



HAL
open science

Performance issues with IEEE 802.11 in ad hoc networking

Claude Chaudet, Dominique Dhoutaut, Isabelle Guérin-Lassous

► **To cite this version:**

Claude Chaudet, Dominique Dhoutaut, Isabelle Guérin-Lassous. Performance issues with IEEE 802.11 in ad hoc networking. [Research Report] RR-5304, INRIA. 2004, pp.17. inria-00070696

HAL Id: inria-00070696

<https://inria.hal.science/inria-00070696>

Submitted on 19 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

*Performance issues with IEEE 802.11 in ad hoc
networking*

Claude Chaudet — Dominique Dhoutaut — Isabelle Guérin Lassous

No 5303

Septembre 2004

THÈME 1

A large blue rectangular area containing the text 'Rapport de recherche' in a white serif font. To the left of the text is a large, light gray stylized 'R' logo. A horizontal gray brushstroke is positioned below the text.

*Rapport
de recherche*



Performance issues with IEEE 802.11 in ad hoc networking

Claude Chaudet , Dominique Dhoutaut , Isabelle Guérin Lassous

Thème 1 — Réseaux et systèmes
Projet Ares

Rapport de recherche no 5303 — Septembre 2004 — 15 pages

Abstract: Although 802.11 is initially designed for infrastructure based networks, the Distributed Coordination Function (DCF) provided in 802.11 allows mobiles to access the radio medium without the need of a base station. Thus, studies in wireless multi-hops networks, also called ad hoc networks often rely on the use of the IEEE 802.11 standard for the physical and MAC layers. But, since a couple of years, the use of 802.11 in ad hoc networks has been called into question. Different scenarios show serious performance issues. The performances offered by 802.11 are often low and directly impact the performances of higher layers protocols.

In this article, we provide a summarize of the different performance issues that have been extracted so far. We identify the main causes of performance losses and we classify the different basic encountered situations according to these causes. We then discuss the known impacts on complex ad hoc topologies and on higher layers protocols. We hope that this classification will help to design more appropriate MAC protocols for ad hoc networks.

Key-words: IEEE 802.11b, ad hoc networks, performance evaluation, simulation

Problèmes de performances de la norme IEEE 802.11 dans un contexte ad hoc

Résumé : La norme IEEE 802.11 est le standard actuel en matière de réseaux sans fil. C'est sur ce protocole, créé pour des réseaux avec station de base, que la majorité de la recherche en matière de réseaux *ad hoc* s'articule. Cependant, depuis quelques temps, la pertinence de l'utilisation de ce protocole d'accès au médium dans un tel contexte est remise en question. En effet, dans certaines configurations, de sérieux problèmes de performances peuvent survenir, impactant le bon fonctionnement des protocoles de niveau supérieur.

Cet article présente une compilation des problèmes de performances connus dans de tels réseaux basés sur la fonction de coordination distribuée (DCF) du standard IEEE 802.11. Nous identifions et classons les différentes configurations problématiques puis examinons l'impact de ce type de problèmes sur des configurations *ad hoc* plus complexes.

Mots-clés : IEEE 802.11b, évaluation de performance, Réseaux ad hoc, simulation

1 Introduction

Mobile ad hoc networks are wireless and infrastructureless networks. They are formed by the gathering of wireless capable terminals and do not require any base station to operate. Communications between out of range peers is assured by a dedicated routing protocol. The terminals are mobile and these networks are based on the self-organizing principle. The randomness of the topology coupled to mobility make these systems uneasy to analyze.

IEEE 802.11 [?], also known as Wi-Fi, has become the standard for wireless networking. This protocol defines both physical and medium access layers. If this standard has initially been designed for base station managed networks, it includes a distributed medium access protocol, called the Distributed Coordination Function (DCF) that requires no base station. With the DCF, ad hoc networks can rely on 802.11 as soon as a routing protocol is provided.

Most of the designed protocols for ad hoc networks assume that 802.11 is used on low layers. But some works have put under light a certain number of performance issues arising with 802.11 used in an ad hoc context that lead to the following question: Is IEEE 802.11 really well suited for ad hoc networks? All the works on the subject accuse the DCF mode. But each study evaluates its specific scenario and there is no work that summarizes all encountered performance issues. Moreover, some studies evaluate 802.11 via specific applications and it is difficult to isolate the performances of 802.11 from the ones of higher layers.

This article presents an overview of all known performance problems that can appear in these networks with the 802.11 DCF. We identify the main causes of performance losses and classify the known configurations according to these causes. Section 2 presents in details the DCF mode. Section 3 compiles a classified list of the elementary issues and presents some simulation results. Section 4 discusses the influence of these problems on complex ad hoc topologies and on higher layers protocols.

2 IEEE 802.11 DCF

The IEEE 802.11 DCF is part of the Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) family. To prevent collisions, emitters have to wait for the channel to become free before trying to send a frame. When a frame is ready to be emitted, if the medium is free when the frame is dequeued, it is emitted after a small fixed period of time called *DIFS* during which the medium shall stay idle. If the medium is or becomes busy during this interval, the sender draws a random number called *backoff* in an interval called *Contention Window*. When the medium becomes idle again, the mobile waits for *DIFS* before starting to decrement its backoff, slot by slot. If the medium becomes busy during this interval, the process is stopped and will be resumed later after a new *DIFS* and with the remaining number of backoff slots when the medium becomes idle once again. As soon as the backoff reaches 0, the frame is emitted.

This random mechanism decreases the collision probability but does not prevent every collision. When a collision happens, the contention window size of involved emitters is

doubled and the same frame is re-emitted with the same process as described above. If another collision happens, the contention window size is doubled again, until it reaches its maximum value defined by the standard. After a fixed number of retransmissions, the frame is dropped and the contention window size is reset. A successful transmission also resets the contention window size.

As collision detection is not possible due to the half-duplex characteristic of the wireless interfaces, each unicast frame has to be acknowledged. When a receiver successfully receives a frame, it waits for a *SIFS* time and then emits the acknowledgment. *SIFS* is much shorter than *DIFS* in order to give priority to acknowledgments over data frames. The lack of acknowledgment is treated by the emitter as a collision and the frame is retransmitted.

Consequently, acknowledgment frames are as important as data frames. Therefore, to prevent accidental jamming of acknowledgments, emitters sensing a signal on the medium that is identified to be a 802.11 frame but unable to decode the frame contents will defer their transmission of a fixed time called *EIFS* instead of *DIFS*, when the medium becomes idle. This timing is much larger than *DIFS* and sufficient to protect acknowledgment transmission. The reception of an error-free frame during the *EIFS* wait cancels this long wait mechanism (the station may transmit after a *DIFS* time following this error-free frame).

To prevent hidden nodes situation in which two independent emitters simultaneously send a frame to the same receiver, an optional Request To Send (RTS) - Clear To Send (CTS) exchange can be used. Before transmitting a frame, the emitter asks the receiver if the medium is free in its vicinity by emitting a RTS frame. If no interfering transmission is present, the receiver answers by a CTS frame and the transmission can begin. These frames are overheard by neighbors of both emitter and receiver nodes and cause a medium reservation for the duration of the transmission. Each of the nodes hearing one of these messages acts as if the medium was busy for the whole transmission duration. This mechanism is called *Virtual Carrier Sense*.

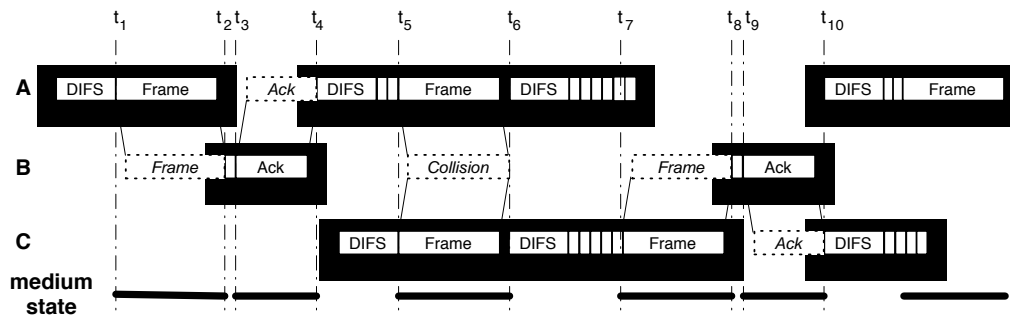


Figure 1: IEEE 802.11 protocol operation

Figure 1 summarizes a part of these operations. This figure represents the MAC operations in a network of three nodes in mutual communication range. At the beginning, node *A* has a frame to send. As the medium is free, it can emit its frame after a *DIFS* period, which is followed by the acknowledgment from *B*. *A* then wishes to emit another frame. If the new frame has been enqueued while the previous exchange was taking place, *A* has to draw a backoff. In this example, a backoff of 2 slots is drawn. In parallel, *C* also wishes to send a frame to *B*. As the medium was sensed free, it emits as soon as *DIFS* is elapsed. Unfortunately, *A* and *C* choose the exact same moment to start their emission and a collision happens at *B*. As no acknowledgment is sent back, both emitters re-emit their frame, drawing a backoff in a doubled contention window. This time, *C* wins the contention and its emission stops the backoff decrementation of *A* two slots before it finishes. After the Data-acknowledgment exchange is finished, *A* restarts the backoff process with its two remaining slots.

3 Performance issues

Performance study of this standard has been the subject of many contributions. Simulation based studies are the most frequent and can help to understand the behavior of the MAC layer even though physical layer is often over-simplified. Analytical modeling is also an active field but the designed models always study fully connected (one hop) networks or at least fully synchronized networks. Finally, experimentations represent the most reliable way for performance analysis but experimental conditions are difficult to control and results are often difficult to analyze. Most of the works that present performance issues with 802.11 in ad hoc context give simulation results. The presented performance issues can be classified in three main categories:

- The configurations that result in an achievable throughput decrease for all involved mobiles, and thus in an overall throughput decrease. That means that a part of the network capacity is not used and wasted.
- The configurations that lead to long term fairness issues in which some flows suffer from starvation while other flows capture the whole channel bandwidth.
- The configurations that result in short term fairness issues where the flows are emitted by burst.

Usually, performance issues do not appear until the medium capacity is overloaded. But, with greedy applications such as file or video transfer, situations where capacity is overloaded can frequently arise in ad hoc networks since the radio medium is scarce and shared. Considering the 802.11 DCF, the causes that prevent an emitter for accessing the radio medium in normal conditions, *i.e.* that reduce its throughput can be classified in three categories:

- The emitter can not transmit because the medium is busy for a long time.

- Its random wait phase (given by the backoff) is long due to its contention window increase after repeated collisions.
- Its fixed wait phase is longer than in usual conditions since it uses *EIFS* in place of *DIFS* due to the presence of distant transmissions.

All the known performance problems come from one of these three causes. The network topology represented in Figure 2 regroups all types of basic configurations that lead to performance issues. On the figure, plain lines represent the ability for two mobiles to communicate directly at a high rate. Dashed lines represent a communicating link between two mobiles but unable to use the full capacity due to signal quality (in this example the used rate is 2Mb/s). Dotted lines represent the carrier sensing links between mobiles. In other words, when one mobile emits, the other senses the medium busy. To simplify the study, we will consider that mobiles in carrier sense range trigger the *EIFS* mechanism when the other mobile emits a frame. Actually these two zones are not exactly the same but are similar. Depending on the active flows in this network, different types of performance issues can appear. In the next paragraphs, each situation will be depicted and analyzed according to the identified causes. In each situation, the flows try to saturate the medium capacity. We simulate some of the situations. We use the NS2 simulator with modified parameters of the 802.11 DCF in order to match the rate and the parameters of the 802.11b wireless cards. In order to isolate the performances of 802.11 from the ones of higher layers, we use AODV as the routing protocol with large timeouts in order to get a static routing without route breakages. Moreover the emitters send UDP packets since UDP does not add any complexity in analysis compared to TCP (as we will see in Section 4).

3.1 Busy medium

Carrier sensing mechanism prevents mobiles from transmitting when another transmission is active. In [?], the authors point out that the presence of slow terminals in a one-hop network slows down every other terminal: during the transmission of a slow terminal the medium is busy for a longer period than during the transmission of a fast terminal. As 802.11 tends to provide per-packet fairness in one-hop networks, each emitter sends one frame in turn statistically. Therefore, the throughput of every node is decreased by the slowest emitters. This situation arises on Figure 2 if flows (A, C) and (B, D) are active simultaneously. As soon as multiple rates are present in an ad hoc networks, a loss of performance can appear due to the existence of slow transmissions. Note that with actual wireless cards, the control frames (RTS, CTS, etc.) and the broadcast frames are transmitted at a slower rate than unicast frames.

In [?], a scenario involving three pairs of emitters-receivers is depicted. Consider on Figure 2 the presence of three flows (C, D) , (E, F) and (G, H) . Emitters C and G are totally independent and evolve asynchronously. When emitter E wishes to transmit a frame, it has to wait for both C and G to leave the medium free simultaneously. Since C and G are independent, their silence periods almost never coincide and E sees the medium busy more

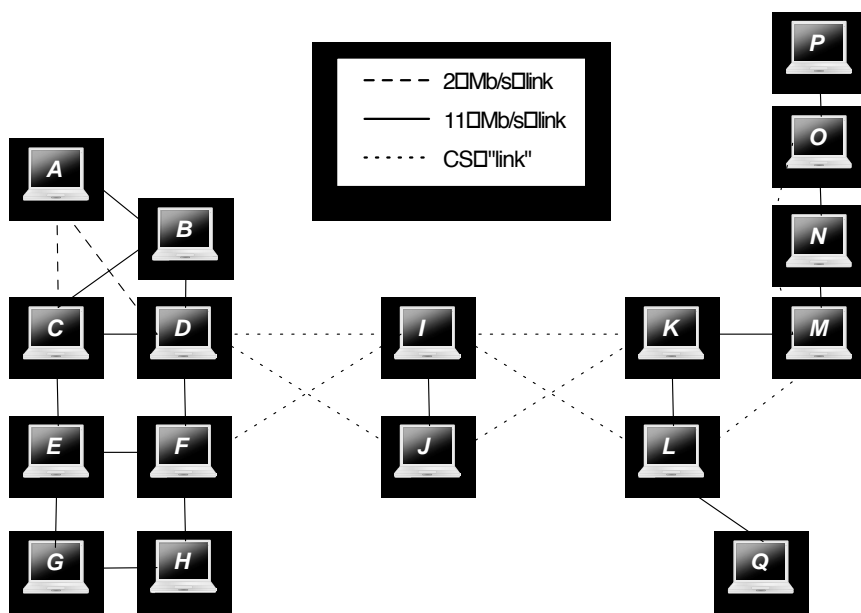


Figure 2: An ad hoc network presenting all performance issue configurations

often than the two others. This basic unbalance in the medium access (one against two) results in a severe fairness issue as the flow between E and F can expect no more than 15 % of the medium capacity whereas the two others can use 75 % of the capacity, as shown in [?], by both analytical modeling and simulation. This share of the radio medium is given in Figure 3 that shows the throughput achieved by each of the emitters according to the packets size. We also see that the throughput of the exterior transmitters increases with the packets size whereas the throughput of the central emitter remains constant. That means that modifying packets size while keeping the same size on the three emitters improves the overall performance of the network at the cost of an unfairness increasing.




Figure 3: Achieved throughput for the three pairs configuration

If the first issue may be solved by providing a fair time share of the medium rather than a fair access of the medium, the second issue is more tedious. The exterior emitters are not aware of their mutual emission and they can not see that they prevent another emitter to transmit.

3.2 Long random wait

Another cause of performance loss is the increase of backoff time due to repeated collisions. When an emitter does not receive an expected acknowledgment or by a CTS, it doubles its contention window size and will statistically have to wait longer than other nodes before transmitting. This can be the cause of many unbalanced situations.

In [?], the authors point out that even though RTS-CTS exchange can solve hidden nodes situations and results in a fair medium access in the long term, the behavior of the medium access is unfair in the short term. On Figure 2, let's consider that flows (K, L) and (Q, L) are active, as the RTS-CTS exchange. Frames from K and Q will collide at L until both

emitters' backoffs are large enough to allow the successful transmission of a RTS. When one emitter succeeds, the CTS that follows blocks the other emitter from further decrementing its backoff which is statistically large as contention window sizes have increased just before. On the other hand, the successful emitter resets its contention window and draws a small backoff at the next round. Thus, it is likely to access the medium again and so on until collisions happen again. This results in a bursty behavior for both emitters. The simulations we have carried out shows that the average number of packets that are consecutively sent by the emitter that has won the contention is nine regardless of the packets size.

In [?], the authors describe another scenario that presents long time fairness issues due to backoff increase. This case is more severe than the previous issue as one flow is almost starved. This situation arises, on Figure 2, when both (C, D) and (F, H) flows are active. In this situation, node C always senses the medium free and therefore alternates backoff decrementation periods and emissions. Transmission from node F to node H always succeeds if it manages to start as collisions never happen at the receiver. F sees the medium free unless virtual carrier sense blocks it or D transmits an acknowledgment to C . Therefore, if RTS-CTS mechanism is not active, the chance for C to successfully transmit a frame depends on the frame size, as it should fit in a backoff interval of F that never increases its contention window. With RTS-CTS mechanism, the chance for C to be able to reserve the medium access is increased but nevertheless, if the CTS is emitted during the time F sends a frame, F will not understand the CTS and will soon emit another frame. This results in a great unfairness as shown on Figure 4. It gives the share of the medium in this situation according to the packets size and with and without RTS-CTS. Without RTS-CTS, the unfairness is high and gets worse when the packets size increases: from 800 bytes, C almost never experiments successful transmissions. It is consistent with the previous explanations. With RTS-CTS, this share remains constant and corresponds to 10 % of the used bandwidth for the successful transmissions of C and to 90 % for F .

In [?], the authors propose to modify the algorithm that manages the contention window size of the DCF in order to not favour in the following the terminal that has just won the contention. Thus, they reduce the short term unfairness. In [?], the authors propose a fair MAC protocol based on the DCF: by listening the packets sent on the medium, each terminal is able to determine if it accesses to the medium more often than its neighbors or vice versa. With this knowledge, the terminals can adapt their contention window size in order to limit or to increase their throughput. The difficulty is to define the thresholds and the modification algorithms of the contention windows size. In these two solutions, the major issue is to achieve a better fairness without too large losses of overall performance.

3.3 Long fixed wait

In order to avoid jamming acknowledgments from distant receivers to distant emitters, nodes are supposed to wait for *EIFS* before decrementing any backoff slot when they sense a signal that can be identified as being a frame but that they cannot fully decode. On Figure 2, if flows (M, N) and (L, Q) are active simultaneously, M does not understand frames from L , and therefore cannot benefit from virtual carrier sensing. If M starts to transfer as soon as





Figure 4: Share of the medium in the two pairs configuration

the medium is free, for example drawing a very small backoff, the resulting signal over noise ratio at L will prevent the acknowledgment from Q from being correctly received. *EIFS* leaves enough time for the acknowledgment to be transmitted successfully in these cases.

In [?], the authors, present an issue, called *Large EIFS problem*, leading to an unfair behavior. If we consider simultaneously flows (M, N) and (N, O) on Figure 2, acknowledgments from O trigger EIFS wait at node M . Therefore, the situation becomes unbalanced between the two contending emitters as one has to wait for *DIFS* while the other needs to wait *EIFS* before decrementing any backoff slot. This kind of situation is very likely to arise in ad hoc networks in which routing is performed resulting in multi-hop communications. Figure 5 gives the achieved throughput for M and N according to the packets size. We see that the throughput of each emitter increases with the packets size but the ratio between the two emitters remains the same: N sends around 7 times more packets than M .

Let's now get back to the three pairs scenario as depicted in Section 3.1, but this time emitters are not in communication range of each other but in carrier sense range, triggering EIFS between each other. For example, this situation arises when flows (D, F) , (I, J) and (K, L) are present simultaneously on Figure 2. The situation is worsened by the use of EIFS as the central pair not only has to wait for both exterior pairs to be silent simultaneously, but it also has to wait for them to be silent simultaneously for a time longer than EIFS in order to be able to decrement a single backoff slot. This leads, as pointed out in [?] and [?] to a medium share of about 5% versus 95%.

If we only consider two of the three previous pairs, for example (D, F) and (I, J) on Figure 2, long term fairness is preserved but we can observe a short time unfairness. The simulations we have carried out shows that the average number of packets that are consec-

Figure 5: Throughput in the large *EIFS* problem

utively sent by the emitter that has won the contention is around four whatever the packets size may be.

EIFS seems therefore to represent an overall performance loss or an inequity increase in many situations. The solution proposed in [?] can not be applied since the knowledge comes from packets that can be decoded which is not the case in the situations of this section. The authors of [?] modify the DCF to reduce the unfairness in the three pairs scenario: the contention window size of each transmitter evolves according to the number of packets that have been sent in the past. If a terminal has sent many packets, it increases its contention window size in order to free the medium for other potential transmitters. This increase is periodically reset in order to not penalize the emitters too much and to not waste too much bandwidth in needless waitings. The authors of [?] propose to measure the length of the frames detected but not decoded and to base the deferring according to these measures. Such a solution requires to use different lengths for the different frame types. As for the previous solutions, the major issue is to get both a good fairness between flows and a good use of the capacity of the network.

4 Influence on complex topologies and on higher layers protocols

The scenarios presented here are quite elementary. One could argue that a combination of several scenarios could resolve partially the problems mentioned here. That may be true, but the analysis is complex and the experiments have not yet come to maturity. Some studies of complex configurations have been published.

Let's now consider the configuration where flows (M, N) and (P, O) are active on Figure 2. In this situation presented in [?] and called *Small EIFS problem*, a significant performance decrease in terms of overall bandwidth is coupled to a short term fairness issue. On one hand, collisions happen since *EIFS* triggered on the *CTS* is not large enough to protect the associated data frame sent by an independent emitter, what results in an increase in the contention windows size. On the other hand, the *EIFS* triggered by an emitter on the acknowledgment of the receiver of the other pair slows down this emitter. It then wastes its time in unnecessary wait. The simulations we have carried out show that the overall performance loss increases with the packets size and ranges from 9% to 25% of the capacity. Moreover, the average number of packets that are consecutively sent by the emitter that has won the contention is almost 28.

In [?] and [?], the authors study the behavior of a chain of mobiles (for example in Figure 2, imagine a flow from M to P and routed by N and O). This configuration happens as soon as two mobiles not in communication range wish to communicate. Combination of multiple problems lead to a global performance problem. They show that achievable throughput is lower than expected: the throughput stabilizes around 250 – 300 kb/s with a 6-hops chain and with a network capacity of 2 Mb/s compared to the expected 400 kb/s . This fact may be explained by the multiple collisions that occur on the second and third mobiles on the chain (N and O on the previous example) and the triggering of *EIFS* due to distant emissions. Thus, very few packets reach the fourth terminal of the chain. But when these packets are transmitted by this mobile, they are likely to reach the destination since few packets are sent on this part of the chain and the medium is not overloaded.

In Figure 6, we compare by simulation the throughput on a 6-hops chain with and without RTS-CTS exchange. It is interesting to note that the throughput decreases when RTS-CTS are activated. In that case, the spatial reuse possibility is reduced: without RTS-CTS, parallel communications between the first and second terminals and the fourth and fifth terminals are possible, whereas with RTS-CTS, if the fourth terminal is already sending packets to the fifth mobile, the second terminal can not answer to the RTS of the source since it is blocked by the transmission of the fourth mobile. The spatial reuse possibility is the reduced and it leads to a notable decrease in throughput.

In [?], the authors show that the achieve throughput on chain configurations in real ad hoc networks is very low and unstable: as soon as the chain has four hops, the throughput drops to 150 – 200 kb/s , which is lower than the one obtained by simulation. The same explanations as for the simulation results coupled to the high instability of the radio links may be given.

In [?], the authors show that the same inefficiencies of 802.11 apply on regular topologies like mesh and on random topologies with random traffic.

All these MAC performance issues can have negative impact on higher layer behaviors. The results on the chain configurations shows that some terminals on the chain wait for a long time before successfully transmitting a packet, what implies multiple route breakages if the timeouts of the routing protocols are not appropriate. A substantial performance loss may thus come from repeated recomputations of the routes even if the network is static.




Figure 6: Throughput on a six-hops chain

Other problems appear for routing protocols. In [?], the authors point out that broadcasted frames that uses a different modulation as data frames, travel farther than data frames. As ad hoc routing protocols often use broadcasted frames for route discovery, there are some situations in which these protocols may fail. For example on Figure 2, based on broadcasted information, A may consider D as a neighbor on one hand, but D can not received the data frames sent at 11 Mb/s by A on the other hand. That means that routing tables based on broadcasted data are not necessarily exploitable on higher rates.

The performance issues of 802.11 have also an impact on the TCP protocols as mentioned in different studies. The problems are the following:

- The instability of the TCP connections: since the MAC links are unstable due the presented issues of 802.11, the TCP flows are often stopped and have to restart with the slow-start phase each time.
- The inequality between some TCP flows: the TCP mechanism increases the fairness issues encountered with 802.11 in some configurations, what results in a strong unfairness between the TCP flows.

Lastly, fairness issues often lead to a bursty behavior. This provokes an increase in the delays, and a high jitter that are unsuited to multimedia or real-time applications or even interactive applications.

5 Conclusion

In this paper, we summarize the performance issues that arise with 802.11 in ad hoc networks, and especially with the Distributed Coordination Function of 802.11. After having identified the main causes of performance loss, we classify the main known elementary topologies that give performance issues. The problems result in an overall throughput decrease or in unfairness situations in the long and short terms. Some simulation results are presented for some issues. Then, we discuss the impact of these problems on complex ad hoc topologies and on higher layers protocols.

This summary shows that the DCF of 802.11 is not well adapted to ad hoc networks and that the performances offered by 802.11 to higher layers protocols are poor, even on simple topologies and on usual topologies (like the chain communication). More efforts and researches are required to design a MAC layer adapted to ad hoc networks. One possible way of research is to consider MAC protocols families different from the CSMA/CA family. Some solutions investigate the use of TDMA or CDMA based protocols. But due to the dynamicity and the lack of synchronization, their implementations are not straightforward and much work remains to achieve a usable and efficient solution. Another possible way of research is to modify the DCF in order to obtain better performances. Some solutions have already been proposed, but they are specific to some configurations and can not solve the whole range of performance issues encountered. Work still remains to design a global solution that proves to be efficient in every situation. We hope that our classification helps to that end.

Lastly, due to the complexity of these networks, an isolated MAC layer may not be the right choice and information from lower and higher layers may help the MAC layer to efficiently adapt to the encountered problems. That is probably the most promising way of research.



Unité de recherche INRIA Rhône-Alpes
655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (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 Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (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