

*On the robustness and stability of Connected
Dominating Sets*

Cédric Adjih — Emmanuel Baccelli — Thomas Clausen — Philippe Jacquet

N° 5609

Septembre 1997

Thème COM



*Rapport
de recherche*

On the robustness and stability of Connected Dominating Sets

Cédric Adjih*, Emmanuel Baccelli†, Thomas Clausen‡, Philippe Jacquet§

Thème COM — Systèmes communicants
Projets Hipercom

Rapport de recherche n° 5609 — Septembre 1997 — 23 pages

Abstract: In this paper, we investigate the effects of mobility, collisions and obsolete information on the performance of connected dominating set (CDS). In particular, we show that neighbor-designated CDS, such as multipoint relay (MPR) in [2, 3], are in general more robust than self-selected CDS such as rule k CDS in [1]. This robustness is especially crucial for applications such as wireless OSPF, where third party topology information may experience arbitrary delay.

Key-words: Mobile networks, connected dominating set, multipoint relays, stability, collisions

* cedric.adjih@inria.fr

† emmanuel.baccelli@inria.fr

‡ Ecole Polytechnique thomas.clausen@polytechnique.fr

§ philippe.jacquet@inria.fr

Sur la robustesse et la stabilité des ensembles dominants connectés

Résumé : Dans ce rapport nous étudions les effets de la mobilité, des collisions et des informations obsolètes sur les performances des ensembles dominants connectés (CDS). En particulier nous montrons que les CDS à désignation de voisins tels que les relais multipoints (MPR) décrits dans [2, 3] sont en général bien plus robustes que les CDS auto-sélectionnant, tels que "rule k " dans [1]. Ceci est particulièrement crucial dans des applications tels que OSPF mobile où les informations de topologies en provenance de tiers peut prendre des délais arbitraires.

Mots-clés : Réseaux mobiles, ensembles dominants connectés, relais multipoint, stabilité, collisions

1 Introduction

Mobile ad hoc networks rely on radio transmissions that generally feature much less capacity than wired media. On the other side constant mobility requires a very frequent update of topology and routing information. For comparison, wired OSPF requires an update every 20 minutes while its wireless counterpart would need an update every ten seconds, if not every second. Topology control therefore burdens the network, and without proper solutions, the network may collapse as the traffic generated in order to manage the routing function on its own is overloading.

An important component of a solution to this problem is an optimization of the broadcast overhead. To this end, the task of relaying a broadcast packet is given to a subset of nodes in the network. The smallest the relay set is, the less costly will the broadcast be, and more bandwidth is available for user data communications. However, a broadcast aims at all nodes in the network receiving the broadcast data. Therefore one should not reduce the relay set too much: in other words, the relay set must form a connected dominating set (CDS).

Finding the smallest relay set is an NP-hard problem. On the other hand, a relay set that is too small may not be robust enough in face of lost transmissions. In the following, we analyse the performance a few heuristics and their resilience to errors and network unstabilities.

There are two kinds of heuristics, aiming to form connected dominating sets: self-selected dominating sets on one hand and neighbor-designated dominating sets on the other hand.

In self selected dominating sets, the relay nodes *select themselves* as part of the connected dominating set. Based on neighborhood information (one-hop, two-hop, or more) they decide that they belong to the relay set. Classical examples of such an approach are the algorithms of Wu Li [1].

On the other hand, in neighbor-designated dominating set, some nodes *are selected* as members of the connected dominating set by their neighbors. The classical example of such an approach is the MultiPoint Relaying [3] used in Optimized Link State Routing (OLSR) [4]. On top of this selection, a self-pruning strategy is used by each node in order to determine whether it must relay a received broadcast packet, or not.

Two types of self-pruning strategies can be used: (i) an off-line strategy, determined once and for all at the time of MPR selection, or (ii) an on-line strategy, determined during the broadcast, as done in MPR flooding [3].

2 Heuristics

In both heuristics, it is assumed that every node knows its neighbor nodes, via periodic hello exchange. It is also assumed that nodes periodically advertize their neighbor list, either in hellos or in "LSA-like" packets, as in OSPF. Therefore the nodes are also aware of their two-hop neighborhood (but this information may come with some delay). For a node A , we call $\mathcal{N}(A)$ its one-hop neighborhood and $\mathcal{N}_2(A)$ its two-hop neighborhood. Both heuristics

have been proven to provide an effective dominating set (all the nodes are covered), when neighborhood information is exact, and transmissions are error-free – without collisions.

2.1 Self-Selected Dominating Set

In this paper, the heuristic we will consider for self selection is the typical rule k of Wu and Li. This heuristic makes a node A select itself as part of the relay set by detecting all its neighbor nodes that have an I.D. which is greater than node A 's I.D. We call this set $S(A)$. The node A selects itself as part of the relay set if the following condition is fulfilled:

- the set $S(A)$ is not a dominating set of the neighborhood of A .

The self-selection is performed everytime there is a notification of change in the neighborhood or in the two-hop neighborhood. We note by the way, that a node A does not have to signal its self selection to its neighbors.

2.2 Neighbor-Designated Dominating Sets

In this paper, the heuristic we will consider for neighbor selection is the MPR selection as done in OLSR [4]. With this heuristic a node A creates an MPR set – a subset of its neighborhood – that must be a dominating set of the two-hop neighborhood of A . Let us denote $\text{MPR}(A)$ the MPR set of node A . MPR selection is then performed with the following algorithm:

1. We start with $\text{MPR}(A) = \emptyset$
2. We add to $\text{MPR}(A)$ all the neighbor nodes that are the unique connector from A to another node in $\mathcal{N}_2(A)$.
3. Until all nodes in $\mathcal{N}_2(A)$ are covered by $\text{MPR}(A)$, we add to $\text{MPR}(A)$ the neighbor of A that covers the most yet-uncovered nodes in $\mathcal{N}_2(A)$.

This MPR selection algorithm has proven to be very efficient and is actually the foundation of the OLSR protocol [4]. Notice that with this approach, a node has to signal MPR selection to its neighbor nodes. This can be done via an additional field in hello messages, for example. A node A which has selected B as MPR node is then called an MPR selector of node B .

2.2.1 Self-Pruning Strategies

There are two categories of self-pruning strategies: on one hand an on-line strategy that applies at the time of the broadcast operations, and on the other hand an off-line strategy that applies beforehand.

The on-line MPR self-pruning strategy is the following, and was first described in [4]:

- A node forwards a broadcast packet if and only if it receives its first copy from an MPR selector.

The off-line self-pruning strategy is the following and is described in detail in [2]. A node A belongs to the relay set if one of the following condition is fulfilled:

- Node A 's I.D. is smaller than all of the I.D. of its neighbor nodes;
- Node A is selected as MPR by the node which has the smallest I.D. among its neighbor nodes.

Note by the way that the on-line self-pruning is equivalent to the off-line strategy, where the I.D. is the sequence order of the retransmissions of the broadcast. The first emitter would have I.D. 0, the second emitter (MPR of the first emitter) would have I.D. 1, etc. Notice that there will be only one node (node 0) that is a local minimum, that makes the on-line self-pruning more efficient than the off-line strategy (less relay nodes).

In the literature the on-line pruning MPR flooding is simply known as the MPR-flooding, and the off-line pruning is also known as the MPR-CDS flooding.

2.3 Previous Analysis

The idea of improving the reliability of overhead-optimized flooding in face of mobility is not new. Dai and Wu have investigated an algorithm called "SBA" that they studied facing mobility, along with CDS "Wu-Li", and CDS "Dai-Wu rule k " [5]. Their study confirms that algorithms based on dynamic self-pruning (like MPR flooding), are naturally more reliable.

To be completely fair with MPR flooding, Wu also considered a "robust" version of MPR flooding, based on the idea developed in SBA [8], which is roughly the same as MPR-COVERAGE=2 in OLSR [4], except that the second coverage transmits only if the first coverage fails to transmit.

Wu and Dai also studied the issue of obsolete neighbor information in another paper [10], and introduced several interesting theories. For example, two sets of neighbors are introduced (i) close ones, and (ii) not-so-close ones. Dominating sets are then created, first using "close nodes", which are less likely to move away fast, and which form a more robust, and denser dominating set. Then, "not so close nodes" are used to complete connectivity. Theoretically, this kind of CDS should be less dense than the previously mentioned "Dai-Wu rule k " CDS, but maybe more robust. Another idea is to explicitly consider parameters such as the delays in topology acquisition, due to the hello mechanism for instance.

In [6] the authors analyze and compare the performance of MPR flooding and self-selected CDS in the case of MANET and OLSR [4].

In this report, we investigate the resilience of each kind of algorithm, *ie* self-selected or neighbour-designated. The simulation scenarii use various perturbations, such as transmission loss or information obsolescence due to mobility. We will first investigate the resilience of these algorithms in face of transmission errors and loss, in Section 4. Then, in Section 5, we will investigate the effects of mobility and obsolete topology information. And finally,

we will generalize the problem of obsolete information in Section 6 by studying versions of self-selected and neighbor-designated algorithms that operate on a purely random basis.

3 The Simulation Model

Let us take a unit disk graph model on a $L \times H$ rectangle (L and H are expressed as multiples of the radio range unit), with N nodes randomly dispatched. We simulate a simplified medium access control scheme: when a broadcast occurs, at any time there is at most one emitter in the network, randomly chosen among the nodes waiting to retransmit. Figure 1 shows the relay set size for different network size in a 2×2 square map. Figure 2 shows the ratio of the relay set over the network size for different network size in a 4×4 square. No collisions, and no transmission errors are considered. The different algorithms show similar performance. As expected, on-line pruning MPR performs better than off-line pruning, and self-selected CDS stands in between. Figure 3 displays the same parameters but on a 6×6 square.

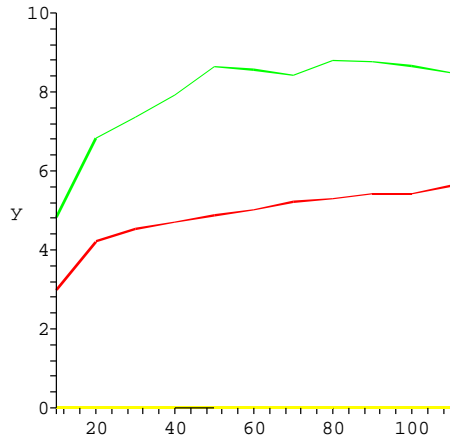


Figure 1: Relay set size for self-selected CDS (middle, green), on-line pruning MPR (bottom, red) on 2×2 square, versus the network size

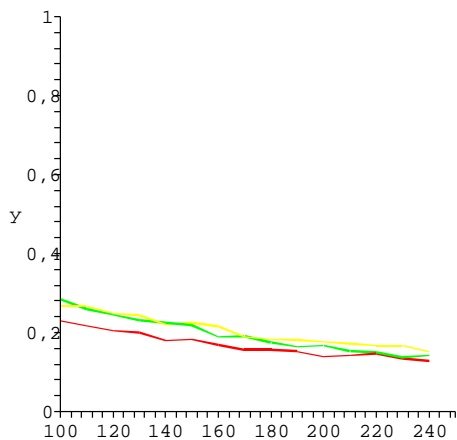


Figure 2: fraction of the network belonging to the relay set for self-selected CDS (middle, green), on-line pruning MPR (bottom, red), off-line pruning MPR (top, yellow) on 4×4 square, versus the network size

4 Resilience to transmission errors

In the previous simulations, we assumed perfect transmission, without any errors. In real cases, collisions and errors occur due to bursty traffic, fading, mobility or simultaneous transmissions. We will show that the neighbor-designated MPR are more resilient than self-selected CDS in face of these errors. The reason for this seems to be two-fold: (i) explicit selection in the neighborhood implies that any neighbor has a relay node in its neighborhood, and (ii) neighbor designation with on-line self-pruning features a significant reserve of potential relays – if some fail others are automatically activated to fill the gap.

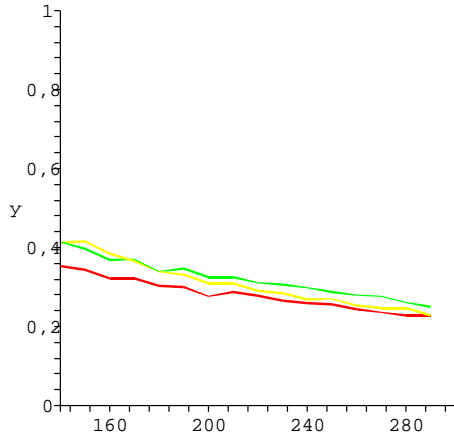


Figure 3: fraction of the network belonging to the relay set for self-selected CDS (middle, yellow), on-line pruning MPR (bottom, red), off-line pruning MPR (top, green) on 6×6 square, versus the network size

4.1 Fixed Transmission Error Rate

We have simulated broadcast flooding with the introducing of an additional parameter p . Quantity p is the reception probability, *i.e.* the probability that a given transmission is correctly received by a given neighbor node. When $p = 1$ we have error free transmission. When $p = 0$ no transmission can be received by any node. We assume that the neighbor information exchange is not subject to error (due to redundant hellos, etc).

Figure 4 displays the fraction of the network receiving the broadcast versus the reception probability p in a network with 100 nodes on a 4×4 square, for the different CDS types. Notice that although self-selected CDS involve more relay nodes than neighbor-designated MPR (see figure 2), its fraction of correct reception drops much faster than with neighbor-designated MPR. To be more precise, when p is close to 1 (which should be the case for stable networks) the simulations suggest that the drop is in $O(1 - p)$ for self-selected CDS and in $O((1 - p)^2)$ for neighbor-designated MPR. This confirms the assumption that self-pruning

adjusts itself when some relay nodes fade. Off-line self-pruning MPR behaves similarly to basic MPR flooding, with some advantage probably due to a larger relay set.

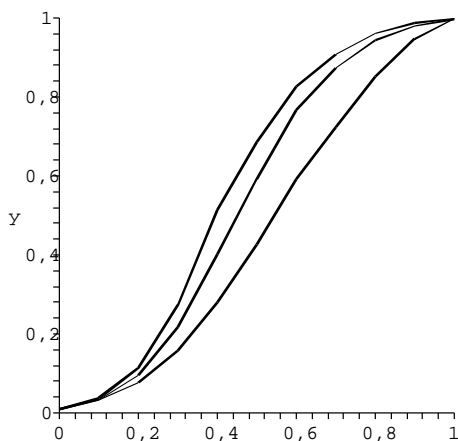


Figure 4: fraction of the network receiving the broadcast versus the receive probability p in a network with 100 nodes on a 4×4 square, for self-selected CDS (bottom), on-line pruning MPR (middle), off-line pruning MPR (top)

Notice that the fraction of correct reception is not exactly zero when $p = 0$, this is due to an artifact of the simulation that always include the initial emitter in the correct reception set.

4.2 Collision rate

Although the simulation model does not include collisions, we have emulated these events by introducing an error rate that depends on the number of contending nodes in the flooding. A packet is received correctly with probability $p = \exp(-\lambda k)$, where k is the number of emitters waiting to transmit, and $\lambda \geq 0$ is a parameter that characterizes the propension of a node to collide with another node. Parameter λ contains altogether the overlap probability of the reception areas, in space and time. We assume that the packets are transmitted with

an exponential jitter. When the jitter increases, quantity λ should decrease, assuming that data traffic and other flooding transmissions don't overlap. In case of overlap with data traffic we should have $p = p_0 \exp(-\lambda k)$.

Figure 5 and 6 show the reception rate for different network scenarii. In figure 6 we have included the performance of the classical full flooding, with which every node retransmits or intends to retransmit the broadcast packet. Notice that the overwhelming mass of retransmitters is counterbalanced by the huge collision rate and the improvement in performance is not very important. In Figure 7 we show that actually the huge number of collision finally renders full flooding inferior to MPR flooding.

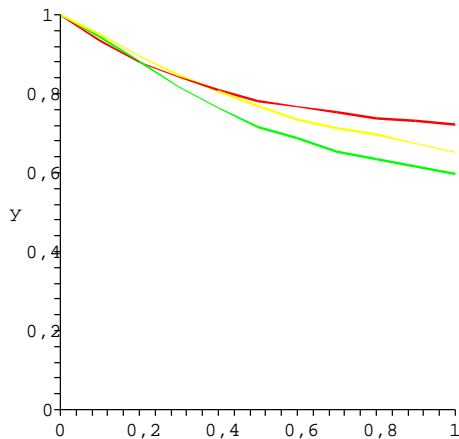


Figure 5: fraction of the network receiving the broadcast versus the collision rate λ in a network with 100 nodes on a 4×4 square, for self-selected CDS (green), on-line pruning MPR (red), off-line pruning MPR (yellow).

The collision rates may not be negligible. In figure 8 we show a real network simulation which evaluates the packet delivery ratio versus the number of broadcast data flows. The network is based on IEEE 802.11 and is simulated in ns-2, with 100 nodes on a 1400×1400 sqm, a typical radio range of 350 m, and random waypoint mobility: 1-2m/s, rest-time 7-15s. The collision rate is here replaced by the number of simultaneous broadcast flows in the network.

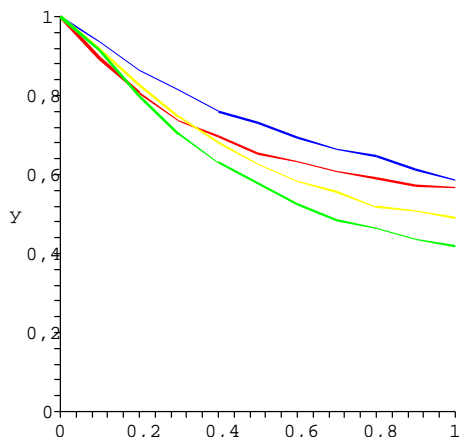


Figure 6: fraction of the network receiving the broadcast versus the collision rate λ in a network with 200 nodes on a 6×6 square, for self-selected CDS (green), on-line pruning MPR (red), off-line pruning MPR (yellow), full flooding (blue).

5 Robustness in Face of Mobility

In this section we investigate the behavior of CDS selection algorithms when the neighborhood information is outdated due to the nodes' mobility. With this respect, two-hop neighborhood information is more fragile than one-hop neighbor information since the former is relayed by LSA packets that are less frequent than hello packets, or by specific hellos that contain the lists of these neighbors. Missing some hellos may cause the loss of the link, but missing packets that contains the two-hop information may outdate the two-hop information without causing the loss of the one hop information. Therefore two-hop information is more fragile than one-hop information.

In the following we address the case where two-hop information is partially obsolete. In this case the CDS algorithms may not converge properly. Self-selected CDS and the neighbor-designated CDS behave differently in this respect. Their behaviors also differ with

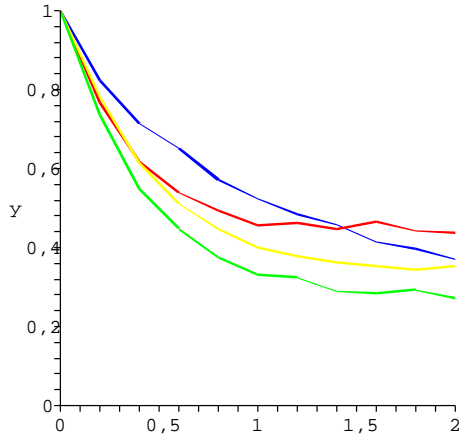


Figure 7: fraction of the network receiving the broadcast versus the collision rate λ in a network with 400 nodes on a 8×8 square, for self-selected CDS (green), on-line pruning MPR (red), off-line pruning MPR (yellow), full flooding (blue).

various information obsolescence scenarii. We will consider two obsolescence scenarii, where two-hop information at time t comes from:

Outdated two-hop-neighborhood: advertized neighborhoods at time t are made of the nodes which have been neighbors during the time interval $(t - T, t)$.

uniformly delayed two-hop neighborhood: advertized neighborhood at time t are exactly made of nodes which were neighbor nodes at time $t - T$;

In the following we assume that nodes are mobile, and that every node follows a random walk. The average speed is of one unit of distance per time, therefore defining the time unit. Of course, nodes don't travel that distance in one time unit, because they are in permanent zig-zag.

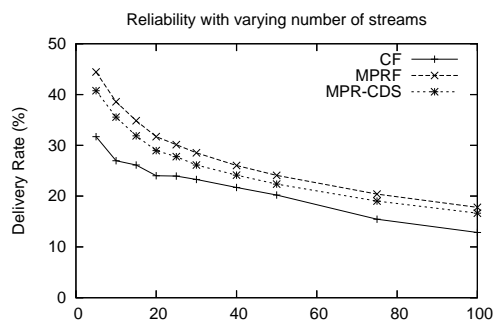


Figure 8: fraction of the network receiving the broadcast versus the number of simultaneous broadcast flows in a network with 100 nodes on a 4×4 square, for on-line pruning MPR (MPRF), off-line pruning MPR (MPR-CDS), full flooding (CF).

5.1 Outdated Two-Hop Information

The simulations of outdated two-hop information show an important degradation of the fraction of the nodes that experience a correct reception when the CDS is self-selected (see figure 9). When the outdateding delay increases the reception ratio tends to 0 (see figure 10 for T up to 10).

The relay size shows a difference here: the self-selected CDS relay set reduces in size while the neighbor-designated relay set remains stable (see figure 11). One possible explanation is that the estimated neighborhood of nodes are larger with window delayed information than with exact neighborhood. Consequently the self-selected CDS will be smaller than the CDS with exact two-hop information. The explicitly selected MPR are apparently less sensitive to this fact and their relay set remains relatively stable.

5.2 Uniformly Delayed Two-Hop Neighborhood

The simulations show little degradation of the fraction of correct reception in the network, with self-selected CDS, when the two-hop neighbor information is uniformly delayed (see figure 12). When such a delay increases the reception ratio tends to be back to 1.

In fact, the performance significantly differs here if we look at the size of the relay sets. The inconsistencies of the two-hop information make the self-selected CDS react by increasing its size while the neighbor-designated MPR remains more stable (see figure 13).

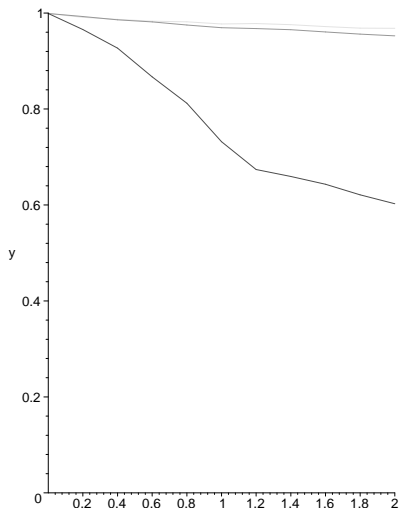


Figure 9: Outdated delayed information: fraction of the network receiving the broadcast versus the outdated window T in a network with 100 nodes on a 4×4 square, for self-selected CDS (bottom), on-line pruning MPR (top, black), off-line pruning MPR (top, grey)

We also ran, out of curiosity, the simulation up to $T = 100$, when the two-hop inconsistency dramatically increases. In this case every node selects itself in the relay set if a the self-selected strategy is used, while with neighbor-designation the relay set keeps significantly below that size. However, in each cases, the correct reception ratio stays very close to 1 (see figure 15).

6 Random CDS and Fundamental Performance

The main conclusion that we can derive from the previous simulations is that self-selected CDS are much less stable than neighbor-designated CDS. Either the self-selected relay set (i) shrinks too much and does not deliver the broadcast properly, or (ii) expands so much that the flooding of a single broadcast will involve much more bandwidth than with a neighbor-designated CDS.

In fact a more realistic radio transmission model will likely exhibit collision storms and a much reduced correct reception ratio. And even when the relay sets are similar in size, the performance may differ significantly.

In this section we give a closer look at the inherent mechanisms used by the self-selection and the neighbor-designation. We investigate their impact on performance. To this end we will assume that self-selected CDS operate as follows:

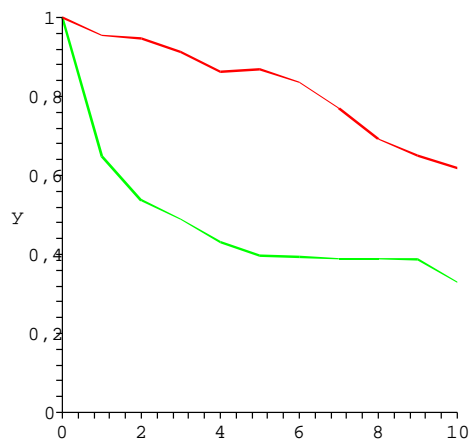


Figure 10: Outdated information: fraction of the network receiving the broadcast versus the outdated window T in a network with 100 nodes on a 4×4 square, for self-selected CDS (bottom, red), on-line pruning MPR (top, green).

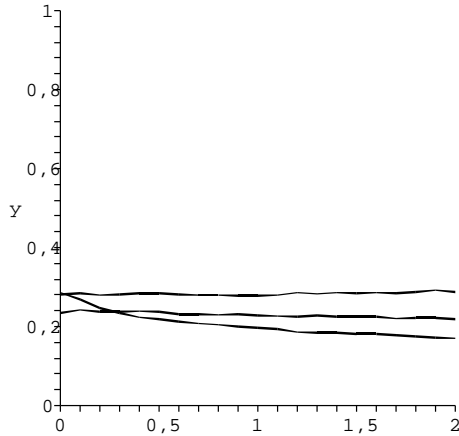


Figure 11: Outdated information: Relay set size versus the outdated window T in a network with 100 nodes on a 4×4 square, for self-selected CDS (bottom), on-line pruning MPR (middle), off-line pruning MPR (top)

- Each node selects itself with probability $\frac{\nu}{M}$

where M is the average neighborhood size of the nodes in the network. Therefore the density of relay per unit disk is ν . The counter part for neighbor-designated MPR selection is the following:

- Each node randomly selects ν neighbor nodes as its MPR

With on-line self-pruning MPR selection, the average number of relay nodes per unit disk is ν , while for the off-line self-pruning it is $\nu + 1$. Figures 16 and 17 show the performance of these random CDS versus the relay set density. It clearly shows that for the same relay set size, the MPR CDS outperforms the self-selected CDS. Furthermore, the self-selected CDS performance is much less stable than those of self-pruning MPR.

Of course, random CDS are not realistic, and may be considered as extreme cases where the neighbor information is completely random. However this gives a better insight, coupled

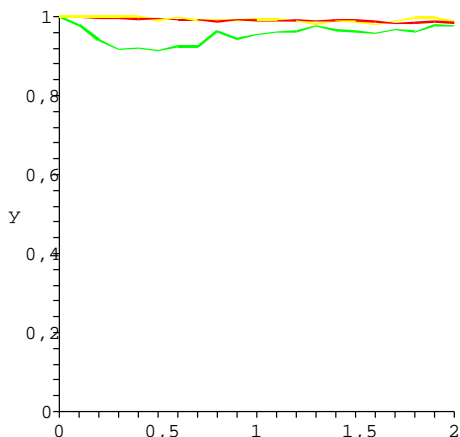


Figure 12: Uniformly delayed information: fraction of the network receiving the broadcast versus the outdated delay T in a network with 100 nodes on a 4×4 square, for self-selected CDS (bottom, green), on-line pruning MPR (top, red), off-line pruning MPR (top, yellow)

with the previous performance simulation results in a context of mobility or transmission errors.

7 Queueing in Stations

The fact that the routing is retransmitted by a limited subset of the nodes may lead to unfairness, if the subset does not vary frequently enough with time and sources. Figure 18 displays the packet delivery rate when LSA packets are flooded with an off-line pruning MPR flooding or with an on-line pruning MPR algorithm. On-line pruning algorithm leads to a permanent diversity in the subset of relay nodes: the relay set varies with the source of the broadcast source. If the medium access control introduces a certain randomness in the order of retransmissions, then the relay node subset also varies with time, even when the network is static. Consequently the burden of forwarding is more fairly shared with on-line

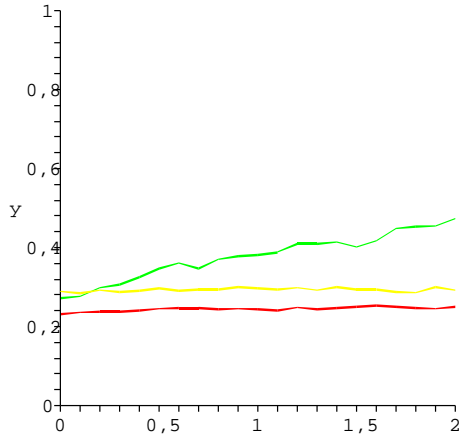


Figure 13: Uniformly delayed information: Relay set size versus the outdated delay T in a network with 100 nodes on a 4×4 square, for self-selected CDS (top, green), on-line pruning MPR (bottom, red), off-line pruning MPR (middle, yellow)

MPR. On the contrary, off-line pruning algorithms and self selected CDS don't vary their relay sets with different broadcast sources, and consequently the burden of forwarding is not equally shared between all the nodes in the network. In particular this may lead to packet loss because the queues in relay node are more frequently congested.

Figure 18 shows the indirect loss rate (*i.e.* the loss rate due to route failure because of missed LSA). In fact the LSA delivery rate should be much lower. The packet delivery rate varies with mobility, as mobility leads to frequent topology changes and triggers numerous extra LSA transmissions, increasing queue congestions.

8 Conclusion

Connected dominating sets are a critical feature for mobile wireless networking. If too many nodes retransmit topology update information, the network may simply collapse,

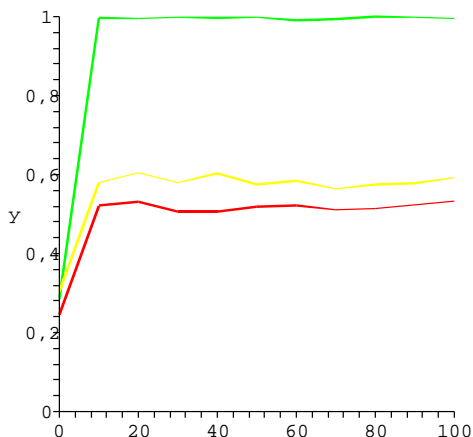


Figure 14: Uniformly delayed information: Relay set size versus the outdated delay T in a network with 100 nodes on a 4×4 square, for self-selected CDS (top, green), on-line pruning MPR (bottom, red), off-line pruning MPR (middle, yellow)

overloaded with control traffic on its own. Therefore a mechanism is needed to limit the set of broadcast forwarders in the network. Using a connected dominating set is such a solution. However there are various ways to build a connected dominating set. Among them, based on two-hop neighborhood information, there are two main approaches: (i) self-selected CDS and (ii) neighbor-designated CDS such as MPR in OLSR. We have shown that neighbor-designated connected dominating sets perform much better in general than self-selected connected dominating sets. In particular they are more resilient to transmission errors and neighbor information outdated. The unstabilities of self-selected CDS either lead to a too small CDS that does not offer enough coverage of the network, or to a much too large CDS that yields too many redundant retransmissions of broadcast packets.

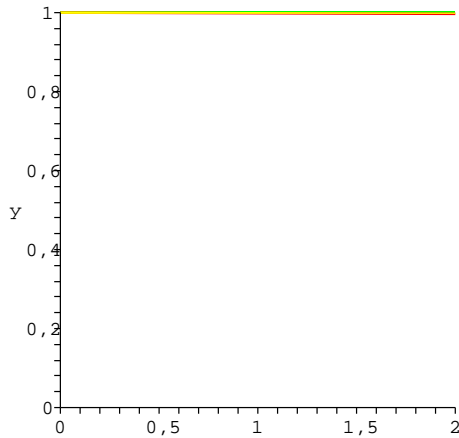


Figure 15: Uniformly delayed information fraction of the network receiving the broadcast versus the outdated delay T in a network with 100 nodes on a 4×4 square, for self-selected CDS (green), on-line pruning MPR (red), off-line pruning MPR (yellow)

References

- [1] J. Wu and H.Li, On calculating connected dominating set for efficient routing in ad hoc wireless networks, Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (1999) 7-14.
- [2] C.Adjih,P.Jacquet, L.Viennot, Computing Connected Dominating Sets with Multipoint Relays, INRIA RR-4597 (2002).
- [3] A. Qayyum, L. Viennot, and A. Laouiti. Multipoint relaying for flooding broadcast message in mobile wireless networks. In Proceedings of HICSS-35, Jan. 2002.
- [4] T. Clausen and P. Jacquet (Ed.), "Optimized Link State Routing Protocol (OLSR)", RFC 3626, October 2003.

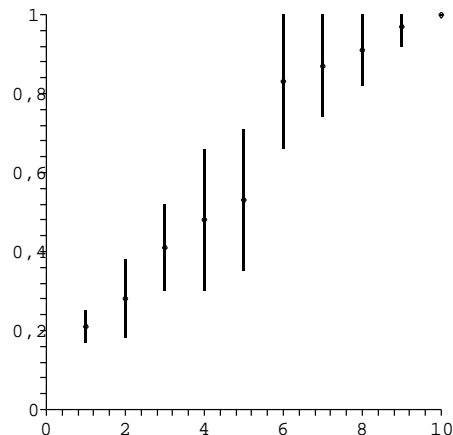


Figure 16: Random selection: Average correct reception ratio for the random self-selected CDS versus relay set density in a network with 100 nodes on a 4×4 square, with standard deviation.

- [5] J. Wu, F. Dai, "Performance Analysis of Broadcast Protocols in Ad Hoc Networks Based on Self Pruning," Proceedings of WCNC, 2004.
- [6] J. Harri, C. Bonnet, F. Filali, "OLSR and MPR, mutual dependence and performance," Proceedings of MedhocNet'05, 2005.
- [7] F. Dai and J. Wu "Double-Covered Broadcast (DCS): A Simple Reliable Broadcast Algorithm in MANETs," Proceedings of IEEE INFOCOM, 2004.
- [8] W. Lou and J. Wu, "Double-Covered Broadcast (DCS): A Simple Reliable Broadcast Algorithm in MANETs," Proceedings of IEEE INFOCOM, 2004.
- [9] J. Wu, F. Dai, "On Constructing k-Connected k-Dominating Set in Wireless Networks," Proceedings of IPDPS, 2005.
- [10] J. Wu, F. Dai, "Efficient Broadcasting with Guaranteed Coverage in Mobile Ad Hoc Networks," to appear in IEEE Transactions on Mobile Computing.

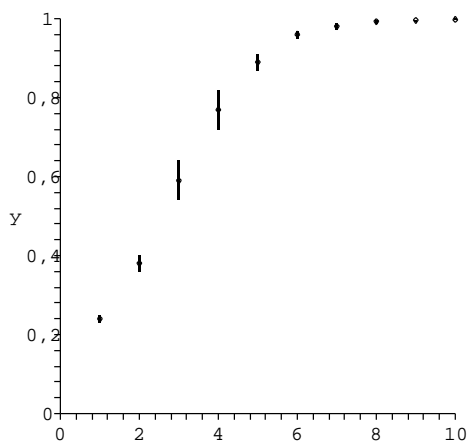


Figure 17: Random selection: Average correct reception ratio for the random self-pruning MPR versus relay set density in a network with 100 nodes on a 4×4 square, with standard deviation.

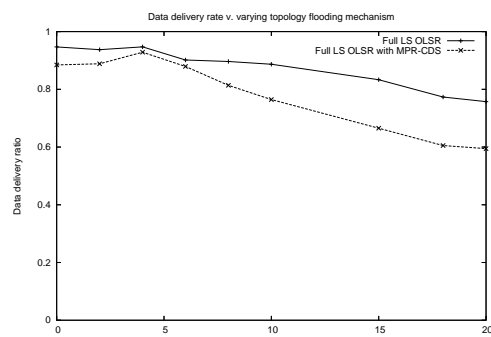


Figure 18: Random selection: Average data delivery rate versus mobility when LSA are flooded by on-line self-pruning MPR (MPRF), off-line MPR (MPR-CDS), ns-2 simulation of a network with 100 nodes on a 4×4 square.



Unité de recherche INRIA Rocquencourt
Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Unité de recherche INRIA Futurs : Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

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