INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Multipoint relaying: An efficient technique for flooding in mobile wireless networks

Amir Qayyum, Laurent Viennot, Anis Laouiti

No 3898

March 2000

THÈME 1

Rapport de recherche

ISSN 0249-6399 ISRN INRIA/RR--3898--FR+ENG

Multipoint relaying: An efficient technique for flooding in mobile wireless networks

Amir Qayyum, Laurent Viennot, Anis Laouiti *

Thème 1 — Réseaux et systèmes

Projet HIPERCOM

Rapport de recherche n ° 3898 — March 2000 — 16 pages

Abstract: In this paper we discuss the mechanism of multipoint relays (MPRs) to efficiently do the flooding of broadcast messages in the mobile wireless networks. Multipoint relaying is a technique to reduce the number of redundant re-transmissions while diffusing a broadcast message in the network. We discuss the principle and the functioning of MPRs, and propose a heuristic to select these MPRs in a mobile wireless environment. We also analyze the complexity of this heuristic and prove that the computation of a multipoint relay set with minimal size is NP-complete. Finally, we present some simulation results to show the efficiency of multipoint relays.

Key-words: multipoint relays, mobile wireless networks, flooding of broadcast messages

(Résumé : tsvp)

* Email: `Amir.Qayyum@inria.fr`, `Laurent.Viennot@inria.fr`, `Anis.Laouiti@inria.fr`

Relais multipoint: Une technique efficace pour la diffusion dans les réseaux mobiles sans fil

Résumé : Dans cet article nous discutons le fonctionnement des relais multipoints (MPR) pour diffuser efficacement des trafics de type broadcast dans les réseaux mobiles sans fil. Le relais multipoint est une technique pour réduire le nombre de re-transmissions redondantes dans le réseau lors de la diffusion d'un message broadcast. Nous discutons le principe et le fonctionnement des MPR, et proposons une heuristique pour choisir ces MPR dans un environnement mobile sans fil. Nous analysons également la complexité de cet heuristique et montrons que le calcul d'un ensemble de relais multipoint de taille minimale est NP-complet. En conclusion, nous présentons quelques résultats de simulation pour montrer l'efficacité des relais multipoints.

Mots-clé : relais multipoint, réseau mobile sans fil, diffusion de trafic broadcast

1 Introduction

The research relating to the mechanisms and the protocols used in the wired networks is becoming mature. As a result, these mechanisms and protocols are now classified according to their relative domains of application, based on the performance results obtained in those specific areas.

For the mobile wireless networks, the research is still in its earlier stage. The things are not yet standardized and there is less consensus about the applicability of different existing techniques and algorithms to these new type of networks. To obtain a satisfactory performance from these techniques or the algorithms, they must be adequate to this new and challenging mobile wireless environment.

1.1 Requirements of a mobile wireless environment

When we talk of “mobile” “wireless” networks, each of these two words put before us a list of requirements, and the daunting task is to fulfill them to their best. The **mobility** implies the limited lifetime of neighborhood or topology information received at any time, because of the movement of the nodes. This implies that the information be updated regularly, otherwise it becomes invalid. More frequently the information is updated, more the mobility of the nodes can be handled correctly and efficiently.

The **wireless** nature of the medium implies the limited bandwidth capacity available in a frequency band. It is further reduced because of the high bit error rate in the radio transmission. This makes it a scarce and hence a precious resource in the wireless world. Every attention is paid to consume it very wisely. Hence, while designing a protocol using wireless links, the main task is to reduce the unnecessary use of this bandwidth.

Therefore, the requirements of these two environments are completely opposite to each other. The mobility requires more traffic to be send in the network to keep the other nodes informed of the changes, and at the same time, the wireless medium does not have the capacity to be used abundantly for the unnecessary traffic. Hence, the compromise is to manage the mobility of the nodes while using minimum of the bandwidth resources.

1.2 Flooding of broadcast messages in the network

In most of the cases, the type of control traffic that is generated to manage the mobility of the nodes in a network is the information that a node declares about its relative movement, its new position, or its new neighborhood, etc. Some times, this information is useful only in the neighborhood of the node which is declaring the information. In these situations, we are not concerned about the information propagating in whole of the network to reach every node. But in many cases, not only the immediate neighbors of the declaring node, but the other, far away nodes also need to know the topological changes occurring anywhere in the network. In these situations, lot of message passing is required in the network to keep the information consistent and valid at each node, by regularly announcing the changes due to the mobility, or failure of links, etc.

These announcements about the changes are destined to each node of the network. But mostly, all the nodes of the network are not in the radio range of each other to communicate directly. So there must be a mechanism to reach the far away nodes to keep them informed of the latest changes. Here comes the concept of *intermediate nodes* which serve as relays to pass the messages between the source and the destination.

When a message is for a specific destination, the determination of the intermediate nodes is simple: all the nodes which form the path (if it exists) from the source up to the destination are the *intermediate nodes*. These nodes agree upon a mechanism to re-transmit the message, on their turn, so that the message is successfully transfered to the destination. Different routing protocols designate, in their own manner, these intermediate nodes for this unicast packet forwarding.

The problem arises when the packet is not destined to a specific node, rather it is a broadcast message for all the nodes in the network. Now the task of determining the intermediate nodes who will forward the packet is not very easy. How should the nodes behave so that the message is reached to every node in the network ? A simple solution is that each node re-transmit the message, when it receives it the first time. Fig 1 shows an example where a packet originated by node S is diffused up to 3-hops with 24 retransmissions. The packet is retransmitted by the intermediate nodes to be diffused in the network. We call this common technique as "pure flooding". It is simple, easy to implement, and gives a high probability that each node, which is not isolated from the network, will receive the broadcast message, but it consumes a large amount of bandwidth because of so many redundant retransmissions.

In certain conditions, and particularly in the "wireless" networks, the availability of the limited resources in terms of bandwidth capacity requires us to restrict the traffic as much as possible. If this constraint of wireless medium is not considered

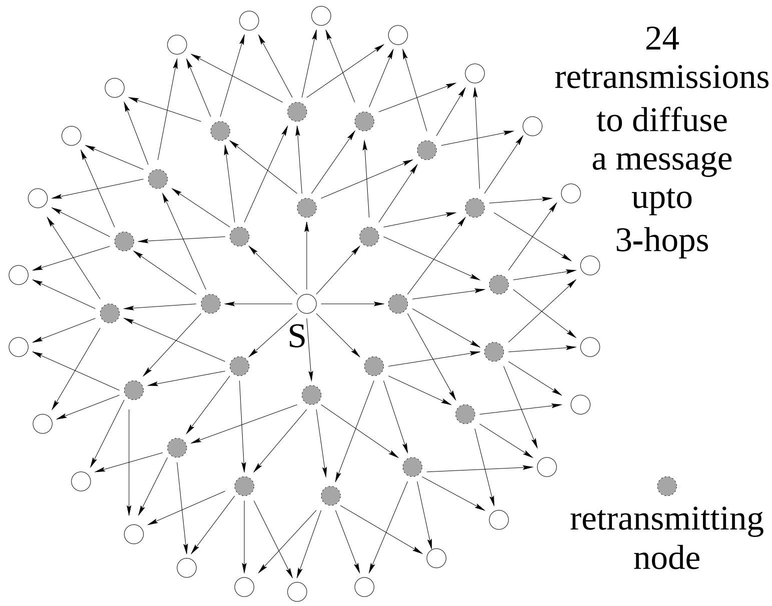

Figure 1: Diffusion of broadcast message using pure flooding

while designing an algorithm, the network may suffer from performance degradations due to high overloads or congestion, which may attain practically a halt of system, while the flooding of broadcast packets is launched in the network. On one hand, the flooding of broadcast messages is necessary, specially in the mobile environment to keep the mobile nodes remain in contact by regularly diffusing the updates. But on the other hand, it is not appreciate able either to affect the actual working of the system due to this additional control traffic.

Every protocol uses some king of flooding of the control messages, for its functioning [4], [6] so it is very advantageous to optimize its resource consumption. There are many techniques described in the literature to limit the flooding of broadcast traffic [7], [2]. Each technique has its own area of application and each has its own advantages and disadvantages. Here, we will discuss the mechanism of "multipoint relaying" as one of the possible solution.

2 Multipoint relaying

The concept of "multipoint relaying" is to reduce the number of **duplicate retransmissions** while forwarding a broadcast packet. This technique restrict the number of re-transmitters as much as possible by efficiently selecting a small subset of neighbors which covers (in terms of one-hop radio range) the same network region

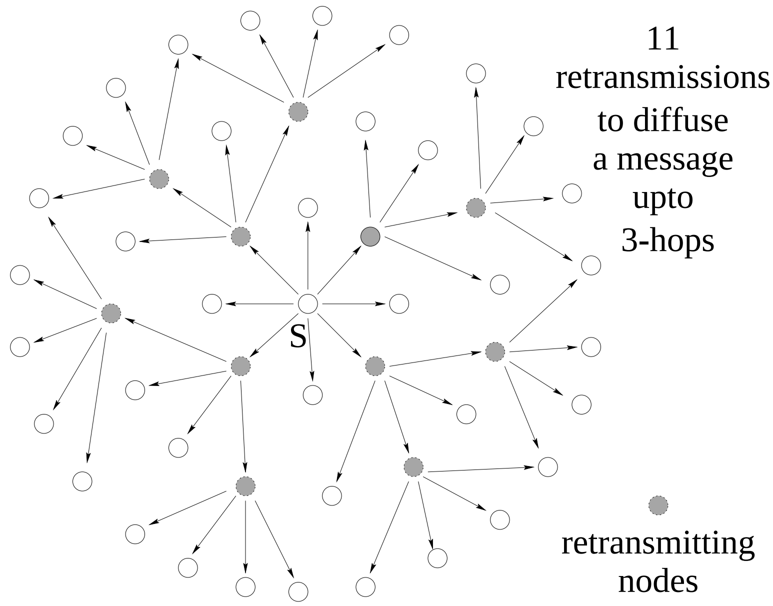

Figure 2: Diffusion of broadcast message using multipoint relays

which the complete set of neighbors does. This small subset of neighbors is called **multipoint relays** of a given network node. The scheme of multipoint relays (or MPRs) provides an adequate solution to reduce the flooding of broadcast messages in the network, while attaining the same goal of transferring the message to every node in the network with a high probability. Fig 2 shows an example where a broadcast message of node S is diffused in the network using the multipoint relays. In this case, it took only 11 retransmissions for a message to reach up to 3-hops.

Multipoint relaying technique works in a distributed way, keeping in view the mobile and dispersed nature of the network nodes. Each node calculates its own set of multipoint relays, which is completely independent of other nodes' selection of their MPRs. Each node reacts when its neighborhood nodes change and accordingly modifies its MPR set to cover its two-hop neighbors.

An important aspect for the utilization of the multipoint relays is the manner in which these multipoint relays are selected by each node. Obviously, the goal is to achieve the maximum performance by selecting an optimal set of these MPRs by each node. But this task is not a trivial one. If the mechanism of selecting the MPRs is too simple, it may not select efficiently the MPRs in the dynamic and complex situations, and the expected performance gain would not be achieved. On the other hand, if the algorithm of MPR selection is very long and complicated to provide a *near to optimal* MPR set, it may become difficult to implement it or it may generate its own control traffic (to gather information for its functioning) that becomes comparable to the saving in the flooding of messages. So, there must be a compromise in designing such an algorithm for the selection of multipoint relays: it should be easy to implement, and it should give near to optimal MPR set in the "majority" of cases.

The information needed to calculate the multipoint relays is the set of one-hop neighbors and the two-hop neighbors, i.e. the neighbors of the one-hop neighbors. To get the information about the one-hop neighbors, most protocols use some form of HELLO messages, that are sent locally by each node to declare its presence. In a mobile environment, these messages are sent periodically as a *keep alive* signals to refresh the information. To obtain the information of two-hop neighbors, one solution is that each node attach the list of its own neighbors, while sending its HELLO messages. In this way, each node can independently calculate its one-hop and two-hop neighbor set. Once a node has this information, it can select the minimal number of one-hop neighbors which *covers* all of its two-hop neighbors.

2.1 Heuristic for the selection of multipoint relays

We propose here one heuristic for the selection of multipoint relays. To select the multipoint relays for the node x, lets call the the set of one-hop neighbors of node x as $N(x)$, and the set of its two-hop neighbors as $N^2(x)$. Let the selected multipoint relay set of node x be $MPR(x)$.

1. Start with an empty multipoint relay set $MPR(x)$
2. First select those one-hop neighbor nodes in $N(x)$ as the multipoint relays which are the only neighbor of some node in $N^2(x)$, and add these one-hop neighbor nodes to the multipoint relay set $MPR(x)$
3. While there still exist some node in $N^2(x)$ which is not covered by the multipoint relay set $MPR(x)$:
   (a) For each node in $N(x)$ which is not in $MPR(x)$, compute the number of nodes that it covers among the uncovered nodes in the set $N^2(x)$
   (b) Add that node of $N(x)$ in $MPR(x)$ for which this number is maximum.

To analyze the above heuristic, first notice that the second step permits to select some one-hop neighbor nodes as MPRs which must be in the $MPR(x)$ set, otherwise the $MPR(x)$ will not cover all the two-hop neighbors. So these nodes will be selected as MPRs in the process, sooner or later. Therefore, if the second step is omitted, the multipoint relay set can still be calculated with success, i.e. it will cover all the two-hop neighbors. The presence of step 2 is for optimizing the MPR set. Those nodes which are necessary to cover the two-hop set $N^2(x)$ are all selected in the beginning, which helps to reduce the number of uncovered nodes of $N^2(x)$ to start with the normal recursive procedure of step 3.

3 Complexity analysis on the computation of multipoint relays

This section is devoted to the analysis of the computation of the multipoint relays. We will show that unfortunately, finding a multipoint relay set with minimal size is NP-hard. Nevertheless we will see that the above heuristic is within a $\log n$ factor from optimality. Let us first give a formal definition of the problem.

3.1 Formal definitions

If x is a node of the network, we denote by $N(x)$ the set of its one-hop neighbors. $N(x)$ is called the neighborhood of x. (Here we consider that $x \notin N(x)$.) Let $N^2(x)$ denote the two-hop neighbors of x.

If y is a one-hop neighbor of x, we also say that x covers y. Or we will simply say that y is a neighbor of x. Moreover, if S and T are sets of nodes, we say that S covers T iff every node in T is covered by some node in S. A set $S \subseteq N(x)$ is a multipoint relay set for x if S covers $N^2(x)$, or equivalently $\cup_{y \in N(x)} N(y) - N(x) \subseteq \cup_{y \in S} N(y)$. A multipoint relay set for a node x is optimal if its number of elements is minimal among all the multipoint relay set for x. We call this number the optimal multipoint relay number for x.

3.2 NP-completeness

We prove that the following problem is NP-complete:

Multipoint Relay: *Given a network (i.e. the set of one-hop neighbors for each node), a node x of the network and an integer k, is there a multipoint relay set for x of size less than k ?*

First of all, notice that this problem is easier than the problem of finding an optimal multipoint relay set. If an optimal set is known, simply computing its size and comparing it to k allows to answer the question. Let us now show that the Multipoint Relay Problem is NP-complete.

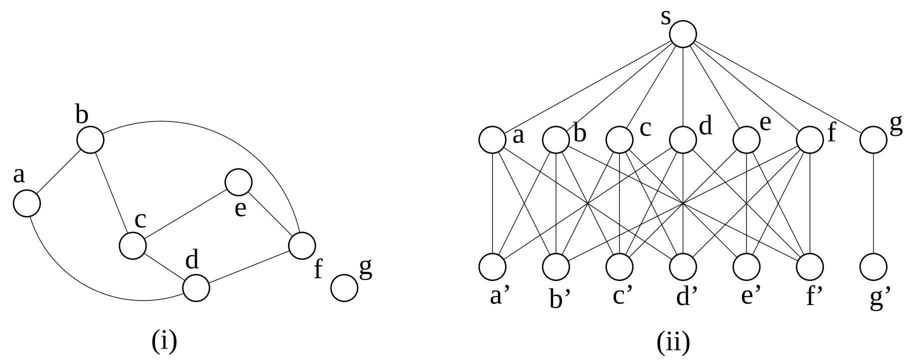

Figure 3: (i) A graph. (ii) The network obtained by the reduction. $\{b, c, g\}$ is a dominating set in (i) and a multipoint relay set for s in (ii).

It is obviously in NP since taking a random set in $N(x)$, one can easily check in polynomial time if it is a multipoint relay set and if its size is less than k. To prove that it is NP-complete, we prove that the following Dominating Set Problem which is known to be NP-complete [5] can be reduced to the Multipoint Relay Problem in polynomial time:

Dominating Set Problem: *Given a graph (i.e. a set of nodes and a set of neighbors for each node) and a number k, is there a dominating set of cardinality less than k ? Where a dominating set is a set S of nodes such that any node of the graph is either in S or in the neighborhood of some node in S.*

Let G be a graph with node set V and let $M(x)$ denote the neighborhood of any $x \in V$. We construct a reduction as follows. Let us make a copy of V and denote with a prime the copies: x' denotes the copy of x for any $x \in V$ and S' denotes the set of copies of the elements of any set $S \subseteq V$ (V' denotes the set of all the copies). Let s be an element not in V nor in V'. Consider a network where the nodes are $\{s\} \cup V \cup V'$ and where the neighborhood are the following (see Figure 3.2 for an example):

$$N(s) = V, \quad N(x) = \{x'\} \cup M(x)' \text{ for } x \in V, \quad N(x') = \{x\} \cup M(x) \text{ for } x \in V$$

Such a data structure can easily be computed in polynomial time. We claim that the answer to the Multipoint Relay Problem for the node s of the computed network with the integer k is valid for the Dominating Set Problem for the considered graph with the same integer k. It is sufficient to prove that any multipoint relay set S for the network is associated to a dominating set of the graph with same cardinality. S is a subset of $N(s) = V$. We show that S itself is a dominating set of the graph. Consider a node $x \in V$ and its copy x'. As S is a multipoint relay set, x' is the neighbor of some node $y \in S$. As $N(y) = \{y'\} \cup M(y)'$ by definition, we have either $x' = y'$ or $x' \in M(y)'$, or equivalently, $x = y$ or $x \in M(y)$. This means that x is in

S or is the neighbor of some node in S. S is thus a dominating set and the proof is achieved.

3.3 Analysis of the Proposed Heuristic

We prove that the heuristic proposed in section 2.1 computes a multipoint relay set of cardinality at most $\log n$ times the optimal multipoint relay number where n is the number of nodes in the network.

We give a proof directly inspired from [3] which is itself inspired from a general proof by Chvátal [1]. The first proof about an analogous heuristic was given in [8].

Let S_1 be the nodes selected in stage 2 of the above algorithm and let $x_1, \dots, x_k$ be the nodes selected in stage 3 (x_i is the ith added node). Let S^* be a solution with minimal cardinality. First notice that $S_1 \subseteq S^*$ since any node in S_1 is the only neighbor of some node in $N^2(s)$. We will show that $|S - S_1| \leq \log n |S^* - S_1|$ which implies that the computed solution is within a factor $\log n$ from the optimal.

Let $N^2{}_1$ be the set of nodes in $N^2(s)$ that are neighbors of some node in S_1. We set $N^{2\prime} = (s)N^2 - N^2{}_1$, $S' = S - S_1$, $S^{*\prime} = S^* - S_1$ and $N'(x) = N(x) \cap N^{2\prime}$ for each node $x \in N$. We associate a cost c_y to each node $y \in N^{2\prime}$. For each x_i chosen by the algorithm, a unit cost is equally divided among the nodes newly covered in N^2. More formally: if x_i is the first neighbor of y added in S by the algorithm, then we set:

$$c_y = \frac{1}{\left|N'(x_i) - \cup_{j=1}^{i-1} N'(x_j)\right|}$$

The costs are linked with the cardinality of the computed solution in the following way:

$$|S'| = \sum_{y \in N^{2\prime}} c_y$$

We are going to show that for any node z in $S^{*\prime}$, we have:

$$\sum_{y \in N'(z)} c_y \leq \log |N'(z)| \tag{1}$$

Notice first that this implies immediately the result. Any node $y \in N^{2\prime}$ is the neighbor of some $x \in S^{*\prime}$ (remember that no node in S_1 is a neighbor of y by definition). We can thus deduce:

$$|S'| = \sum_{y \in N^{2\prime}} c_y \leq \sum_{z \in S^{*\prime}} \sum_{y \in N'(z)} c_y \leq \sum_{z \in S^{*\prime}} \log |N'(z)| \leq |S^{*\prime}| \log n$$

We still have to prove Inequation 1 to conclude. Let z be a node in $S^{*\prime}$ and let

$$u_i = \left|N'(z) - \cup_{j=1}^{i} N'(x_j)\right|, \text{ for each } 0 \leq i \leq k \ (u_0 = |N'(z)|)$$

be the number of neighbors of z in $N^{2'}$ which are still not covered after the choice of $x_1, \ldots, x_i$. Let l be the first index such that $u_l = 0$. When x_i is chosen, $u_{i-1} - u_i$ neighbors of z are then covered. We can thus deduce:

$$\sum_{y \in N'(z)} c_y = \sum_{i=1}^{l} (u_{i-1} - u_i) \frac{1}{\left|N'(x_i) - \cup_{j=1}^{i-1} N'(x_j)\right|}$$

We then notice that the choice of x_i by the algorithm implies:

$$\left|N'(x_i) - \cup_{j=1}^{i-1} N'(x_j)\right| \geq \left|N'(z) - \cup_{j=1}^{i-1} N'(x_j)\right| = u_{i-1}$$

This implies:

$$\sum_{y \in N'(z)} c_y \leq \sum_{i=1}^{l} (u_{i-1} - u_i) \frac{1}{u_{i-1}} \leq \int_{u_l}^{u_0} \frac{dt}{t} \leq \log u_0 \leq \log |N'(z)| \leq \log n$$

The upper bound on the approximation factor follows. Notice that we can get a sharper bound on the approximation factor: it is bounded by $\log \Delta$ where Δ is the maximum number of two-hop nodes a one-hop node may cover. When a vertex covers at most 40 nodes, the approximation factor of the heuristic is bellow 3.7. When a vertex covers at most 100 nodes, the approximation factor of the heuristic is bellow 4.7.

Some simulations have been made to show how multipoint relays computed with this heuristic may practically be useful.

4 Simulations

The objective of the simulations was to compare two types of algorithms for the diffusion of packets in the radio networks; one is the pure flooding technique, and the second is the diffusion of packets using multipoint relays. The simulations aim to evaluate the behavior of these algorithms in the conditions of high error rates, either due to the radio transmission problems or because of the dynamic environment with rapidly changing topologies. We were interested in seeing the impact of these errors on the network with these two techniques. Moreover we wanted to see until which limit, the algorithm of multipoint relays is able to ensure the diffusion and can guarantee good results.

4.1 Simulation model

Our study relates to large networks, in terms of number of nodes. We considered the dense networks, so the nodes have a significant number of links with their neighbors. In order to make sure the existence of a link from a node to all other nodes in the network, we considered the connected networks only, *i.e.* without any partitions or the isolated nodes. The graph of the network was composed of a grid of nodes and their links. All the nodes were placed on the grid, to form a square network region. A radio range radius was defined, and all the nodes which were inside this radius were considered as the direct, one-hop neighbors. For all the simulations, we considered a graph of 1024 nodes placed on a 32x32 grid.

The simulations consisted of varying the probability of error of reception from 0 to 100%, and diffusing a message of a node in all of the network. This procedure was repeated for each node of the network to calculate an average of these values, for each value of the probability of error.

In our simulations, we adopted certain assumptions to appropriately define the area of our study on the problem of impact of error of reception on the diffusion of packets. These assumptions are as follows:

- The messages are broadcast messages which do not require an explicit acknowledgement to confirm the reception. Hence there was no retransmission when error of reception occur.
- There are no asymmetric links. Each link between a pair of nodes is a perfect symmetric link (bi-directional).
- The only traffic which exist in the network is that of the diffusion of the broadcast packet.
- Each node retransmits a packet (if it has to retransmit according to the protocol) only once.
- There is a synchronization among the transmissions. Channel is time-slotted and each transmission takes one slot.
- Each time a node transmits a packet, its one-hop neighbors receive this packet with probability P, P being a percentage which lies between 0 and 100.

For a node to transmit, it was necessary that none of its neighbors up to 2-hops are transmitting. We call this as blocking up to 2-hops. It was used to eliminate the problem of interference when a node receives two radio transmissions at the same time by two of its neighbors, which are not neighbors themselves.

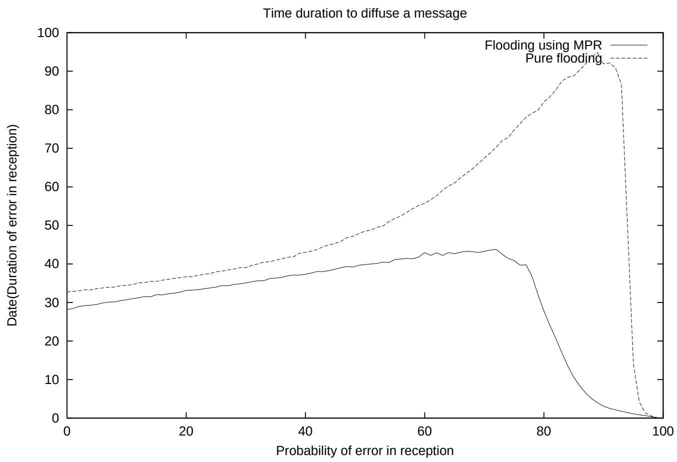

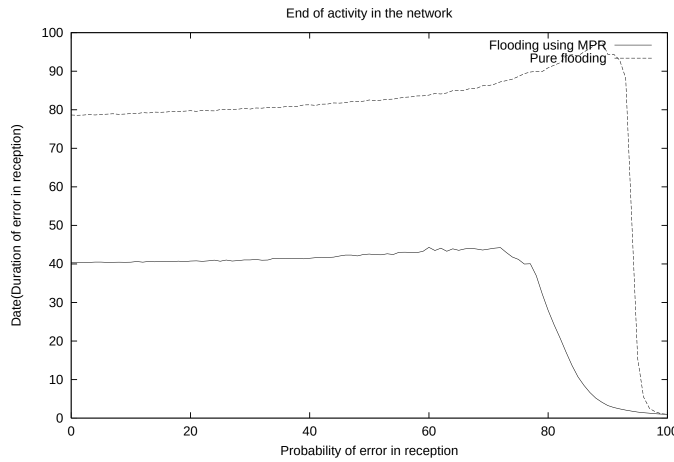

Figure 4: Completion of diffusion, and end of activity in the network

4.2 Simulation results

Here we discuss some of the simulation results that we have obtained.

First of all, in figure 4, the first graph shows that when a message was diffused in the network, how much time it took (in terms of clock ticks) so that all the nodes of the network get that message. It can be seen that pure flooding took more time as compared to multipoint relay technique, to diffuse the message in the network. In the second graph, we compare the time at which the activity ended in the network that was started to diffuse a message. As expected, the pure flooding took almost double the time as compared to multipoint relay technique. This behavior can be explained as the result of packet retransmission by each and every node of the network, even when it is not required. This can be proved by comparing the two graphs side by side, and we can easily observe that when the packet is successfully diffuse in the network, multipoint relaying technique took quite less time to stop further retransmissions, but pure flooding continue to retransmit, as each node must retransmit the packet, once, on its turn.

In figure 5, the first graph shows that how many of the nodes has retransmitted the message, on the average. For multipoint relaying, this figure was quite low as only selected nodes had retransmitted the packet, still achieving the comparable performance (as shown in rest of the graphs). Unless the error rate was too high where all the nodes were not able to receive the packet, in case of pure flooding, obviously it was all the 1024 nodes which retransmitted. As a result of this, the second graph shows that in pure flooding, on average, the nodes have received too

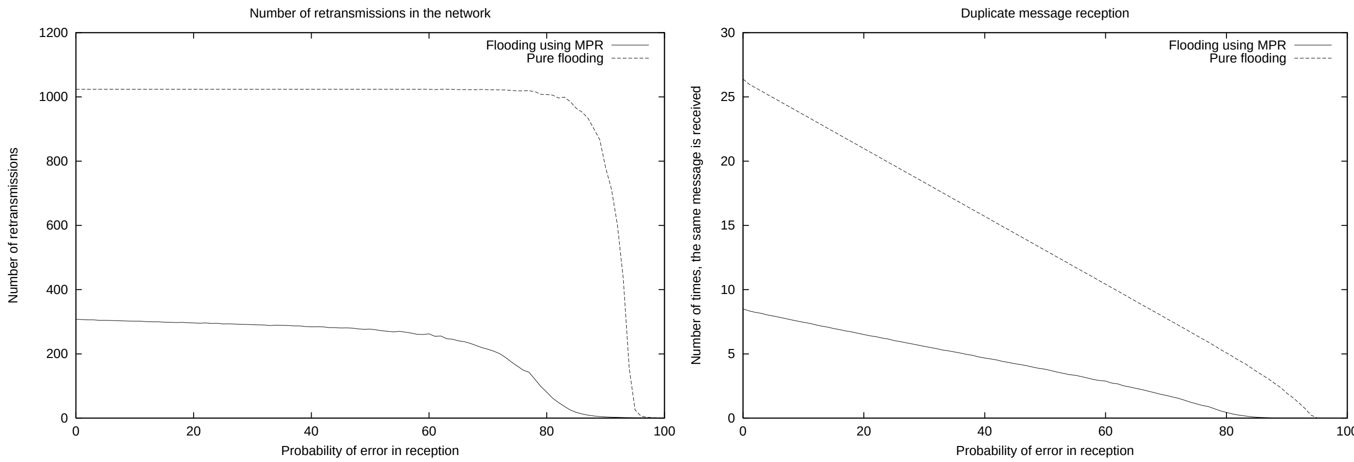

Figure 5: Number of retransmitting nodes, and the duplicate receptions

many duplicate copies of the same message as compared to the case of multipoint relaying.

In optimizing the flooding and reducing the traffic by multipoint relaying, there is a small price to pay, and that is the robustness of the protocol in varying conditions of error rates. In pure flooding, as each node retransmit without exception, there are more chances that the message reach at the maximum of nodes, as compared to multipoint relaying, where only a selected number of nodes propagate the message. Figure 6 shows this fact, by comparing two protocols. We observe that when the error probability is higher than 20 or 25%, the multipoint relaying technique starts loosing the packets, and some nodes do not receive the message because of these errors.

5 Conclusions

We have seen the performance of the two techniques, and according to the simulation results,the multipoint relaying has shown superiority over pure flooding scheme. The results of the simulation show that although the classic technique of pure flooding to diffuse a message in the network is more reliable and robust, it consumes a large amount of bandwidth as its cost. On the other hand, multipoint relaying gives equally good results, with much less control traffic, when the errors of reception remain less than 20%. In general, its a quite realistic assumption to consider these errors as less than 10% in a network. So we can say that in the range of error rate

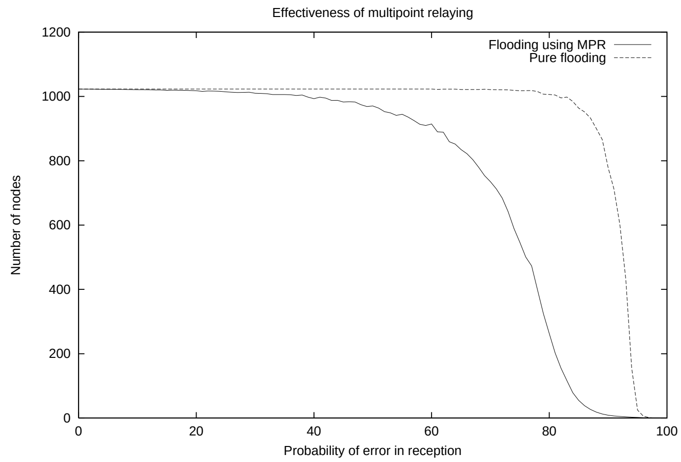

Figure 6: Effectiveness of diffusion in different conditions of error rate

which is most common, the multipoint relaying gives us quite satisfactory results, with a tremendous gain in performance due to quite less traffic.

References

[1] V. Chvátal. A greedy heuristic for the set-covering problem. *Mathematics of Operation Research*, 4(3):233–235, 1979.

[2] ETSI STC-RES10 Committee. Radio equipment and systems: High performance radio local area network type 1, functional specifications. Technical report, ETSI, 1995.

[3] T. Cormen, C. Leiserson, and R. Rivest. *Introduction to Algorithms*, chapter 37.3. MIT Press, 1990.

[4] M. S. Corson, S. Papademetriou, P. Papadopoulos, V. Park, and A. Qayyum. An internet MANET encapsulation protocol (IMEP) specification. Internet Draft, draft-ietf-manet-imep-spec-01.txt, August 21 1998, Work in progress.

[5] M. R. Garey and D. S. Johnson. *Computers and intractability, a guide to the theory of NP-completeness*. W.H. Freeman, 1979.

[6] P. Jacquet, P. Muhlethaler, and A. Qayyum. Optimized link state routing (OLSR) protocol. Internet Draft, draft-ietf-manet-olsr-01.txt, February 7 2000, Work in progress.

[7] Philippe Jacquet, Pascale Minet, Paul Muhlethaler, and Nicolas Rivierre. Increasing reliability in cable-free radio LANs: Low level forwarding in HIPERLAN. *Wireless Personal Communications*, 1997.

[8] S. Sahni and T. Gonzales. P-complete approximation problems. *Journal of the ACM*, 23:555–565, 1976.

Unité de recherche INRIA Lorraine, Technopôle de Nancy-Brabois, Campus scientifique,
615 rue du Jardin Botanique, BP 101, 54600 VILLERS LÈS NANCY
Unité de recherche INRIA Rennes, Irisa, Campus universitaire de Beaulieu, 35042 RENNES Cedex
Unité de recherche INRIA Rhône-Alpes, 655, avenue de l'Europe, 38330 MONTBONNOT ST MARTIN
Unité de recherche INRIA Rocquencourt, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex
Unité de recherche INRIA Sophia-Antipolis, 2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex

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