links, having a high packet error rate, or that transit through heavy loaded areas, presenting a high level of radio interference or a high level of congestion. QoS routing uses metrics from other layers that are commonly called *cross-layer* metrics. A number of such metrics have been proposed: the *expected transmission count* (ETX) proposed by De Couto et al. [6] measures the bidirectional packet loss ratio of links and the metric presented by Iannone et al. [8] combines the packet success rate, the interference level, and the physical bit rate.

OLSR is a proactive link-state protocol that uses an optimization called Multi-Point Relays (MPR) to build an efficient broadcast structure and to reduce the number of link advertisements. Propositions already exist in the literature [2, 7] for QoS routing with OLSR. However, we show in this paper that the current solutions are not suitable for supporting more than one class of service.

The main contributions of this paper are the following. First, it provides a detailed analysis of the existing solutions for QoS routing with OLSR that highlights the proposed modifications to OLSR and their consequences in terms of QoS, network overhead, architecture design and interoperability. Second, it proposes some improvements to support several classes of service. The idea is to chose MPRs only to create a reliable broadcast structure, and to modify link announcements to increase the number of advertised links.

2. Standard OLSR

The Optimized Link State Routing (OLSR) protocol is an augmented version of a pure proactive link-state protocol specially designed for ad-hoc networks. It has been strongly optimized to reduce the network overhead induced by control traffic. In OLSR, each node periodically exchanges information about the network topology with the others in order to maintain its routing tables. OLSR provides hop-by-hop rather than path-based forwarding, and so each node uses its most recent knowledge of the network topology in order to decide on the next hop to which to forward a packet.

One advantage of OLSR from a QoS perspective is that is *proactive*, meaning that it establishes routes before a source asks to forward a packet. This property is attractive for adhoc networks that need to support voice applications, as the connection establishment latency is comparatively small. Another advantage of OLSR is that, as a link-state protocol, route computations are performed with the knowledge of the entire network state. This provides a better support for QoS than do distance-vector protocols.

One principle optimization of OLSR is its use of Multi-Point Relays (MPRs) [11]. Each node selects from its neighborhood a set of nodes called its MPRs. The MPRs are a subset of the node's 1-hop neighbors that is a minimal subset through which one must pass to reach all of its 2-hop neighbors. A node is said to be an MPR if it has been selected in this way by at least one of its 1-hop neighbors. Figure 1 shows an example. The set of node A's 1-hop neighbors is $\{B,C,D\}$. To reach node A's 2-hop neighbors, it is necessary to pass through nodes in one of the following subsets: $\{B,D\}$, $\{C,D\}$, or $\{B,C,D\}$. Since the subset $\{C,D\}$, which is highlighted, is a minimal subset, it is an MPR set. The heuristic for MPR selection is Algorithm 1.

Algorithm 1: $\mathrm{OLSR\_MPR\_selection}\left(i,N_i^1,N_i^2\right)$

```
Input  : A node $i$, $i$'s 1-hop neighborhood $N_i^1$, $i$'s 2-hop neighborhood $N_i^2$.
Output: $M_i$, the MPR set of $i$.
begin
    Add to $M_i$ the nodes in $N_i^1$ which are the only nodes to provide
    reachability to a node in $N_i^2$.
    Remove the nodes from $N_i^2$ which are now covered by a node in $M_i$.
    while $N_i^2 \neq \emptyset$ do
        For each node in $N_i^1$, calculate the reachability, i.e., the number of
        nodes in $N_i^2$ that it can reach.
        Add to $M_i$ the node that provides the highest reachability. In case
        of multiple possibilities, select the node that has the highest
        number of 1-hop neighbors.
        Remove the nodes from $N_i^2$ that are now covered by a node in
        $M_i$.
end
```

MPRs provide two optimizations:

- A virtual tree structure rooted at each node that allows broadcast operations to be performed while reducing the number of messages transmitted. Only MPRs can relay broadcast messages. All OLSR broadcasts of control packets take advantage of this mechanism.
- A limitation in the number of links whose states are advertised, which brings about a reduction in routing protocol overhead. Only MPRs can advertise links, using broadcast messages called Topology Control (TC) messages. Furthermore, not all the outgoing links of an MPR are advertised. The subset of links advertised consists of the links to the *MPR selector set*. This is the set of 1-hop neighbors of a node that have selected this node as an MPR.

From a practical point of view, each node periodically broadcasts control messages, called HELLO messages, to its 1-hop neighbors. Each HELLO message contains the list of 1-hop neighbors of the node. In this way, a node obtains knowledge of its 2-hop neighborhood, allowing it to run the MPR selection algorithm. Also, a HELLO message notifies a node if it has been selected as an MPR by its neighbor.

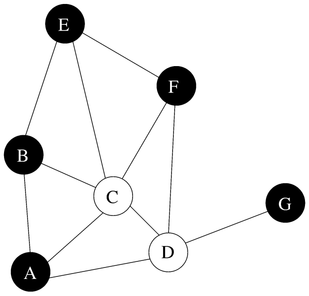

Figure 1. Nodes C and D are MPRs for A.

Having received TC messages and being aware of their 2-hop neighborhoods, nodes are able to compute their routing tables. The route computation algorithm is based on the construction of a spanning tree to provide all possible shortest hop count paths.

3. QoS routing integration in OLSR

Existing solutions for QoS routing with OLSR can be divided into two strategies. One is strictly based on MPR selection. The other, called QOLSR [2] (for QoS OLSR), uses MPR selection and also employs a revised version of TC messages to make QoS information available at the scale of the whole network.

3.1. MPR selection only

As described in Sec. 2, the heuristic for MPR selection aims at selecting a minimal set of 1-hop neighbors that can be used to reach all the 2-hop neighbors. This heuristic can be modified for QoS purposes. Changing the set of MPRs that are selected modifies the set of links that are advertised across the network. Only these advertised links are involved in the route computation process. The strategy here is to chose MPRs such that good quality links are advertised instead of poor quality ones. As an example, Ge et al. [7] changed the basic heuristic for MPR selection to make OLSR find maximum bandwidth paths. The basic idea is that, when there are more than one 1-hop neighbors that cover the same 2-hop neighbors, the one that has a link with the current node with the largest bandwidth is selected.

This strategy has some advantages. It does not break interoperability with standard versions of OLSR and with other versions based on the same strategy but using other heuristics. It also avoids the exchange of additional control information between nodes. Finally, it leads to the creation of shortest paths in terms of hop count. The basic principle of MPR selection, which is to chose an MPR set such that all the 2-hop neighbors are covered, guarantees that the route computation algorithm finds shortest paths.

This strategy also has a number of drawbacks. First of all, it can only take one metric, or one specific combination of metrics, into account. It does not allow for the provision of several classes of service. And, even with a single metric, it leads to the creation of paths that are not necessarily the best paths for that metric, but rather to shortest hop count paths that are better with regard to that metric. Then, since a link, let's say $\overrightarrow{AB}$, is advertised in the network only if B is an MPR of A, a certain number of links are not advertised. As a consequence, some good links may not be considered in the route computation. We study this point in greater detail in Sec 4. Finally, the QoS mechanism based on MPRs goes against the possibility of using asymmetric cross-layer metrics. A node selects its MPRs based on interesting properties of the links that it has with them while it is the reverse links that are advertised in the network with TC messages. Cross-layer metrics may be asymmetric: a metric for the link $\overrightarrow{AB}$ may not be the same as for the link $\overrightarrow{BA}$. Note that Ge et al. [7] assume the use of a symmetric metric.

3.2. QOLSR

Badis et al. [1] have proposed QOLSR, a QoS routing extension to OLSR. It is based on the use of a special heuristic for MPR selection and a modification of TC messages to spread QoS information throughout the network.

For MPR selection, the strategy is to find MPRs that maximize the available bandwidth and minimize the delay toward the 2-hop neighbors. In that way the route computation is performed on a better set of links than in regular OLSR. To apply such a heuristic, nodes need some knowledge about the 2-hop neighborhood. HELLO messages are modified to support the exchange of QoS information between 1-hop neighbors. Each node announces the available bandwidth and the delay for each of its 1-hop neighbors. It can optionally announce other QoS metrics, using an extensible QoS field. Having received such HELLO messages, nodes can chose their MPRs using the heuristic of Algorithm 2.

In QOLSR, TC messages are similar to those of OLSR but they carry QoS information associated with links. The available bandwidth and delay of announced links are embedded in TC messages, and the basic route computation algorithm uses this information. TC messages can also contain additional QoS metrics in the same way that HELLO messages do, and the route computation algorithm can also use this information. Each node constructs a weighted graph, with links weighted according to the supplied metric values. Routes can be computed with an algorithm similar to the one used in standard OLSR, or based on Dijskstra's algorithm.

Algorithm 2: $\mathrm{QOLSR\_MPR\_selection}\left(i, N_i^1, N_i^2\right)$

```
Input  : A node $i$, $i$'s 1-hop neighborhood $N_i^1$, $i$'s 2-hop neighborhood $N_i^2$.
Output: $M_i$, the MPR set of $i$.
begin
    while $N_i^2 \neq \emptyset$ do
        For each node in $N_i^1$, calculate the reachability, i.e., the number of
        nodes in $N_i^2$ it can reach.
        Select a node $z$ from $N_i^2$.
        Add to $M_i$, if not yet present, the node in $N_i^1$ that provides the
        shortest-widest path (path with maximum available bandwidth and
        minimum delay) to reach $z$.
        In case of a tie in the above step, select the node that reaches the
        maximum number of nodes in $N_i^2$.
        Remove from $N_i^2$ the nodes that are now covered by a node in
        $M_i$.
end
```

QOLSR does not define any specific route computation algorithm. However, it suggests one direction. As noted by Qayyum et al. [11], the problem of MPR selection is NP-complete when several metrics are used. To reduce complexity, QOLSR proposes to use the idea introduced by Wang et al. [12] that is to chose a path with maximum bottleneck bandwidth (a widest path). In the presence of more than one widest path, the one with shortest propagation delay is selected. These kinds of paths are called *shortest-widest* paths.

QOLSR has a couple of advantages. First, it can handle several metrics. Although by default QOLSR computes routes with shortest-widest links, other metrics can also be used. Second, QOLSR finds true shortest-widest paths. This is an advantage when compared to the MPR selection method of Sec 3.1.

However, QOLSR also has a number of drawbacks. First, it breaks backward compatibility with other OLSR versions. IETF compliant versions of OLSR cannot understand QOLSR's TC and HELLO messages. Second, bandwidth and delay, which are used as basic metrics, are very difficult to measure when using the IEEE 802.11 MAC layer. To our knowledge, no reliable solutions have been found even if some means of estimating have been suggested [9]. Third, QOLSR suffers from a lack of flexibility due to the obligation to use bandwidth and delay as basic metrics. This somewhat reduces the diversity of possible paths that can be selected with the help of other metrics. Fourth, QOLSR may encounter self interference, causing oscillations in the MPR set, for the two following reasons: the MPR selection is based on metrics that vary a lot, and the same links are used for data and for control. Finally, some links' QoS information might not be broadcast. We have pointed out this limitation for the MPR selection scheme in Sec. 3.1, and it is the same for QOLSR because the advertisements are still based only on the MPRs.

4. Consequences of OLSR optimisations on QoS routing

We have implemented a stand alone simulator to study MPR selection in OLSR. It does not implement the MAC and physical layers, but allows graph-based analysis of algorithms. We generated random static topologies of 500 nodes in a square playground having 1000 meter long edges. We measured density by the average number of nodes per radio coverage area, and we varied the density by varying the radio range.

4.1. Properties of missed links

Recall that the link $\overrightarrow{AB}$ is announced, by the default mode in OLSR and QOLSR, only if B has chosen A as an MPR. Fig. 2 shows simulation results that help us analyse the consequence of this optimization.

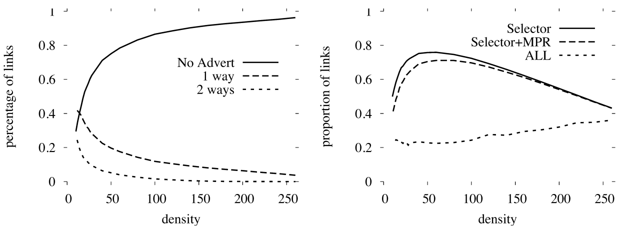

(a) Link directions advertised (b) Useful links missed

Figure 2. Link advertisements.

Fig. 2(a) plots the proportion of links for which: there are no TC message advertisements (*No Advert*), just one direction is advertised (*1 way*), or both directions are advertised (*2 ways*). It shows that the number of links that are not advertised with TC messages increases rapidly with the density. However, this does not mean that the network necessarily suffers. At high density, many nodes are at 1 or 2 hops away from the others and so there is much QoS information available from passing HELLO messages. To understand the effects more clearly, we plot Fig. 2(b), which shows the average proportion of useful links missed by nodes. We define a useful link as a link not in the 2-hop neighborhood of a node. This is a link that a node discovers by receiving TC messages, and that will be useful for route computation. In this plot we compare the different strategies that OLSR proposes to increase the number of links advertised. This is controlled by the parameter TC_REDUNDANCY. MPRs can announce: only their MPR selector set, which is the default mode (*Selector*), their MPR selector set and their MPR set (*Selector+MPR*), or their whole neighborhood (*ALL*). Fig. 2(b) shows that a significant proportion of useful links are significant and that there is a peak at density 50. The first two methods are not so different: the addition of the MPR set does not significantly reduce the number of missed links. Even when MPRs announce their whole neighborhood, there is still a large portion of unadvertised links. Note that the nature of the missing links may depend on the heuristic used for MPR selection.

One alternative solution to changing the TC_REDUNDANCY setting in order to reduce the proportion of non-advertised links would be to use OLSR's MPR_REDUNDANCY parameter. It defines the minimum number of nodes that have to be selected as MPRs, if possible, to reach each 2-hop neighbor. Increasing the redundancy of MPRs allows more links to be advertised, thus reducing the proportion of useful links that are missed, as shown by Fig. 3(a) for several network densities d. However, increasing the MPR redundancy dramatically reduces the efficiency of the broadcast structure. Fig. 3(b) presents the efficiency of the broadcast compared to regular flooding in terms of the number of packets transmitted. Broadcast efficiency is 1 with no MPR redundancy and is 0 when broadcast is achieved through flooding. Increasing MPR redundancy can bring about rapid reductions in broadcast efficiency, especially under lower network densities.

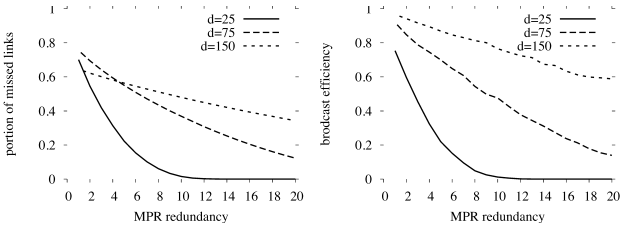

(a) Useful links missed (b) Broadcast efficiency

Figure 3. MPR redundancy.

4.2. Shortcomings of links advertised

Looking at the way MPRs are selected and the consequences for the advertisement of links in OLSR, we see that the currently proposed mechanisms for MPR selection work against the possibility of doing QoS routing with asymmetric metrics. The fact that B chooses A as an MPR because of the interesting properties of $\overrightarrow{BA}$, will lead to the announcement of the link $\overrightarrow{AB}$, which is not necessarily useful.

Another problem with link advertisements based on MPRs is that MPRs are generally nodes that present longer links with their selector compared to those with their regular 1-hop neighbors, as pointed out by Busson et al. [3]. This comes from the fact that MPRs are chosen to reach the maximum number of 2-hop neighbors. We have conducted similar simulations on random topologies by varying the network density. Fig. 4(a) confirms that MPRs are farther than regular nodes. Thus, the link announcement mechanism in OLSR, triggered by the way MPRs are selected, tends to propagate longer links, which would tend to be poorer quality links.

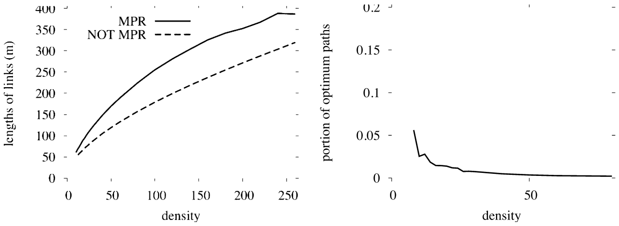

(a) Average length of links with 1-hop neighbors.

(b) Support for multiple classes of service.

Figure 4.

4.3. Support of several classes of service

We conducted an experiment that highlights the fact that the optimizations in OLSR do not offer efficient support for QoS routing with several classes of service. We set two different metrics on the links, m_1 and m_2, taking integer values in $[0 : 100]$. After performing the MPR selection with m_1 using Algo. 2, we analyse the subgraph advertised by OLSR to evaluate the proportion of paths that are optimum regarding m_2. Fig.4(b) shows that, at very low density, some of the paths are m_2-optimum. We believe this is because of the lack of diversity in the paths that exist between pairs of nodes and the very low amount of non advertised links by OLSR. However, the figure shows that the proportion of optimum paths decreases dramatically with the density. We see that two different metrics can not in general be optimised using this MPR selection mechanism.

5. QoS routing with several classes of service

This section focuses on adaptations of QOLSR to support several classes of service. We propose improvements that allow an increase in the amount of useful QoS information advertised in the network while keeping the network overhead low. Here we have conducted a graph-based evaluation that will need to be validated with networking simulations. We refer to our proposal as QOLSR$^+$.

5.1. Design proposition

QOLSR$^+$ incorporates the following modifications:

- MPRs are chosen only to create a stable and reliable broadcast structure. Further studies have to be performed to decide which metric or combination of metrics can be employed to create such a structure.
- TC messages advertise bi-directional QoS information. Since nodes periodically obtain, by HELLO messages, QoS information about the reverse direction of their out-going links, they are able to announce them bi-directionally.
- Only MPRs can send TC messages to advertise all their 1-hop neighbors. If a neighbor is also an MPR, the link is advertised only by the node with the higher IP address.
- Links are just announced once. A mechanism is used to avoid a link being advertised by both end nodes. Simply, the one chosen is the one that has the highest main IP address.

5.2. Evaluation

We evaluate this solution with respect to the number of useful links missed and the network overhead induced. In order to analyse all the strategies for link announcements, we compared the three options seen in Sec. 4 with and without the bi-directional advertisement mechanism of QOLSR$^+$, denoted in the plots by *B*. Recall that QOLSR$^+$ uses the option ALL B. We also compared this proposition to a very simple one, called *NOMPR*, where simply every node sends TC messages with bi-directional QoS information advertisements. We ran simulations to evaluate the different approaches with the parameters used in Sec. 4.

Fig. 5 shows the simulation results. We see from Fig. 5(a) that QOLSR$^+$ still misses a portion of useful links in contrast to NOMPR, which advertises all the links. However, this portion is very low. Less than 10% for densities under 100. Fig. 5(b) shows that solutions like QOSLR$^+$, that use MPRs to trigger link advertisements, lead to a lower number of TC messages sent when the density increases. This is in comparison to NOMPR, which engenders a high constant number of topology messages starting at density 30. This property makes QOLSR$^+$ more scalable than NOMPR. We also plot, respectively in Fig. 5(c) and Fig. 5(d), the average and the variance of TC message sizes. We see that QOLSR$^+$ leads to an average size always lower

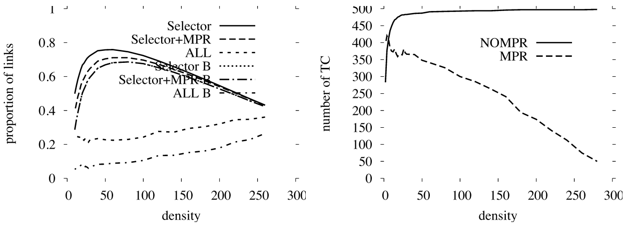

(a) Useful links missed

(b) Number of TCs

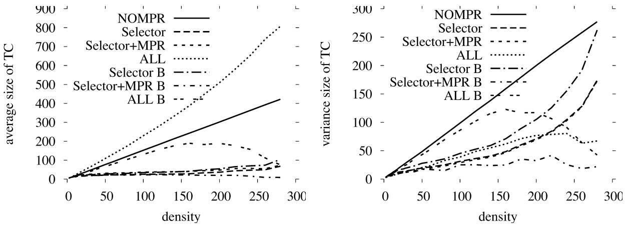

(c) Average TC size

(d) Variance in TC size

Figure 5. QOLSR$^+$ performance.

than NOMPR and QOLSR$^+$ without bi-directional advertisements. The advantage of bi-directional advertisements is clearly seen at high density where the average TC size explodes with unidirectional advertisements while it decreases for QOLSR$^+$. Regarding the variance, it is again lower in all the cases for QOLSR$^+$ compared to NOMPR. The variance for QOLSR$^+$ without bi-directional advertisements explodes at high density while it decreases for QOLSR$^+$. All these results are encouraging and show that we can improve the scalability of the protocol if we tolerate having a very small portion of links not advertised.

6. Conclusion and future work

In this paper, we have analyzed OLSR's basic mechanisms to better understand their impact on the possible integration of QoS routing considerations. We have shown that existing mechanisms to support QoS routing in OLSR are not able to support several classes of service. We have proposed QOLSR$^+$, consisting of possible improvements to QOLSR. We showed encouraging preliminary results using a graph-based evaluation that demonstrates that the number of links advertised can be increased while keeping a low network overhead.

In future work, we plan to investigate, via networking simulations, solutions to the problem addressed in this paper and to conduct studies to determine which metric to use for MPR selection to create a reliable, stable and efficient broadcast structure. Other options will also be envisioned to support several classes of service. For example, by maintaining N MPR sets, one for each of the N QoS classes, and simply merging them using the union operator. All the interesting links would be advertised while only one kind of MPR will be used for broadcast. We will also investigate the problem of link $\overrightarrow{AB}$ being advertised because of the good properties of $\overrightarrow{BA}$. A possible solution involves *reverse link* advertisements, in which TC messages contain in-going links of MPRs with their selector set. This solution may be evaluated with networking simulations using real metrics. We leave this evaluation for future work as it is hard to attribute realistic values to a given topology in a graph-based evaluation.

7. Acknowledgments

We would like to thank the funders of this work: LiP6, Thales Communications, Euronetlab, and the ANRT, which provided the CIFRE grant 135/2004. We also thank Hervé Aiache and Sandrine Masson from Thales Communications for their valuable comments.

References

[1] H. Badis and K. A. Agha. Internet draft draft-badis-manet-qolsr-01.txt: Quality of service for ad hoc Optimized Link State Routing Protocol (QOLSR). IETF MANET working group, September 2005. Work in progress.

[2] H. Badis, A. Munaretto, K. A. Agha, and G. Pujolle. QoS for ad hoc networking based on multiple-metric: Bandwidth and delay. In *Proc. MWCN*, 2003.

[3] A. Busson, N. Mitton, and E. Fleury. An analysis of the MPR selection in OLSR and consequences. In *Proc. MED-HOC-NET*, 2005.

[4] T. Clausen and P. Jacquet. RFC 3626: Optimized link state routing protocol (OLSR). IETF, October 2003.

[5] S. Corson. RFC 2501: Mobile ad hoc networking (MANET): Routing protocol performance issues and evaluation considerations. IETF, January 1999.

[6] D. S. J. De Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *Proc. MobiCom*, 2003.

[7] Y. Ge, T. Kunz, and L. Lamont. Quality of service routing in ad-hoc networks using OLSR. In *Proc. HICSS*, 2003.

[8] L. Iannone, R. Khalili, K. Salamatian, and S. Fdida. Cross-layer routing in wireless mesh networks. In *Proc. ISWCS*, 2004.

[9] L. L. M. Déziel. Implementation of an IEEE 802.11 link available bandwidth algorithm to allow cross-layering. In *Proc. WiMob*, 2005.

[10] C. Perkins, E. Belding-Royer, and S. Das. RFC 3561: Ad hoc on-demand distance vector (AODV) routing. IETF, July 2003.

[11] A. Qayyum, L. Viennot, and A. Laouiti. Multipoint relaying technique for flooding broadcast message in mobile wireless networks. In *Proc. HICSS*, 2002.

[12] Z. Wang and J. Crowcroft. Quality of service routing for supporting multimedia applications. *IEEE Journal on Selected Areas in Communications*, 14:1265–1275, September 1996.