

Experiments of some performance issues with IEEE 802.11b in ad hoc networks

Claude Chaudet, Dominique Dhoutaut, Isabelle Guérin-Lassous

► **To cite this version:**

Claude Chaudet, Dominique Dhoutaut, Isabelle Guérin-Lassous. Experiments of some performance issues with IEEE 802.11b in ad hoc networks. [Research Report] RR-5303, INRIA. 2004, pp.17. inria-00070697

HAL Id: inria-00070697

<https://hal.inria.fr/inria-00070697>

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.

*Experiments of some performance issues with IEEE
802.11b in ad hoc networks*

Claude Chaudet — Dominique Dhoutaut — Isabelle Guérin Lassous

N° 5304

Septembre 2004

THÈME 1



*Rapport
de recherche*

Experiments of some performance issues with IEEE 802.11b in ad hoc networks

Claude Chaudet , Dominique Dhoutaut , Isabelle Guérin Lassous

Thème 1 — Réseaux et systèmes
Projet Ares

Rapport de recherche n° 5304 — Septembre 2004 — 17 pages

Abstract: Most of the studies in ad hoc networks assume that 802.11 is the used protocol on low layers. For a couple of years, some studies have extracted different performance issues with 802.11 in ad hoc networks. These problems result in a great unfairness in the long and the short term and in an overall performance loss. Most of these studies are based on simulation results and these performance issues have almost never been tackled from the angle of real experiments.

In this paper, we experiment different ad hoc configurations that present some performance problems. We focus on the fairness issue. We show that some problems appear even if the effects are sometimes softened, whereas some others are nonexistent.

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

Expérimentations de problèmes de performances de la norme IEEE 802.11b pour les réseaux ad hoc

Résumé : La majorité des études effectuées sur les réseaux *ad hoc* présupposent que la norme IEEE 802.11b est le protocole d'accès au médium utilisé. Depuis quelques années, certaines contributions ont toutefois mis en lumière certains problèmes de performances lorsque 802.11 est utilisée dans un contexte *ad hoc*. Ces problèmes conduisent à un comportement inéquitable des émetteurs, à long terme comme à court terme et à une perte globale de performance. La plupart de ces études sont basées sur des résultats de simulations et n'ont jamais été vérifiées en environnement réel.

Cet article présente certaines topologies présentant des problèmes de performance et plus spécifiquement d'équité d'accès au médium. Certains des problèmes soulevés par la communauté sont vérifiables par expérimentation bien que parfois atténués alors que d'autres sont inexistantes.

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

1 Introduction

Ad hoc networks have become more and more popular and some works in this field have come to maturity. The scale of researches has been extended to other themes than routing protocols design, like quality of service, security, location or services discovery. Most of these researches often assume that there exist efficient protocols on the lower layers than the one where the problem is addressed.

Actual ad hoc networks often rely on wireless cards that are commercially available. But these cards mainly follow the IEEE 802.11 standard [11], and especially the 802.11b extension (at least in France where the 802.11a version is forbidden). Recently, different studies have shown that 802.11 was not adapted to ad hoc networks, and they mainly have accused the MAC protocol of 802.11, known as the Distributed Coordination Function (DCF) (note that the DCF principle is the same whatever the version of 802.11 may be). These studies have extracted different performance issues like a decreasing of the overall performance and a great unfairness between the different flows.

Most of these performance issues have been proved by simulation results, and sometimes verified by theoretical models. But, as far as we know, no study has approached these performance problems from the angle of real experiments. However, simulators need to make some compromises between efficiency and reality and their simplifications may lead to results that are far from the ones obtained on real networks. Some of these differences have already been noticed in some articles that investigate the performances of ad hoc networks via experiments.

In the light of these observations, we think that it is useful to know in what extent the performance issues proved by simulations exist in real ad hoc networks: do they still exist in reality and what are their quantitative features compared with the simulation results? This article presents the experiments of some configurations that show performance issues with 802.11. As the available wireless cards are mostly 802.11b, we use this extension of 802.11 in the experiments. All the experimental results are compared with the simulation results carried out with the simulator NS2¹. We start this article by very briefly describing the 802.11 DCF mechanism in Section 2. In Section 3 the related works concerning performance issues with 802.11 in ad hoc networks on one hand and the experimental evaluations of ad hoc networks on the other hand are presented. Then, in Section 4, we analyze the experimental results obtained on some topologies that present performance issues.

2 IEEE 802.11 DCF

To make the reading of this paper easier, we briefly describe the DCF mechanism. For a complete description see the standard [11] or one of the numerous books on the subject, like [8] for instance.

The DCF is a Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) protocol: emitters have to wait for the channel to become free before trying to send a frame.

¹www.isi.edu/nsnam/ns/

When a frame is ready to be emitted, if the medium is free, it is emitted after a small 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 is free. 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, each implied emitter re-transmits its frame with a new backoff drawn in a new contention window of larger size. This size is modified by the Binary Exponential Backoff algorithm: the contention window size is doubled each time a collision happens. 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. The lack of acknowledgment is treated by the emitter as a collision. 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 wait for a fixed time called *EIFS* instead of *DIFS*, when the medium becomes idle. *EIFS* is much larger than *DIFS* and sufficient to protect acknowledgment.

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 it is the case, the receiver answers by a CTS frame and the transmission can begin. These frames are received by neighbors of both emitter and receiver nodes and cause a medium reservation for the duration of the transmission, *i.e.* they consider that the medium is busy for the whole transmission duration.

3 Related works

Performance study of this standard has been the subject of many contributions. The first studies dealt with the evaluation of one-hop networks where each terminal is in communication range with each other. The reader can refer, for instance, to [3, 17, 19, 4] for such evaluations.

In [2], the authors show the short term unfairness that arises in hidden terminals configuration with the RTS-CTS exchange. This configuration is theoretically studied in [13]. In [1], the authors describe a scenario composed of two pairs of communicating mobiles that presents long time fairness issues due to a contention window size increase for some mobiles. In [9], the authors point out that the presence of slow terminals in a one-hop network slows down every other terminal that sends data at higher rates. In [6], the authors put a high

long term unfairness to light with the most simple unbalance configuration where one emitter competes with two other emitters whereas these two terminals only compete with one. In that case, the fairness issue does not come from the contention window size increase but from the lack of free medium for the emitter that competes with the two others. In [14], the authors present another unfair scenario caused this time by the triggering of *EIFS* due to interfering transmissions. In [14] also, they show that the triggering of *EIFS* on control frames (like CTS and acknowledgment) and not on data frames strongly decreases the overall throughput of the network and also induces a short time unfairness.

All these situations are basic topologies. The works of [12] and [7] are focused on more complex topologies like regular topologies (chains, lattices) and random topologies. They show that the obtained throughputs are smaller than expected. These bad performances come from a combination of multiple problems that appear in basic situations.

These performance issues have also some impacts on higher layers protocols. Concerning routing, the low throughput and the short time unfairness may lead to repeated route breakages if the timeouts are not appropriate. Some studies (see [10, 20] for example) have proved that TCP protocols are also affected by the 802.11 performance issues: the TCP connections are very unstable and a strong unfairness can arise between TCP flows.

Most of these studies are based on simulation results and very few use real experiments (except for [2], [9] and [7]). The first experimental studies evaluate the performances of specific ad hoc routing protocols, like in [16] and [18]. APE [15] and Forwarding [5] are completely dedicated to experiments, since they provide a software for the deployment of scenarios on real ad hoc networks and the evaluation of the results. APE main concern is the evaluation of routing protocols whereas the aim of Forwarding is to test the MAC protocols themselves.

4 Experimental results

In this article, we approach the performance problems from the angle of real experiments. As shown in some studies, the experiments may sometimes give different behaviors from the ones observed in simulators. For instance, in [7], the throughput achieved on a chain of mobiles (situation that always happens as soon as two mobiles not in communication range wish to communicate) is smaller than the one achieved on NS2. Therefore, it is legitimate to wonder if the performance issues obtained by simulations are still present in real ad hoc networks and in what extent.

All the performance issues appear when the medium is overloaded, *i.e.* when the network capacity is saturated. They can be classified into three categories:

- An overall performance degradation: a part of the overall bandwidth is not used but the implied mobiles can not increase their throughput by using this free bandwidth. They spend time in useless waits.
- A long time unfairness: some flows capture the medium whereas others very scarcely succeed in transmitting packets.

- A short time unfairness: the flows are fairly transmitted in the long term, but the packets are emitted by burst. Each time an emitter accesses the medium, it sends several consecutive packets before the other competitive terminals succeed in transmitting packets.

In this paper, we focus our evaluation on fairness issue. Before describing the configurations we have experimented and the associated results, we briefly describe the experimental framework we have used.

4.1 Experimental framework

Since the goal is to evaluate the performances of the DCF, we have used the *Forwarding* software for experiments, rather than APE for instance. Forwarding is devoted to the experimental evaluation of MAC protocols and thus allows to isolate the performance of the MAC protocols for higher layers protocols. It uses a static routing according to predefined rules and CBR flows. It provides functions to deploy and monitor scenarios. For a more detailed description of Forwarding, the reader can refer to [5].

We also compare our experimental results with simulation results obtained with NS2. In order to be consistent with the experiments, we also use a static routing in NS2: we set high values in the timeouts of AODV to prevent from route breakages and then route repairs, and we put the initial phase of route sets up aside in order to not take it into account in the performance results. The used rates are 2 *Mb/s* and 11 *Mb/s* both in experimentations and simulations. The 2 *Mb/s* rate allows to test some scenarios with broadcasted packets that give very useful information (like the communication area and the carrier sensing area for instance) that can not be obtained with unicast packets. To obtain simulations at 11 *Mb/s*, we have modified the DCF implemented in NS2 in order to match the rate and the parameters of 802.11b.

4.2 Long term unfairness

Different scenarios present long term fairness issues. In this section, we investigate three configurations: a two pairs scenario presented in [1] where the fairness issue comes from hidden terminals, a three pairs scenario described in [6] where a terminal almost never senses the medium free and a scenario called *Large EIFS problem* in [14] where an emitter has to wait longer than another due the *EIFS* triggering. If these three configurations lead to the same fairness issue, the causes of this problem are completely different.

4.2.1 Another impact of the hidden nodes configuration

The topology of this scenario is given on Figure 1. *A* sends packets to *B* and *C* sends packets to *D*. *B* and *C* are in communication range, but *A* is independent of *C* and *D* and *B* of *D*. *A* and *C* are then hidden and although *C* does not send packets to *B*, *B* receives packets from both *A* and *C*. In this situation, *A* always senses the medium free since it

does not see C emitting and alternates between a wait phase and an emission phase. Let's start by considering that there is no RTS-CTS. C can emit if the medium is free, *i.e.* B is not sending an acknowledgment. As soon as C emits, the transmission will be successful since no collision happens at D and B is blocked by the transmission of C . While C is emitting, A is transmitting too. The only chance for A to successfully transmit a packet is that its frame size fits into a backoff interval of C . This is unlikely to happen, all the more if the packets size is large. Then, many collisions happen at B , what results in a contention window size increase for A whereas C keeps its initial contention window size. If the RTS-CTS are activated, the chance for A to reserve the medium is increased, but nevertheless if the CTS is emitted by B while C is sending a frame, C will not understand the CTS and will soon emit another frame, that will result in collision at B .

In this experiment, we have put the mobiles outside around a building in order to have an obstacle that hides the two emitters from one another and the second receiver from the mobiles of the first pair. The used rate is 11 Mb/s without and with a RTS-CTS exchange. Figure 1 compares experimental and simulation results according to different packets sizes. Between the seconds 100 to 200, packets of 250 bytes are sent without RTS-CTS. Between 300 to 400, packets of 500 bytes are sent without RTS-CTS as well. Between 500 to 600, RTS-CTS is still not used with packets of 1000 bytes. The two following experiments use the RTS-CTS: between 700 and 800, it corresponds to packets of 500 bytes and between 900 and 1000 to packets of 1000 bytes. The y-axis gives the throughput in bits per seconds. We first see that the fairness issue arises both in simulation and in experiments: the second pair uses most of the bandwidth whereas the first one achieves a low rate. Nevertheless, some differences between simulation and experimentation can be noticed. The unfairness is worsen with the simulations, especially without the use of RTS-CTS. If we more precisely analyze the logs of the experiment, we see that B can sometimes receive some packets of A even if C is transmitting, what results in a higher throughput for the first pair. In the case of RTS-CTS, this difference is no longer so marked. Indeed, the RTS-CTS exchange implies the medium reservation, what prevents successful transmissions in the first pair while the emitter of the other pair is transmitting, as we have observed without RTS-CTS. We can also note that, without RTS-CTS, the increase in packets size (between 250 and 500 bytes) does not result in an increase in the fairness issue. It can also be explained by the possible parallel successful transmissions between the two pairs: the number of such emissions is almost the same regardless of the packets size, what results in a higher throughput for the two emitters. Finally, note that the throughputs are much more unstable in the experiments than in simulations.

4.2.2 A simple unbalance in the medium access

This situation is described in Figure 2. It is the most simple unbalance in the medium access since one mobile has to compete with two other terminals and these two mobiles compete with only one mobile. This topology consists of three pairs of mobiles where the two exterior emitters are completely independent and evolve asynchronously and where the central emitter is in the carrier sense range of the two exterior transmitters and vice versa.

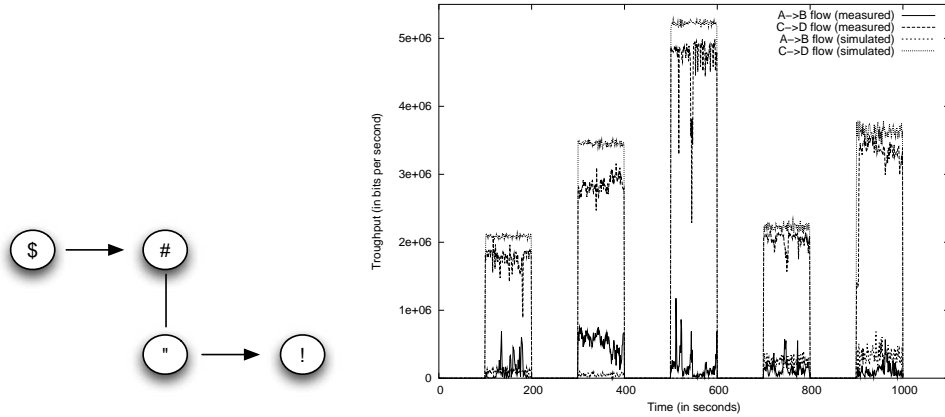


Figure 1: Another impact of the hidden nodes configuration: the scenario and the simulation and experimental results

To emit a packet, the central has to wait for both exterior pairs to be silent simultaneously, what scarcely happens since the exterior emitters are independent. Moreover this situation is worsened by the carrier sense mechanism since this silent period has to be larger than *EIFS* to allow the central emitter to decrement at least one slot of backoff.

This experiment has been carried in our laboratory, the mobiles of the experiment being the only ones active. The rate is 2 Mb/s without RTS-CTS. On Figure 2, we compare the experimental and the simulation results for different packet sizes. Between the seconds 100 to 200 it corresponds to a packets size of 700 bytes, between 300 and 400 to 800 bytes, between 500 and 600 to 1000 bytes and between 700 and 800 to 1200 bytes. The y-axis gives the throughput in bits per seconds. We see that the general scheme is the same for simulation and experimentation: the exterior pair uses most of the bandwidth whereas the rate of the central pair is low (note that for lisibility only the result for one exterior pair is given, but the rate of the other exterior pair is the same). We can also notice, that the fairness issue decreases in the experiments since the central pair gets a better throughput at the expense of the exterior pairs. If we more deeply study the logs, spatial reuse sometimes appears between the central emitter and the exterior emitters, *i.e.*, the central emitter can sometimes emit while the exterior terminals are transmitting. Anyway, it scarcely happens and it can not explained this difference between simulation and experimentation on its own. One possible explanation is that the *EIFS* is not triggered and that only the *DIFS* is used even if the emitter detects a signal that it can not decode. Removing the *EIFS* and using *DIFS* in place reduces the fixed wait, especially for the central emitter that often waits and thus increases its chance to decrement its backoff. Therefore, the unfairness is decreased as we observe with the experimental results.

It is also interesting to note that if using large packets on the different emitters increases the overall throughput, it increases the fairness issue since only the exterior pairs see their throughput increase.




Figure 2: A simple unbalance in the medium access: the scenario and the simulation and experimental results

4.2.3 The Large *EIFS* problem

In this simple configuration, given in Figure 3, there are two concurrent flows: *A* sends packets to *B* and *B* sends packets to *C*. Moreover *A* and *C* are in carrier sense range, represented by the dotted line in the figure. That means that when *C* sends an acknowledgment to *B*, *A* will trigger an *EIFS* for its next communication, what slows it down in comparison with *B* that always waits a smaller fixed time (*i.e.* *DIFS*).

The experiments have been carried out in our laboratory. The used rate is 11 Mb/s without RTS-CTS. Figure 3 compares the simulation and the experimental results for two packets sizes: between the seconds 100 and 200, packets of 500 bytes are sent and between the seconds 300 and 400, packets of 1000 bytes are transmitted. The y-axis gives the throughput in bits per seconds. Unlike the two previous experiments, there is a strong difference between simulations and experimentations: the two flows share the bandwidth and get the same rate in experiments whereas the second flow has a throughput 7 times higher than the first one in simulations. We think that these results corroborate the results obtained in the simple unbalance topology: the *EIFS* are not triggered in the used wireless cards since this mechanism is the only cause of unfairness in the Large *EIFS* problem evaluated in this section. Therefore, *A* waits for a *DIFS* like *B* what results in a fair share of the medium.

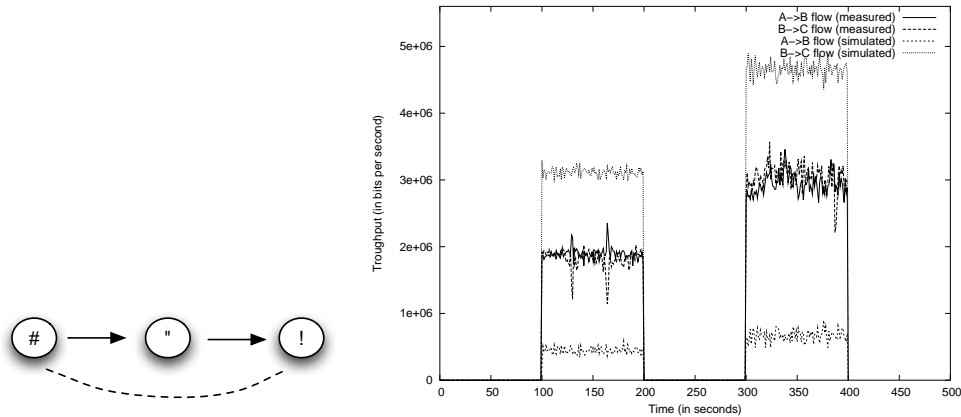


Figure 3: The Large *EIFS* problem: the scenario and the simulation and experimental results

4.3 Short term unfairness

Some other scenarios are fair in the long term, but present a short term fairness issue. The first scenario that presents such an unfairness is the hidden terminals configuration. In [2], the authors point out this short term unfairness and argue that the binary exponential backoff algorithm is not appropriate. Recently, the authors of [13] have theoretically studied this scenario. This fairness issue comes from the contention window size increase. We have evaluated this configuration via experimentations. In this section, we also describe a simple scenario consisting of two pairs of mobiles where this short term unfairness has a different cause: a large fixed wait (with *EIFS*). Once again, we have some scenarios with the same problems but induced by different causes.

4.3.1 The hidden terminals configuration

In this famous configuration, two terminals that can not detect each other try to send packets to a same receiver (see Figure 4). If the RTS-CTS mechanism is not activated, their emissions result in collisions at the receiver and thus in an increase of their contention window sizes. A successful transmission may happen in the frame size of one emitter fits into the backoff interval of the other one. When there are the RTS-CTS, collisions happen until one RTS is successfully transmitted. Then the medium is reserved and the winner resets its contention window size whereas the other mobile keeps its large contention window size. Therefore, the winner is likely to succeed in transmitting its next packet, and so on until a new collision happens. This phenomenon leads in a short term fairness issue.

We present different results for this situation we have deeply experimented. We have tested it in our laboratory. The used rate is 2 Mb/s. Figure 4 give the simulation results for different packets sizes. The y-axis gives the throughput in bits per second.




Figure 4: The hidden terminals: the scenario and the simulation results

Figure 5 gives the experimental results with broadcasted packets according to the packets size: the size ranges from 100 bytes to 1600 bytes and is increased by 100 every period of 50 seconds. The y-axis is the throughput in bits per second. In each period of 50 seconds, one emitter transmits during a couple of seconds, then the two emitters transmit at the same time, and finally the first emitter stops and the second emitter keeps sending packets. These isolated measures allow to get the maximal rate with the associated packet size. This experiment is interesting because it gives an idea of what happens at a receiver when two hidden mobiles broadcast packets (what may happen with control information in routing protocols for instance). The broadcasted packets are not acknowledge, therefore the emitters never increase their contention window size. We expect then many collisions on the receiver that can never correctly receive a packet as soon as the frame size does not fit into the initial contention window size. Surprisingly, we see that the receiver receives some packets regardless of the packets size. It receives almost the same number of packets for the two transmitters. Sometimes this number is different, but it seems to come from the people movement in the laboratory (we succeeded in stopping the other mobiles of the laboratory but not the movements of the team!). When we analyze the logs, we can observe three scenarios: there is a collision at the receiver, the two emitters transmit at the same time but the receiver decodes one of the packets, one emitter senses the medium busy by the other transmitter. This last point rarely happens but it happens. This combination of scenarios results in a reasonable number of successful transmissions since 58% of the medium capacity is used for such emissions. Finally, we can note that this number decreases for packets of

1500 and 1600 bytes. This is explained by the fragmentation mechanism that sends two packets in place of one. Therefore a frame has less chance to be correctly received.




Figure 5: The hidden terminals: experimental results with broadcasted packets

Figure 7 gives the experimental results with unicast packets without RTS-CTS. The axes are the same as in the previous figure. First, we see that the received packets fairly come from the two emitters (except at some periods). Since the unicast packets are acknowledged, when collisions happen, the contention window size is increased. Therefore, the emitters should wait during longer periods than with broadcasted packets, what should result in a lower throughput for the two emitters. Figure 7 shows that it is not the case since more packets are received. It seems that the difference in the contention windows size when an emitter succeeds in correctly transmitting and the synchronization of the two emitters that results from this transmission allow this emitter to send more consecutive packets and cancels this longer wait phase when collisions happen. Anyway, this balance decreases with the packets size, since at 100 bytes, the successful transmissions represent 72% of the medium capacity whereas at 1000 bytes, they only represent 53%. Indeed, when the frame size increases, the emitter that is in its backoff decrementation phase has more time to decrement its backoff, what leads to more collisions.

Now, if we compare the experimental results with the simulation results (Figure 4), we see that the experimental throughputs are higher and do not decrease with the packets size as in simulations. It can be explained by the phenomena observed with broadcasted packets: collisions do not always happen even if the two transmitters emit at the same time and carrier sensing is effective sometimes. Thus more packets are correctly received whatever the packets size may be in experiments. This can be more easily understood with the chronogram of Figure 6. This chronogram represents on an horizontal time scale the packets sent by the two emitters and the packets received by the receiver: a small rectangle corresponds to a sent or received packet. The upper and lowest lines give the packets sent

by each of the emitters (the upper line for the emitter A and the lowest line for the emitter C). The central one corresponds to the packets received by B . For each packet, its emitter is indicated by a specific color. We can see at the position (A) of the chronogram that when A and C transmit at the same time, there can be a collision and thus no reception at node B . But we also see that sometimes one of the two simultaneous packets is however received, like at the positions (B) and (C). And finally, sometimes C can hear A and thus defers its own transmission in the position (D).




Figure 6: Partial chronogram of the hidden nodes experiment




Figure 7: The hidden terminals: experimental results with unicast packets without RTS-CTS

Figure 7 gives the experimental results with unicast packets with RTS-CTS. The axes are the same as in the previous figures. First, we see that the throughputs are not equal even on periods of 50 seconds. This experiment has been carried out several times and the same kind of results have always been obtained. We have not been able to find an explanation so far. The signal of one emitter received by the receiver may be more powerful but in that case, we do not understand why it does not appear on the previous experiments. This difference in throughput is around 200 kb/s in average. Anyway, the aggregated throughput is high since it is most equal to the medium capacity regardless of the packets size.

If the experimental results are a little bit higher than the ones obtained with simulations, the differences are nevertheless reduced compared to previous scenario without the RTS-CTS (see Figure 4 for the simulation results). The medium reservation is effective both

in simulations and experiments and the phenomena observed with the broadcasted packets increase the probability to successfully transmit.

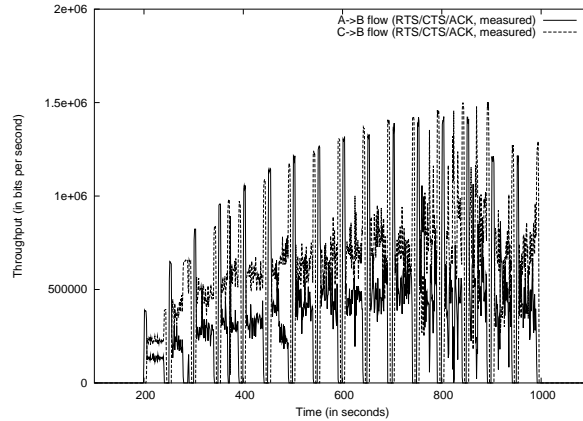


Figure 8: The hidden terminals: experimental results with unicast packets with RTS-CTS

We have also evaluated the size of the bursts, *i.e.* the number of consecutive packets sent by an emitter that has first won the contention. Surprisingly, the measured average size of the bursts is almost the same without and with a RTS-CTS exchange and is around 4. We expected a higher value with RTS-CTS since this exchange synchronizes the mobiles that will enter in the contention phase at the same time. Therefore, the emitter that lost the medium access should take more time to decrement its backoff unlike in the scenario without RTS-CTS where this emitter will often decrease its backoff during the emission of the other source. One possible explanation is that this parallel backoff decrementation during a transmission results in multiple collisions, what reduces the average size of the bursts. With RTS-CTS, simulations exhibit bursts of 9 to 10 packets long in average, regardless packet size. Without RTS-CTS, the simulated bursts length depends on packet size: from 30 packets long in average for 100 bytes packets to less than 4 consecutive packets for 1400 bytes packets. This difference between simulation and experimentation can, once again, be explained with the possible transmissions that arise in experiments and not in simulation, resulting in a smaller increase of the contention window size.

4.3.2 A two pairs scenario

In this configuration, we consider the scenario of Figure 2 and we remove one of the exterior pairs. Thus, we obtain a scenario with two pairs of mobiles where each pair is in the carrier sense range of the other. That means that when a pair exchanges a packet, the emitter of the other pair will have to wait for a *EIFS* after the acknowledgment in the concurrent pair, whereas the concurrent emitter will wait a *DIFS*. Then this last mobile is more likely to access the medium than the other transmitter.

Whereas the simulation results give an average burst of four packets, the experiments do not reveal the short time unfairness since each emitter alternates in the medium access. Once again, it makes us think that the *EIFS* is not triggered by the used wireless cards.

5 Conclusion

In this paper, we have evaluated via experimentations some performance issues encountered in the literature. We focus our evaluation on the fairness issue. In particular, we have shown that some fairness issues appear like in simulations but the effects are softened by the instability of the radio medium. The communication and the carrier sense areas are not fixed like in simulations and are sometimes affected by the environment. Thus, independent mobiles share the medium at some periods or one of the concurrent transmissions can be correctly received. More successful communications are then possible even for the mobiles penalized. We have also shown that with the used cards, the *EIFS* was not triggered and a *DIFS* was used in place. The fairness issues caused by this mechanism are therefore nonexistent. Finally, the experiments on the hidden terminals configurations reveal that the short time unfairness is not so high since the average burst is around 4. Moreover, the number of successful received packets are surprisingly high when the RTS-CTS are not activated.

Even if the unfairness is not so strong in experiments than in simulations, some configurations (like in Figure 2 and in Figure 1) show a big difference in the flows' throughput. This fact may have an impact on protocols on higher layers and we think that the design of appropriate MAC protocols is still of real interest. Some enhancements have already been proposed in that end, but they concern specific solutions that often solve isolated configurations and not the whole set of performance issues. A more general solution is expected.

Finally, these results also show that some efforts should be focused on the implementation on more accurate propagation models in simulators in order to obtain simulation results closer to reality.

References

- [1] Brahim Bensaou, Yu Wang, and Chi Chung Ko. Fair Medium Access in 802.11 Based Wireless Ad-Hoc Networks. In *Proceedings of the First International Workshop on Mobile Ad Hoc Networking and Computing (MobiHoc 2000)*, Boston, Massachusetts, USA, August 2000.
- [2] Vaduvur Bharghavan, Alan Demers, Scott Shenker, and Lixia Zhang. MACAW: a media access protocol for wireless LAN's. In *Proceedings of the conference on Communications architectures, protocols and applications (ACM Sigcomm '94)*, pages 212–225, Londres, Royaume-Uni, August 1994.

- [3] G. Bianchi. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE JSAC*, 18(3):1265–1275, March 2000.
- [4] F. Cali, M. Conti, and E. Gregori. IEEE 802.11 Wireless LAN: Capacity Analysis and Protocol Enhancement. In *Proceedings of Infocom*, 1998.
- [5] D. Dhoutaut and I. Guérin-Lassous. Experiments with 802.11b in ad hoc configurations. In *Proceedings of 14th IEEE International Symposium Personal, Indoor and Mobile Radio Communications*, Beijing, China, September 2003.
- [6] Dominique Dhoutaut and Isabelle Guérin Lassous. Impact of Heavy Traffic Beyond Communication Range in Multi-Hops Ad Hoc Networks. In *Proceedings of the Third International Network Conference (INC 2002)*, Plymouth, Royaume-Uni, July 2002.
- [7] Dominique Dhoutaut and Isabelle Guérin Lassous. Performance of a multi-hops configuration with 802.11: from simulation to experimentation. In *Proceedings of the fifteenth IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC 2004)*, Barcelona, Spain, September 2004. IEEE Press.
- [8] M. S. Gast. *802.11 Wireless Networks*. O’Reilly, 2002.
- [9] Martin Heusse, Franck Rousseau, Gilles Berger-Sabbatel, and Andrzej Duda. Performance Anomaly of 802.11b. In *Proceedings of the Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies (Infocom 2003)*, pages 836–843, San Francisco, Californie, USA, April 2003.
- [10] G. Holland and N. H. Vaidya. Analysis of TCP performance over mobile ad hoc networks. In *Proc. of MobiCom*, pages 219–230, 1999.
- [11] IEEE Standard for Information Technology – Telecommunications and Information Exchange between Systems. Local and Metropolitan Area Network – Specific Requirements – Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, 1997.
- [12] Jinyang Li, Charles Blake, Douglas S.J. De Couto, Hu Imm Lee, and Robert Morris. Capacity of Ad Hoc wireless networks. In *Proceedings of the Seventh annual international conference on Mobile computing and networking (MobiCom 2001)*, pages 61–69, Rome, Italie, July 2001.
- [13] Zhifei Li, Sukumar Nandi, and Anil Gupta. Modeling the Short-Term Unfairness of IEEE 802.11 in Presence of Hidden Terminals. In *Proceedings of the Third IFIP-TC6 Networking Conference (Networking 2004)*, pages 613–625, Athènes, Grèce, May 2004.
- [14] Zhifei Li, Sukumar Nandi, and Anil Gupta. Improving MAC Performance in Wireless Ad Hoc Networks Using Enhanced Carrier Sensing (ECS). In *Proceedings of the Third IFIP-TC6 Networking Conference (Networking 2004)*, pages 600–612, Athènes, Grèce, May 2004.

-
- [15] H. Lundgren, D. Lundberg, J. Nielsen, E. Nordstrom, and C. Tschudin. A Large-scale Testbed for Reproducible Ad hoc Protocol Evaluations. In *Proceedings of the 3rd annual IEEE Wireless Communications and Networking Conference*, 2002.
 - [16] David A. Maltz, Josh Broch, and David B. Johnson. Quantitative Lessons From a Full-Scale Multi-Hop Wireless Ad Hoc Network Testbed. In *Proceedings of the IEEE Wireless Communications and Networking Conference*, Chicago, USA, September 2000.
 - [17] Y. C. Tay and K. C. Chua. A Capacity Analysis for the IEEE 802.11 MAC Protocol. *Wireless Networks*, 7:159–171, 2001.
 - [18] C.-K. Toh, Richard Chen, Minar Delwar, and Donald Allen. Experimenting with an Ad Hoc Wireless Network on Campus: Insights and Experiences. *ACM SIGMETRICS Performance Evaluation Review*, 28(3):21–29, 2001.
 - [19] V. Vishnevsky and A. Lyakhov. 802.11 LANs: Saturation Throughput in the Presence of Noise. In E. et als Gregori, editor, *Proceedings of Networking*, volume 2345 of *LNCS*, pages 1008–1019. Springer-Verlag, 2002.
 - [20] S. Xu and T. Saadawi. Does the IEEE 802.11 MAC Protocol Work Well in Multihop Wireless Ad Hoc Networks? *IEEE Communications Magazine*, pages 130–137, June 2001.



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