How to Reduce and Stabilize MPR sets in OLSR networks

Leonardo Maccari
DISI – University of Trento
Trento, Italy
Email: leonardo.maccari@disi.unitn.it

Renato Lo Cigno
DISI – University of Trento
Trento, Italy
Email: locigno@disi.unitn.it

***Index Terms*—MPR Selection, Ad-Hoc networks, Mesh Networks, Wireless Routing, Connected Dominating Sets.**

***Abstract*—MPR selection is one of the most important and critical functions of OLSR. The OLSR standard specifies an algorithm that has good local properties in terms of number of MPR selected but does not use available information in order to reduce the global number of MPR nodes. MPR selection affects many network properties, from the actual logical topology, to the routing efficiency, to the protocol overhead and the broadcast/multicast delivery. This paper proposes and evaluates two simple modifications to the MPR selection strategy, which are oriented to global properties rather than local 'optimality'. The results presented show that even marginal modifications of the heuristic lead to a performance improvement, with, for instance, a reduction of up to 15% in the number of control messages required to maintain the topology, a relevant gain specially when obtained without introducing any overhead in control messages.**

I. INTRODUCTION

Multi Point Relays (MPR) have been introduced in 2000 [1] and were later on used as one of OLSR most important components, being the basis of its broadcasting system and the single feature that reduces most the protocol overhead. Since the number of Topology Control (TC) messages generated in the network is proportional to the number of MPR nodes, many efforts have been done all these years to analyse the MPR selection algorithm. The algorithm is completely distributed, each node performs the selection locally with the information coming from its 2-hop neighbors. The problem has been shown for a single node to be equivalent to a classical covering set problem, that is NP-complete. Nevertheless OLSR introduces a heuristic that works quite well in the general case giving solutions close to the optimum. Maybe for this reason the attention has been concentrated on other open issues, such as performing the choice of MPRs in order to enhance redundancy or using MPRs as a building block to create a connected dominating set (CDS). The set of MPR nodes can be a starting point to find a smaller CDS by reduction, thus approximating a minimum-CDS. Related literature is discussed in detail in Sect. VI.

In this paper we approach the problem from another point of view and we show that local optimization is not a sufficient condition to achieve the principal goal of MPR selection, that is, to locally choose MPRs in order to minimize their global number as well as the protocol overhead induced by TC broadcasts. We propose a solution that, with a very simple modification to the original algorithm, performs better than the so-long used heuristic without introducing any overhead (neither as transmission overhead nor as computational complexity). For this reason it is suitable for mobile networks. We show through simulations that the heuristic can be modified in order to reduce the global number of MPR from 8% to 15% in different conditions. With a second technique we introduce, we show that with a minimal overhead we can obtain a relevant gain in the route stability of OLSR.

From the analysis of this work two key observations emerge that are of general interest when investigating MPR selection algorithms:

- The main goal of the algorithm should not be the minimization of the number of MPR a node chooses for itself, but the global number of MPR nodes in the network;
- Studying this topic with computer simulations may hide some of the difficulties introduced by real-world scenarios. Some common pitfall are as discussed in Sect. III

The rest of the paper is organized as follows. Sect. II introduces the role of MPR nodes and the algorithms used in OLSR to select them; Sect. III exploits some example scenarios to lay the ground for the techniques we introduce in Sect. IV, which explains our proposals and the line of reasoning that leads to them. Sect. V presents results and comparison with standard OLSR through simulations, while Sect. VI compares them with the state of the art and discusses our findings and contribution. Sect. VII and Appendix A conclude the paper with a final discussion and details on the simulations and mobility models.

II. MULTI-POINT RELAYS

We recall here only the basic concepts related to MPR selection, since it is a subject very well treated in the literature (see for instance [2] for a survey).

In a wireless multi-hop network sending broadcast messages is a fundamental function in order to build the topology of the network and the routing tables. The most naive way is to instruct every node to rebroadcast every message they receive (excluding copies). In a network of n nodes this will generate a total of n packets for each broadcast message. OLSR uses MPRs to reduce this overhead. As defined in [1] the MPR set $M(x)$ of a node x is an arbitrary subset of its symmetric 1-hop neighborhood $N_1(x)$ which satisfies the following condition: every node in the 2-hop neighborhood $N_2(x)$ of x must have at least a symmetric link towards a node in $M(x)$. More formally:

$$\{x \cup N_2(x)\} \subseteq \bigcup_{u \in M(x)} N_1(u)$$

We say that if $u \in M(x)$ then u "*covers*" some of the 2-hop neighbors of x. Clearly $M(x) = N_1(x)$ is always a solution that corresponds to naive broadcasting.

Once x has selected its MPRs it will communicate them that it has become one of their *MPR selectors*. If x wants to send a broadcast message to the network this message will be forwarded only by its MPRs and will reach all the 2-hop neighbors with only a fraction of retransmissions compared to brute-force flooding. If every MPR in the network retransmits the broadcast messages received by its selectors (possibly avoiding duplicates), broadcast messages will reach all the nodes in the network.

In OLSR, selected MPRs start behaving as follows:

- They periodically generate TC messages containing the list of their selectors;
- They rebroadcast the TCs that are received from nodes that are their selectors.

This allows the construction of shortest path routing tables.

If N is the set of nodes in the network we call the union of all the MPR sets chosen by each node the *global MPR set* M_g:

$$M_g = \bigcup_{x \in N} M(x)$$

Reducing the size of M_g reduces the number of TC packets generated and retransmitted in the network, so it is very important that each node chooses not only a minimal set of MPR among its one-hop neighbors, but also tries to select $M(x)$ so that M_g is also minimal. Choosing the minimal $M(x)$ is NP-complete so OLSR introduces the heuristic described in Algorithm 1.

The *willingness* is a configuration parameter for each node that can be used by a node to encourage or discourage its neighbors in electing it MPR. We call $M^\star(x)$ the optimal (minimal) $M(x)$ and $M^1(x)$ the MPR set that is generated only by the first step of the algorithm. It is intuitive that $M^1(x) \subseteq M^\star(x)$, since the nodes in $M^1(x)$ are a forced choice due to the topology, the rest of the MPRs are freely chosen by the heuristic. Let also $S = ||M(x) \setminus M^1(x)||$ and $S^\star = ||M^\star(x) \setminus M^1(x)||$. $S - S^\star$ measures the sub-optimality of the solution produced by the heuristic. In [1] it is shown that the following relation stands:

$$S \leq \log_2(\Delta) S^\star \tag{1}$$

where Δ is the maximum number of nodes in $N_2(x)$ a node in $N_1(x)$ can cover. In practice, the sub-optimization introduced in the size of the solution produced by the heuristic is limited by a factor of $\log(\Delta)$. This result is calculated without considering step 2.b and 2.c since when there is a tie any choice has the same effect on the coverage of nodes in $N_2(x)$. Moreover, if not diversely configured every node has the same willingness so that step 2.b is often unused. Step 2.c instead is aimed at choosing the neighbor y that has a higher degree irrespective of the coverage related to x. This choice has the effect of giving more redundancy to the distribution of the TC messages.

1) Find all nodes w in $N_2(x)$ that have only one neighbor u in $N_1(x)$, insert u in $M(x)$. Those nodes must be inserted to guarantee full 2-hop connectivity.
2) Repeat the following until all nodes in $N_2(x)$ are reachable using nodes in $M(x)$:
   a) Order every node u in $N_1(x) \setminus M(x)$ based on their *reachability*, i.e. the number of nodes in $N_2(x)$ that are covered by u and are not covered by any other node already in $M(x)$. Insert in $M(x)$ the node with highest reachability;
   b) In case of a tie, insert the one with the highest *willingness*;
   c) In case of a further tie compute the degree of u, that is $||N_1(u) \setminus N_1(x)||$ and insert the one with higher degree. In practice, get the one with the highest number of 1-hop neighbors that are not shared with x.

Algorithm 1: The OLSR MPR selection heuristic

In realistic networks the number of symmetric links is limited, thus the heuristic gives a very practical trade-off with the complexity of more precise methods. For this reason the focus of research has moved in two main directions: studying the properties of networks with redundant MPR sets (considering the effects on reliability and QoS) and using MPRs as a starting point to create a minimal CDS. Building a CDS is an alternative approach to identify a set of nodes that can be used as a backbone to perform broadcast message delivery. A CDS is a connected subgraph of the network such that any node in the network is a 1-hop neighbor of at least a node in the CDS. If all the nodes in the CDS forward the broadcast packets they receive, broadcast communication is achieved. CDSs and MPR sets are "relatives", but not the same. In general, M_g is also a CDS, while the reverse is not always true, so that given an M_g it is always possible to find a CDS with size $\leq ||M_g||$. A minimal CDS is a CDS with minimal size and in general M_g is not a minimal CDS. To achieve broadcast, a node in a minimal CDS must always forward the broadcast packets it receives, while an MPR node only forwards the broadcast received from its selectors. Building a minimal CDS is more complex than building the collection of $M(x)$; indeed, a minimal CDS can not be computed only with local information as MPR sets. This implies that once the global MPR set M_g is chosen it can be reduced using information on the global topology (or at least an ordering function applicable to the nodes). The need of network-wide information may be a limit to the applicability

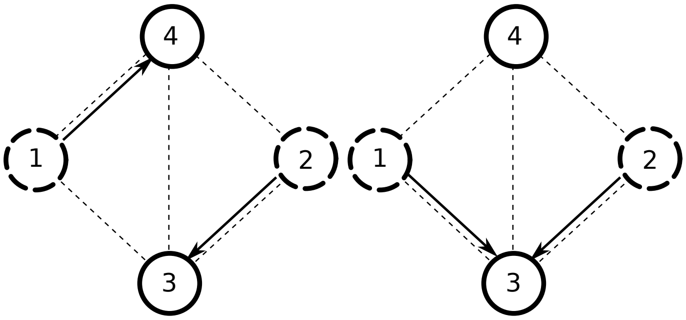

(a) (b)

Figure 1. An example topologies where different MPR selections are possible and lead to different M_gs, MPR have solid borders, other nodes have dashed borders and arrows point from a selector to the MPR

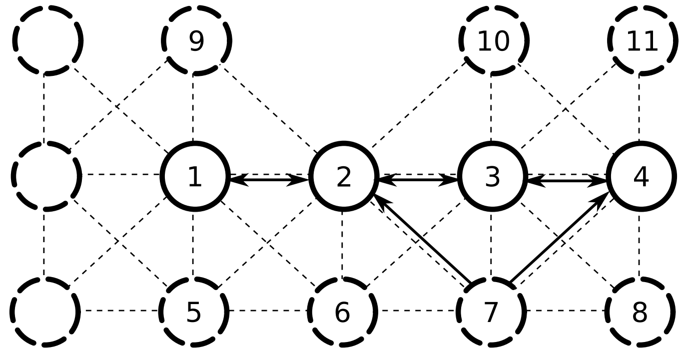

Figure 2. Larger example scenario, $\{1, 2, 3\}$ would be the optimal M_g. Only the MPR chosen by nodes 1,2,3,4,7 are drawn

in mobile networks in which the topology changes frequently, churn is present, and the local information of each node may be out of sync due to packet loss.

III. Revising MPR selection and common pitfalls

One of the most evident limitation of the heuristic described in the previous section is that it does not introduce a deterministic algorithm to resolve the ties. When two potential candidates to be chosen as MPR have the same reachabilty, willingness, and degree, the choice is purely random, and may lead to clearly sub-optimal selections.

Consider the scenario reported in Fig. 1, where dashed lines represent active wireless links and all nodes have the default willingness.

Node 3 and 4 have complete visibility of the network, so they do not need to elect any MPR, node 1 and 2 instead are 2-hop neighbors so they will elect an MPR. To node 1 both node 3 and 4 are equivalent, they have the same reachability, willingness and degree; the same stands for node 2. Since the algorithm does not resolve the tie both the configurations in the figure are allowed. It is perfectly possible that both node 3 and 4 will be chosen as MPR, thus doubling the total number of TC messages *generated* compared to the optimal case. Note however that the number of TC messages *forwarded* is the same in both cases, since node 1 and 2 will not be selectors of each other (no TC is forwarded). Note also that node 3 and 4 chose a minimal MPR set so the heuristic matches the optimal algorithm. This example outlines a first relevant remark: the necessary condition for the global MPR set to be minimal is that each node chooses a minimal local MPR set, but since the local MPR sets are not disjoint, this condition is not sufficient.

A naive modification of the OLSR heuristic would consist in adding a 0-step to Algorithm 1 as follows: 0) initialize $M(x)$ with all the nodes in $N_1(x)$ that have already been chosen as MPR by some other node. This information is available to x at runtime since MPR nodes send TC messages. The logic behind it is that since some nodes in $N_1(x)$ are already MPR, the heuristic can be improved including all those nodes that will generate TC messages anyway. This is a greedy algorithm that works in the simple network of Fig. 1 but fails in more complex networks, mainly due to the correlation between nodes introduced by step 0), which implies that $M(x)$ depends also on other node choices, and hence on the random order in which nodes compute $M(\cdot)$. To show this failure consider Fig. 2 that has an minimal $M_g = \{1, 2, 3\}$.

Imagine that 7 is the first node to choose its MPR set, $N_1(7) = \{2, 3, 4, 6, 8\}$ and $N_2(7) = \{9, 10, 11, 1, 5\}$ so following the new heuristic it will behave as follows:

- Since no node has been elected MPR step 0) has no effect for 7;
- 9 is reachable only through 2 so $M(7) = \{2\}$ for step 1. The only node in $N_2(7)$ that is not covered is now 11;
- Both node 3 and 4 can be used to reach node 11. Their reachability is the same (both have reachability$= 1$), their degree is the same (2).

As a consequence 7 is free to choose any of the two nodes. In case it chooses 3 M_g may be optimal, otherwise 4 will enter M_g. In this latter case, all the other neighbors of 4 will become selectors of 4 due to step 0). Nodes 1,2,3 will be selected in any case since they satisfy step 1 for at least a node but even node 3 will select node 4. We see that step 0 does not always solve the problem observed in Fig. 1, and in other scenarios it can produce more control messages since the number of selectors per MPR is increased. If step 0) is not used M_g may not be minimal anyway since 4 could still become MPR but it would not be chosen by node 3, thus it will never forward any TC message. In practice this simple heuristic does not solve the problem illustrated in Fig. 1: it makes worse. This simple example shows how improving the global performance of a distributed MPR selection strategy is not trivial, and can easily mislead research. Indeed, hundreds of sample topologies can be found where this heuristic works very well, but many more exist where it fails, thus the algorithms definition must follow a more formalized and abstract reasoning to minimize the possibility that performance improvements are indeed limited to some cases or naive topologies.

The greedy algorithm introduced with step 0) does not perform well because it interferes with the original heuristic. With the exclusion of node 7, which is the first to select the MPRs, all the other nodes are forced to select an MPR set not only larger than the minimal one, but also larger than the one selected without considering step 0). A more careful analysis also shows that this solution break the upper bound given by (1), thus destroying all the properties that derive from it.

Before we introduce another, more promising solution we want to describe a pitfall that can partially hide this problem when studied on network simulators. We have said that in case of tie there is no deterministic algorithm to choose the MPR. In the simulator code we used (it is the INET code in Omnet++ simulator that was derived from the NS2 implementation of OLSR and is also present in NS3) the nodes in $N_1(x)$ are stored in an array that is filled when a new node is discovered, it is thus ordered by the age of the nodes as seen by node x. The age of a node is the time passed from the reception of the first HELLO message. OLSR uses a random jitter in order to avoid the collision of HELLO messages, the jitter is limited to a fraction of the time interval between two HELLO messages. Consider the network of Fig. 1 and imagine that the nodes are switched on following their numeric ID. The effect of the random jitter in a real network is irrelevant compared to the time needed to switch on the nodes plus the boot time. As a consequence, both node 1 and 2 will consider node 3 as the oldest node, and the tie will be resolved twice in favour of 3, thus generating a minimal M_g. However, changing the boot sequence for the nodes will produce different effects, not to consider the impact of mobility and churn, so that in real networks the problem we outlined is present and has a non marginal impact. In simulated networks the nodes generally boot at the same instant (it is the case for INET in Omnet++ and NS3). Among node 3 and node 4 the one which chooses the shortest random jitter will be the oldest node for both node 1 and node 2. This in practice completely hides a real problem in the simulation environment. This is a conceptual error that we found in the two widely used open simulation environments that has an impact in the study of the MPR distribution.

Last but not least, during this work we isolated two software bugs that were present in the original OLSR code developed in 2006. Both survived up to now in INET, NS2 and one of them also in NS3 code[1] and caused an inefficient MPR selection.

IV. Proposed Strategies

We propose two strategies, the first one uses a ranking among the MPRs and aims at minimizing the number of MPR nodes in the network and consequently the control traffic, the second one tries to keep a consistent MPR set and is aimed at increasing the stability of the routes and reducing routing tables computations.

The first strategy starts from the intuition that since the nodes have all the information about their 2-hop neighborhood and the MPRs that have been chosen, they have to locally minimize the overlap. This is accomplished ranking the candidate MPR nodes giving high priority to those that have already been chosen by some other node. Actually, after an initial boot phase the MPRs are globally known, whenever x receives a TC message it updates a data structure where it keeps the IP of the MPR and the number of selectors (derived by the TC messages). This provides an ordered list of the MPRs that have been chosen by more neighbors. When building its MPR set we introduce another step between 2.a and 2.b into the heuristic in algorithm 1 that chooses among the candidates the node with highest rank. This way we keep the set minimal, since we do still respect the reachability order (thus we do not change the upper bound), but we try to concentrate the choice on fewer nodes. This strategy has two positive effects, the first is that we have a smaller global MPR set, the second is that we have less MPR nodes that have been chosen by only one node, so less fluctuations on the choice of MPRs. We call this strategy Selector Set Tie Breaker (SSTB).

The second strategy concentrates on minimizing the changes in the MPR selector sets. The recalculation of the routing table is a costly operation that is performed very often in OLSR. Every time a TC that changes the local view of the whole topology for node x is received, x will recalculate the routing table. In real mobile networks this is an issue that must be considered, since it can severely hit the power consumption of the devices. With this strategy we introduce another step between 2.a and 2.b that chooses the node that were already MPRs in the previous run of the algorithm. We call this strategy Stability Driven MPR Choice (SDMC).

The main difference between the two approaches is that with SSTB x tries to maximize the intersection of the MPR sets using a better metric than the one based on a local-only choice. With SDMC instead x may choose a node with a lower degree, trying to keep $M(x)$ stable in time (but potentially performing a worse topological choice).

We have not included in these strategies considerations on link qualities or specific channel models, as we think that MPR selection strategies should first of all be general and robust: tailoring to specific models and scenarios should be done only for scenario-driven customization or when supported by strong experimental evidence brought up by real measurements.

Furthermore we have not included energy consumption considerations, as lack of energy can be mapped by nodes on their *willingness*, which is already part of the standard algorithm.

V. Performance Evaluation

We performed simulations with the Omnet++ network simulator, the network is composed of 60 nodes with WiFi wireless radios. Three scenarios have been considered, the first is a static scenario with random node positions, the second is a mobile scenario with random way point movement, the third uses a realistic mobility model as introduced by [3] and a ray-tracing pathloss model that considers the presence of obstacles in the area. More details of the simulation environment and aggregated results (Table I) can be found in Appendix A.

We chose to use a simple pathloss model and not a more sophisticated model like for instance [4], for several reasons. First of all propagation models should be supported by large measurement campaigns, otherwise they remain models, but

[1]For the details on the software bugs see http://pervacy.eu/MPR-simulator-bug. The debugged code has been filed to the developers of the simulation platforms.

lack validation for the specific scenario. Just to give a hint to the problem, most 802.11 based receivers implement today some form of packet capture, which have a humongus impact on performance, but no standard simulator today include credible models for captures. Furthermore, in the specific problem at stake, considering for instance fast fading in a mobile scenario will only give rise to the need of much longer simulations to average out the effects of fading, without increasing the quality of results. Indeed, the MPR selection strategies should not be affected at all by fast and transient phenomena, as the backbone selection must adapt to topology changes and not to random fluctuations of one communication channel. This is per-se another line of research (stability and robustness of the MPR selection), mixing it with the performance issues discussed in this paper is out of scope.

For each simulation and both strategies we measured:

- The average number of MPR nodes during the simulations, sampled every second,
- The selector set size for every MPR,
- The total number of TC messages generated,
- The total number of TC messages forwarded,
- The total number of times that an MPR receives an HELLO message that changes its selector set
- The fraction of routes in the routing table that have changed sampled every second. This means that every second we scan the routing table and check for every destination if the next hop has changed.

Fig. 3 reports the distribution of the average size of MPR selector set size for the three scenarios and compares the standard OLSR heuristic with SSTB. For each MPR selector set size (x-axis), the number of MPR with that selector size is reported (y-axis). It can be noted that SSTB introduces a higher polarization toward large selector set sizes: there are less MPRs with few selectors and more MPRs with larger selector set. The average size of M_g (not reported in the graphs) is lower than with the standard OLSR procedure; it depends on the topology of the network, but oscillates between -8% and -15%.

In Fig. 4 we report the same graphics for the SDMC strategy and we note that SDMC has the opposite effect, there is an increment in the MPR number lower than 5% and a lower polarization. In practice the choice of trying to keep the MPR set stable in time forces the selectors to use an old MPR even when it is topologically less convenient.

In Fig. 5 we report the gain in terms of control messages that are generated using the two strategies compared to OLSR heuristic for the three scenarios. The graphs report the gain expressed as the fraction of TC generated (that is proportional to the reduction of M_g), the fraction of TC messages forwarded and the fraction of the sum of both values. It can be seen that depending on the scenario, SSTB is able to reduce the control traffic of the network of a value roughly proportional to the decrease of the number of MPR selected in the network. Consequently we note that SDMC generates a larger global MPR set, a price to be paid to have a more stable routing.

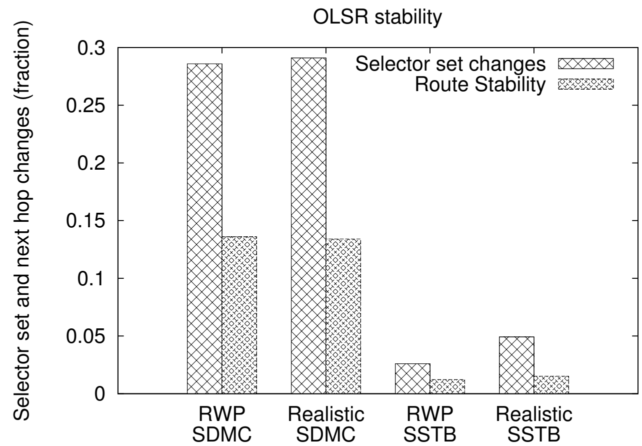

Figure 6. Stability metrics for mobile scenarios

In Fig. 6 we report the comparison of the metrics that deal with the stability of the routing algorithm. Let's clarify it in order to have a better comprehension of their value. Every time an MPR node i receives a HELLO message it increases a counter R_i if the number of its selectors has changed (that is, if the node that sent the HELLO changed its status of selector). In the graph is reported the value $1-\overline{R_s}/\overline{R_{OLSR}}$ where $\overline{R_s}$ is the counter measured with one of the two strategies averaged on all nodes for all the simulation runs. Also, every second we scan the routing table for every node and check for every destination if the next hop has changed. We save the fraction of changed routes in a variable C_i and we report in figure 6 the value of $1-\overline{C_s}/\overline{C_{OLSR}}$. Since the metrics are useful to verify the behavior of the strategies with dynamic topologies, only the results with mobility are reported.

The metrics show that with SDMC the size of the selector set changes almost 30% times less than with standard OLSR. When a node receives a TC and the selector set is changed, this will trigger a recalculation of the routing table, since the topology has changed from the perspective of the node receiving the TC. This does not directly mean that the routing table will change, it may be that a topology change involving two nodes in a remote part of the network will not change the next hop for any of them. Anyway the route recalculation must be done and will have an impact on the CPU load and on energy consumption.

In OLSR when choosing the next hop for a destination there is no mandatory preference for MPR nodes, but it is suggested. If we couple this feature with SDMC, the route stability is positively affected, since at every recalculation we have a decrease of around 13% in the routes for which the next hop address changes. Note that we are not using any quality metric in our experiments, so that the change in the routing tables generated by standard OLSR is not the consequence of a wiser choice for the next hop. It is the result of the fluctuations of the choice of MPR nodes when there are many MPR that have only one or few selectors, so that they easily change their state. Both with SSTB and SDMC this tendency

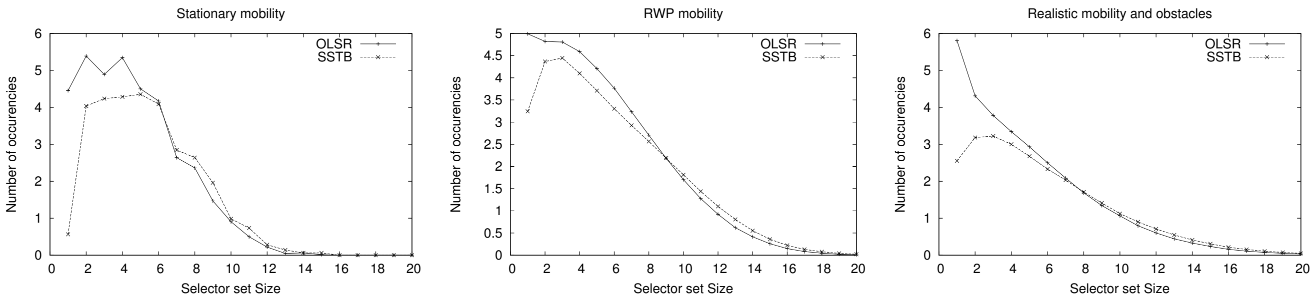

(a) (b) (c)

Figure 3. The MPR selector set size distribution for SSTB compared with the standard heuristic in the three considered scenarios

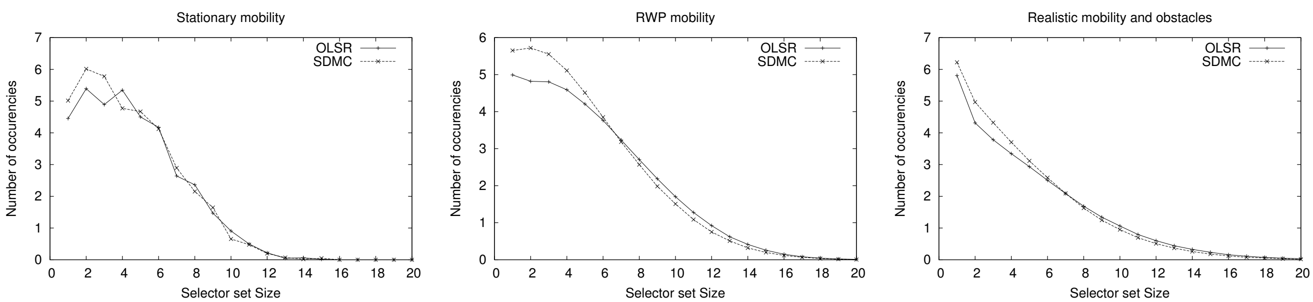

(a) (b) (c)

Figure 4. The MPR selector set size distribution for SDMC compared with the standard heuristic in the three considered scenarios

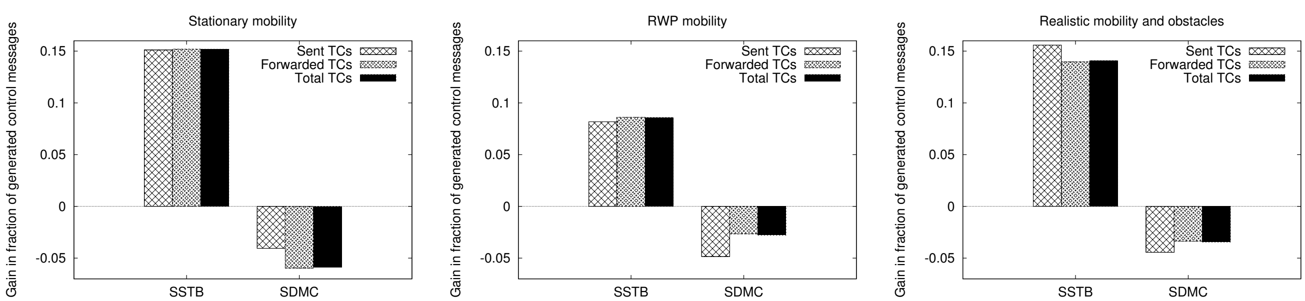

(a) (b) (c)

Figure 5. The variation in TC control messages of SSTB and SDMC compared to OLSR

is limited, in the first case because we have less MPR nodes with little selector set, in the second because we explicitly seek that effect.

VI. Related Works and Discussion

The concept of MPR in OLSR has been introduced in [1] where it is also shown that the problem of locally computing the MPR set is NP-complete. Also the heuristic is introduced together with its upper bound. The performance of MPR-based flooding in terms of number and distribution of MPRs has been largely investigated mostly using theoretical approaches [5], [6]. In [7] four optimizations are reported that try to improve the MPR selection heuristic not in the direction of reducing the MPR number but with a focus on other properties. With that work we share the approach of changing only the tie-break part of the heuristic, which guarantees that the performances are not lowered. The authors also note that when choosing the MPR set, the first step of the algorithm accounts for 68% of the chosen nodes, so there is room for improvement only in the remaining 32%. In this work we have shown that there is room for more improvement since locally optimizing $M(x)$ is not sufficient to minimize M_g.

The good performances of the heuristic have encouraged the researchers to explore other properties, rather than the minimization of the global MPR set. In [8] the concept of MPR is extended in order to take into account energy-preserving strategies, while another set of works introduces QoS metrics in the decision (see [2] for a survey). Another direction is taken by [9] that introduces a basic algorithm to produce a CDS out of M_g pruning some nodes. It also shows that the number of nodes involved in forwarding a broadcast message is very close in both cases, since the reduction of the size of a CDS compared with M_g is compensated by using a source-independent forwarding rule. In the same direction goes [10], which improves the initial idea optimizing the CDS formation. Interestingly the suggested improvement starts from the same observations that we did on the network of Fig. 1 but does not suggest the reduction of the MPR set as we did, instead it focuses on how to reduce the CDS once an over-sized global MPR set has been chosen.

A new improvement and a comparison of the CDS based techniques can be found in [11]. In the scenarios we tested, SSTB produces higher gains than the ones shown by CDS-based techniques in [11] in reducing the forwarding nodes without introducing any real complexity. We did not directly compare our proposal with CDS-based techniques for several reasons. First, we do not strictly depend on the introduction of a global ranking to locally minimize the overlap of MPR sets. We introduced it since this information is already available with OLSR. Second, we do not interfere with how TC messages are forwarded since we respect the MPR approach, which makes our approach compatible with OLSR. Third, we note that the cited CDS-based works deal with static nodes. The improvement introduced in [11], for instance, relays on more complex topological evaluations that add complexity to the original choice and could introduce instability when the topology continuously changes. Fourth, the techniques do not exclude each other: reducing the initial M_g will also improve the construction of a minimal CDS if needed.

More recently, the problem of minimizing M_g set has been studied in [12] and [13], which introduced a centralized algorithm to find the composition of all the possible $M(x)$ with total minimal size. The same authors in [14] propose a QoS based MPR selection scheme that introduces the notion of inefficient EMPR (an MPR with few selectors), but the proposed solution diverges completely from our simple approach. Finally, in [15] it is presented a cooperative MPR selection algorithm in which nodes are split in master and slave roles and the MPR choice is performed accordingly. It is the work that goes closest to our approach but leaves many open points: how to decide the fraction of master nodes among the total, the election of those nodes and consequently the update of this choice with dynamic networks. Our approach is much simpler and needs no coordination.

A further note must be added on choosing a redundant MPR set. In certain cases this may be a choice aimed at increasing the fault-tolerance, or at having richer information to calculate quality routes. In the first case the correct starting point must be a minimal set, in order to be able to fine tune the redundancy with configuration parameters. In the second case the MPRs used for packet forwarding can be a superset of the ones used for flooding, the latter set must be kept minimal.

VII. Conclusions

The selection of MPRs is known to be an NP-complete problem even at the local scale of a single node choosing its own MPRs. In the OLSR standard, the selection is based on a heuristic algorithm that has the property of limiting the number of MPRs selected within a tight, logarithmic bound from the minimum number of MPRs required. This algorithm, however, disregards entirely the minimization at a global level of the number of selected MPRs, and consequently produces a larger overhead for OLSR itself.

We have proposed two different heuristics that, with simple changes from the standard algorithm, are able to control also some global properties of the the entire network. Both modifications simply change the strategy of breaking ties so they do not change the local properties of the standard algorithm and introduce no computational overhead on nodes. The first aims at reducing the protocol overhead, while the second goal is stabilizing the routing tables.

Simulation results on different topologies, with and without mobility, show that in all the considered cases the novel algorithms fulfil their goals with significantly improved performance. In defining the simulation experiments we also identified a conceptual modelling limitation of the Omnet++ version of OLSR, which we corrected and made available to the community as a patch to the standard release.

Future work includes addressing more sophisticated heuristics that can lead, through the proper selection of MPRs, to desired global properties of the logic topology built by OLSR.

Acknowledgements

This work is funded by The Trentino programme of research, training and mobility of post-doctoral researchers, incoming Post-docs 2010, CALL 1, PCOFUND-GA-2008-226070. For more information see http://pervacy.eu.

References

[1] A. Qayyum, L. Viennot, and A. Laouiti, "Multipoint Relaying: An Efficient Technique for Flooding in Mobile Wireless Networks," INRIA, Tech. Rep. RR-3898, 2000. [Online]. Available: http://hal.inria.fr/inria-00072756

[2] O. Liang, Y. A. Sekercioglu, and N. Mani, "A survey of multipoint relay based broadcast schemes in wireless ad hoc networks," *Communications Surveys Tutorials, IEEE*, vol. 8, no. 4, pp. 30 –46, quarter 2006.

[3] M. Musolesi and C. Mascolo, "A community based mobility model for ad hoc network research," in *Proceedings of the 2nd international workshop on Multi-hop ad hoc networks: from theory to reality*, 2006.

[4] J. Gozalvez, M. Sepulcre, and R. Bauza, "Impact of the radio channel modelling on the performance of vanet communication protocols," *Telecommunication Systems*, pp. 1–19, 2010. [Online]. Available: http://dx.doi.org/10.1007/s11235-010-9396-x

[5] P. Jacquet, A. Laouiti, P. Minet, and L. Viennot, "Performance Analysis of OLSR Multipoint Relay Flooding in Two Ad Hoc Wireless Network Models," in *The second IFIP-TC6 NETWORKING Conference*, Pisa, Italie, 2002.

[6] A. Busson, N. Mitton, and E. Fleury, "Analysis of the multi-point relay selection in olsr and implications," *Challenges in Ad Hoc Networking*, pp. 387–396, 2006.

[7] B. Mans and N. Shrestha, "Performance evaluation of approximation algorithms for multipoint relay selection," in *Proceedings of the 3rd Annual Mediterranean Ad Hoc Networking Workshop*, 2004, pp. 480–491.

[8] J. Lipman, P. Boustead, J. Chicharo, and J. Judge, "Resource aware information dissemination in ad hoc networks," in *The 11th IEEE International Conference on Networks (ICON2003)*, October 2003.
[9] C. Adjih, P. Jacquet, and L. Viennot, "Computing connected dominated sets with multipoint relays," INRIA, Tech. Rep. RR-4597, 2002.
[10] J. Wu, "An enhanced approach to determine a small forward node set based on multipoint relays," in *Vehicular Technology Conference, 2003. VTC 2003-Fall. 2003 IEEE 58th*. IEEE, 2003.
[11] Y. Fu, X. Wang, W. Shi, and S. Li, "Connectivity based greedy algorithm with multipoint relaying for mobile ad hoc networks," in *The 4th International Conference on Mobile Ad-hoc and Sensor Networks*. IEEE, 2008, pp. 72–76.
[12] K. Yamada, T. Itokawa, T. Kitasuka, and M. Aritsugi, "Redundant tc message senders in olsr," *IEICE TRANSACTIONS on Information and Systems*, vol. 93, no. 12, pp. 3269–3272, 2010.
[13] T. Kitasuka and S. Tagashira, "Density of multipoint relays in dense wireless multi-hop networks," in *Networking and Computing (ICNC), 2011 Second International Conference on*. IEEE, 2011, pp. 134–140.
[14] T. Koga, S. Tagashira, T. Kitasuka, T. Nakanishi, and A. Fukuda, "Highly efficient multipoint relay selections in link state qos routing protocol for multi-hop wireless networks," in *World of Wireless, Mobile and Multimedia Networks & Workshops, 2009. WoWMoM 2009. IEEE International Symposium on a*. IEEE, 2009.
[15] K. Yamada, T. Itokawa, T. Kitasuka, and M. Aritsugi, "Cooperative mpr selection to reduce topology control packets in olsr," in *TENCON 2010-2010 IEEE Region 10 Conference*. IEEE, 2010.
[16] C. Sommer, D. Eckhoff, R. German, and F. Dressler, "A Computationally Inexpensive Empirical Model of IEEE 802.11p Radio Shadowing in Urban Environments," in *8th IEEE/IFIP Conference on Wireless On demand Network Systems and Services (WONS 2011), Poster Session*, Jan. 2011.

APPENDIX

In this section we detail the simulation environment and give more detailed results. The network is imagined to reproduce an ad-hoc network of mobile devices carried by human with pedestrian mobility.

The simulation scenario is a 350mx350m area, each node is equipped with a IEEE 802.11 stack, the transmission power allows a communication radius of roughly 70m with a dual-slope pathloss model, the average metric to destination is 2.7 hops.

Speed is chosen uniformly between 0.5 and 2 m/s, the mobility model used is a random way point or a social mobility model as proposed in [3]. It defines groups of people that are preferentially bound to certain areas but exchange elements among them. This mobility model derives its properties from social science results and the statistical properties of the traces generated are close to the ones measured in real experiments. In the realistic scenario in the area are present two squared obstacles that affect the pathloss, the ray-tracing algorithm is the one used in [16], thus the average hop-count is higher and network is less connected. For each static scenario 64 simulations of 300s were run, while for each mobile scenario 32 simulations of 2000s were run.

In Table I we report numeric aggregated values of the performance, expressed as a percentage on the same metric measured with standard OLSR.

Network	MPRs	TC	SSC	RS
Stationary-SSTB	15.3	15.1	-	-
Stationary-SDMC	-4.1	-5.8	-	-
RWP-SSTB	8.3	8.5	2.6	1.2
RWP-SDMC	-4.7	-2.7	28.5	13.5
Realistic-SSTB	15.6	14.0	4.9	1.5
Realistic-SDMC	-4.5	-3.4	29.1	13.4

Table I
% GAIN IN PERFORMANCE METRICS. MPRS = M_g SIZE, TC = TOTAL TC NUMBER, SSC = SELECTOR SET CHANGES, RS = ROUTE STABILITY