Performance Anomaly of 802.11b

Martin Heusse, Franck Rousseau, Gilles Berger-Sabbatel, Andrzej Duda

LSR-IMAG Laboratory
Grenoble, France
email: {heusse, rousseau, gberger, duda}@imag.fr

Abstract— **We analyze the performance of the IEEE 802.11b wireless local area networks. We have observed that when some mobile hosts use a lower bit rate than the others, the performance of all hosts is considerably degraded. Such a situation is a common case in wireless local area networks in which a host far away from an Access Point is subject to important signal fading and interference. To cope with this problem, the host changes its modulation type, which degrades its bit rate to some lower value. Typically, 802.11b products degrade the bit rate from 11 Mb/s to 5.5, 2, or 1 Mb/s when repeated unsuccessful frame transmissions are detected. In such a case, a host transmitting for example at 1 Mb/s reduces the throughput of all other hosts transmitting at 11 Mb/s to a low value below 1 Mb/s. The basic CSMA/CA channel access method is at the root of this anomaly: it guarantees an equal long term channel access probability to all hosts. When one host captures the channel for a long time because its bit rate is low, it penalizes other hosts that use the higher rate. We analyze the anomaly theoretically by deriving simple expressions for the useful throughput, validate them by means of simulation, and compare with several performance measurements.**

I. INTRODUCTION

We focus our paper on wireless local area networks such as IEEE 802.11 that have become popular as access networks to the wireless mobile Internet. They can be deployed in hot spots areas and offer performance comparable to wired local area networks. The question of the performance effectively perceived by mobile hosts becomes increasingly important as many new emerging applications such as mobile information access, real-time multimedia communications, networked games, immersion worlds, cooperative work require sufficient bandwidth.

Although IEEE 802.11b provides a means for allocating a part of the radio channel bandwidth to some hosts (PCF - Point Coordination Function), the commonly available access method (DCF - Distributed Coordination Function) uses the CSMA/CA protocol to share the radio channel in a fair way. However, we have observed that in some common situations in a wireless environment, the method results in a considerable performance degradation. In a typical wireless local area network, some hosts may be far away from their access point so that the quality of their radio transmissions is low. In this case current 802.11b products degrade the bit rate from the nominal 11 Mb/s rate to 5.5, 2, or 1 Mb/s – when a host detects repeated unsuccessful frame transmissions, it decreases its bit rate. If there is at least one host with a lower rate, a 802.11 cell presents a *performance anomaly*: the throughput of all hosts transmitting at the higher rate is degraded below the level of the lower rate. Such a behavior penalizes fast hosts and privileges the slow one. The reason for this anomaly is the basic CSMA/CA channel access method which guarantees that the long term channel access probability is equal for all hosts. When one host captures the channel for a long time because its bit rate is low, it penalizes other hosts that use the higher rate.

In this paper, we analyze the performance of the 802.11b DCF access method by deriving simple expressions for the available throughput (Section II). Then we extend the analysis to the case of hosts with different bit rates (Section III). We validate the analysis by means of simulation (Section IV) and measurements. Section V provides several performance measurements and comparisons to illustrate the anomaly. Finally, we briefly discuss related works (Section VI) and present some conclusions (Section VII).

II. PERFORMANCE OF IEEE 802.11B DCF ACCESS METHOD

The IEEE 802.11b standard [1] defines two access methods: the Distributed Coordination Function (DCF) that uses CSMA/CA to allow for contended access to the wireless media and the Point Coordination Function (PCF) providing uncontended access via arbitration by a Point Coordinator, which resides in the Access Point. The DCF method provides a best effort type of service whereas the PCF guarantees a time-bounded service. Both methods may coexist, a contention period following a contention-free period. PCF would be especially well suited for real-time traffic as it permits to allocate the radio channel according to applications requirements, but the PCF method is not implemented in current 802.11 products.

The DCF access method is based on the CSMA/CA principle in which a host wishing to transmit senses the channel, waits for a period of time (DIFS – Distributed Inter Frame Space) and then transmits if the medium is still free. If the packet is correctly received, the receiving host sends an ACK frame after another fixed period of time (SIFS – Short Inter Frame Space). If this ACK frame is not received by the sending host, a collision is assumed to have occurred. The sending host attempts to send the packet again when the channel is free for a DIFS period augmented of a random amount of time.

Let first consider that a single host in a 802.11b cell transmits a single data frame. If we neglect propagation times, the overall transmission time is composed of the transmission
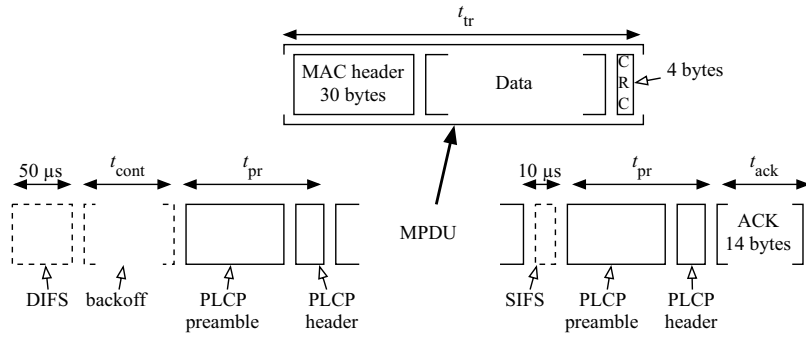
Fig. 1. Successful transmission of a single frame

time and a constant overhead:

$$T = t_{\mathrm{tr}} + t_{\mathrm{ov}} \qquad (1)$$

where the constant overhead

$$t_{\mathrm{ov}} = DIFS + t_{\mathrm{pr}} + SIFS + t_{\mathrm{pr}} + t_{\mathrm{ack}}$$

is composed of the PLCP (Physical Layer Convergence Protocol) preamble and header transmission time t_{pr}, $SIFS = 10~\mu s$, t_{ack} is the MAC acknowledgment transmission time (10 μs if the selected rate is 11 Mb/s, as the ACK length is 112 bits), and $DIFS = 50~\mu s$. t_{tr} is the frame transmission time (see Figure 1). t_{pr} varies according to the bit rate used by the host. When it transmits at 1 Mb/s, the long PLCP header is used and $t_{\mathrm{pr}} = 192~\mu s$. If it uses 2, 5.5 or 11 Mb/s, then $t_{\mathrm{pr}} = 96~\mu s$ (short PLCP header). For bit rates greater than 1 Mb/s and the frame size of 1500 bytes of data (MPDU of total 1534 bytes), proportion p of the useful throughput measured above the MAC layer will be:

$$p = \frac{t_{\mathrm{tr}}}{T} \times \frac{1500}{1534} = 0.70. \qquad (2)$$

So, a single host sending long frames over a 11 Mb/s radio channel will have a maximum useful throughput of 7.74 Mb/s.

If there are multiple hosts attempting to transmit, the channel may be sensed busy and hosts enter a collision avoidance phase: a host executes the exponential backoff algorithm – it waits for a random interval distributed uniformly between $[0, CW] \times SLOT$. The congestion window CW varies between $CW_{\min} = 31$ and $CW_{\max} = 1023$, the value of $SLOT$ is 20 μs (these parameters are for 802.11b). The host that chooses the smallest interval starts transmitting and the others hold counting down until the transmission is over. Each time a host happens to collide, it doubles CW up to $CW_{\max}$.

In the case of a single host that tries to transmit a sustained traffic (the host has always a packet to send), the carrier sense applies also to the host's own transmissions, so that it inserts a random interval between each transmission. This is mandatory, because transmitting frames continuously would prevent any other host from accessing the channel. Thus, Eq. 2 gives an upper bound that can be attained only for a host passing packets to the MAC layer at the moment the previous transmission is completed.

Collisions and the exponential backoff mechanism influence the performance of 802.11b for an increasing number of hosts. The overall frame transmission time experienced by a single host when competing with $N - 1$ other hosts has to be increased by some interval t_{cont} that accounts for the time spent in contention procedures. So the overall transmission time becomes:

$$T(N) = t_{\mathrm{tr}} + t_{\mathrm{ov}} + t_{\mathrm{cont}}(N) \qquad (3)$$

The analytical formula for $t_{\mathrm{cont}}(N)$ is difficult to derive (see [2]) and we propose to use a simple approximation. Considering that the hosts always sense a busy channel when they attempt to transmit and that the number of transmissions that are subject to multiple successive collisions is negligible we find:

$$t_{\mathrm{cont}}(N) \simeq SLOT \times \frac{1 + P_{\mathrm{c}}(N)}{2\,N} \times \frac{CW_{\min}}{2}$$

where $P_{\mathrm{c}}(N)$ is the proportion of collisions experienced for each packet successfully acknowledged at the MAC level ($0 \leqslant P_{\mathrm{c}}(N) < 1$).

A simple expression for $P_{\mathrm{c}}(N)$ can be derived by considering that a host attempting to transmit a frame will eventually experience a collision if the value of the chosen backoff interval corresponds to the residual backoff interval of at least one other host. Such an approximation holds if multiple successive collisions are negligible. So we have

$$P_{\mathrm{c}}(N) = 1 - (1 - 1/CW_{\min})^{N-1}. \qquad (4)$$

Finally, proportion p of the useful throughput that can be obtained by a host depends on the number of hosts and is given by:

$$p(N) = t_{\mathrm{tr}}/T(N) \qquad (5)$$

A. Discussion

The analysis in this section shows that the throughput that can be obtained over the 802.11b WLAN is much smaller than the nominal bit rate of 11 Mb/s, for example, if there are no collisions, one host may expect the maximum throughput of 7.74 Mb/s. Furthermore, the proportion of the useful throughput strongly depends on the number of competing hosts.

III. Performance Anomaly of IEEE 802.11b

Consider now the situation in which N hosts of different bit rate compete for the radio channel: $N-1$ hosts use the high transmission rate $R = 11$ Mb/s and one host transmits at a degraded rate $r = 5.5, 2,$ or 1 Mb/s. In this case, the frame transmission time depends on the rate: $t_{\mathrm{tr}} = s_{\mathrm{d}}/R$ or $t_{\mathrm{tr}} = s_{\mathrm{d}}/r$, where s_{d} is the data frame length in bits. The MAC layer ACK frame is also sent at the rate that depends on the host speed, thus we denote by t_{ov}^R and t_{ov}^r the associated overhead time.

Let T_{f} be the overall transmission time of a "fast" host transmitting at rate R:

$$T_{\mathrm{f}} = t_{\mathrm{ov}}^R + \frac{s_{\mathrm{d}}}{R} + t_{\mathrm{cont}}.$$

Similarly, let T_{s} be the corresponding time for a "slow" host transmitting at rate r:

$$T_{\mathrm{s}} = t_{\mathrm{ov}}^r + \frac{s_{\mathrm{d}}}{r} + t_{\mathrm{cont}}.$$

We can analyze the overall performance by assuming that the hosts alternate transmissions: channel utilization by a fast host is the ratio between the time to send one packet (T_{f}) and the time during which all other hosts transmit once with all possible collisions. This holds on a long term, because the long term channel access probability CSMA/CA is equal for all hosts. However, this is not true for a short term – it has been shown that the short term behavior of CSMA/CA is not fair [3]. So, we will have

$$U_{\mathrm{f}} = \frac{T_{\mathrm{f}}}{(N-1)\, T_{\mathrm{f}} + T_{\mathrm{s}} + P_{\mathrm{c}}(N) \times t_{\mathrm{jam}} \times N} \qquad (6)$$

where t_{jam} is the average time spent in collisions. The duration of a collision depends on the type of the hosts involved in the collision (slow or fast). t_{jam} can be found by considering all possible pairs between $N-1$ fast hosts and a slow one (see Appendix):

$$t_{\mathrm{jam}} = \frac{2}{N}T_{\mathrm{s}} + (1 - \frac{2}{N})\, T_{\mathrm{f}}.$$

The throughput at the MAC layer of each of the $N-1$ fast hosts is

$$X_{\mathrm{f}} = U_{\mathrm{f}} \times p_{\mathrm{f}}(N) \times R$$

where $p_{\mathrm{f}}(N) = \frac{s_{\mathrm{d}}}{R\, T_{\mathrm{f}}}$. Finally, we obtain

$$X_{\mathrm{f}} = \frac{s_{\mathrm{d}}}{(N-1)\, T_{\mathrm{f}} + T_{\mathrm{s}} + P_{\mathrm{c}}(N) \times t_{\mathrm{jam}} \times N} \qquad (7)$$

Similarly, we can express the channel utilization of the slow host as

$$U_{\mathrm{s}} = \frac{T_{\mathrm{s}}}{(N-1)\, T_{\mathrm{f}} + T_{\mathrm{s}} + P_{\mathrm{c}}(N) \times t_{\mathrm{jam}} \times N} \qquad (8)$$

so that the throughput at the MAC layer of the slow host is

$$X_{\mathrm{s}} = U_{\mathrm{s}} \times p_{\mathrm{s}}(N) \times r$$

where $p_{\mathrm{s}}(N) = \frac{s_{\mathrm{d}}}{r\, T_{\mathrm{s}}}$. Finally, we obtain the following:

Result. *The fast hosts transmitting at the higher rate R obtain the same throughput as the slow host transmitting at the lower rate r:*

$$X_{\mathrm{s}} = X_{\mathrm{f}} = X. \qquad (9)$$

A. Discussion

The result of this section show that when one host that uses a lower bit rate competes with other hosts, the throughput of all hosts may be significantly limited – the fast hosts see their throughput decreased roughly to the order of magnitude of the slow host's throughput. When fast and slow hosts share the radio channel of a 802.11b cell, the throughput is bounded by $N \times X$ instead of $R \times p_{\mathrm{f}}(N)$ as we may expect in presence of a large number of fast hosts.

The fair access to the channel provided by CSMA/CA causes a slow host transmitting at 1 Mb/s to capture the channel eleven times longer than hosts emitting at 11 Mb/s. This degrades the overall performance perceived by the users in the considered cell, and this anomaly holds whatever is the proportion of slow hosts.

B. UDP traffic

To be able to compare the analytical expressions with measurements, we should consider an application level scenario that governs the traffic generated at the MAC level. We consider two basic scenarios. In the first one, N hosts (one of these hosts may send at the lower rate) generate a connectionless UDP stream to a host behind the Access Point. As the streams are not acknowledged, the only traffic generated at the MAC level is composed of UDP data packets. In this case, the throughput experienced by each host will be given by Eq. 7.

C. TCP traffic

In the second scenario, there are $N-1$ hosts that send data over TCP connections to a server behind the Access Point. We assume that the server is connected to the Access Point via a 100 Mb/s switched Ethernet, so that the wireless link is the only bottleneck. In this configuration, TCP data segments and ACKs travel only once over the wireless channel. As data segments are acknowledged by the server, we have a total of N hosts competing for the radio channel ($N-1$ hosts and the Access Point that forwards TCP ACKs from the server).

We denote by T_{f}^{ack} (resp. T_{s}^{ack}) the time to transmit a frame containing a TCP ACK to a fast (resp. slow) host at rate R (resp. r). In the TCP scenario, $N-2$ hosts send their packets at higher rate R of 11 Mb/s, one host sends at lower rate r (1, 2, 5.5 Mb/s), and the Access Point sends TCP ACKs at a variable rate that depends on the rate of the host which is the destination of the ACK. As TCP sends ACKs for every other data segment transmitted (we have effectively observed such a behavior in our experiments), the throughput in the TCP scenario can be expressed as:

$$X = \frac{s_{\mathrm{d}}}{(N-2)(T_{\mathrm{f}} + \frac{T_{\mathrm{f}}^{ack}}{2}) + T_{\mathrm{s}} + \frac{T_{\mathrm{s}}^{ack}}{2} + P_{\mathrm{c}}(N)\, t_{\mathrm{jam}} \frac{3}{2}(N-1)} \qquad (10)$$

The derivation of t_{jam} is given in Appendix.

IV. SIMULATION

To validate the analytical expressions given in the previous section and to gain insight into the complex backoff mechanism of CSMA/CA we have developed a simple 802.11b simulator. It simulates the CSMA/CA access method under heavy load, i.e. we assume that the channel is always sensed busy when hosts need to transmit. In this way we can evaluate a worst case bound for $P_c(N)$. All the parameters of the simulation such as $CW_{\min}, CW_{\max}, SLOT$ have values defined in the 802.11b standard.

Figure 2 shows the simulated value of $P_c(N)$ for a varying number of hosts and compares it with Eq. 4 as well as with several measured values. Each point of the graph comes from the simulation of 1000000 transmitted packets.
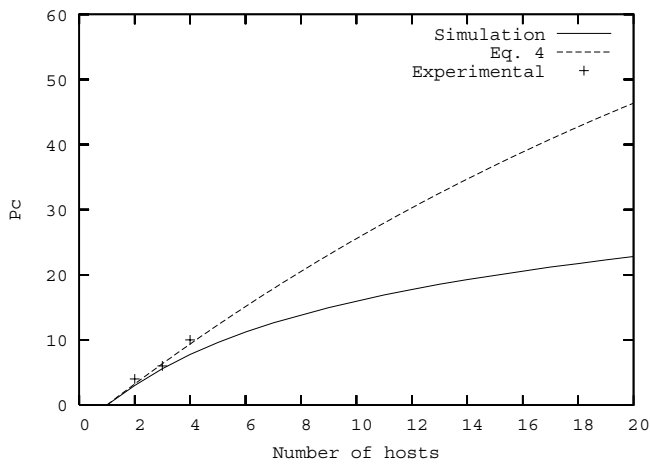
Fig. 2. Proportion of collisions P_c for a varying number of hosts

As shown in the previous section, all hosts obtain the same throughput even if one of them transmits at a lower rate. The throughput is shown in Figure 3 for different rates of the slow host and for a varying number of hosts in the cell. The throughput curves obtained from simulation or expression 7 are superimposed.

Finally, we give the cumulative throughput of all hosts. Although the presence of a slow host decreases the throughput perceived by the other ones, its impact on the cumulative throughput decreases when the number of hosts increases (cf. Figure 4). This is due to the fact that the slow host uses a diminishing proportion of the channel as the number of hosts rises.

A. Discussion

Our simulation results show that the precision of the analytical expressions is fairly good. The expression for the proportion of collisions slightly overestimates the impact of collisions. However, this discrepancy results in a small error in the throughput formulae, because the proportion of collisions is only a factor of the average time spent in collisions t_{jam}.
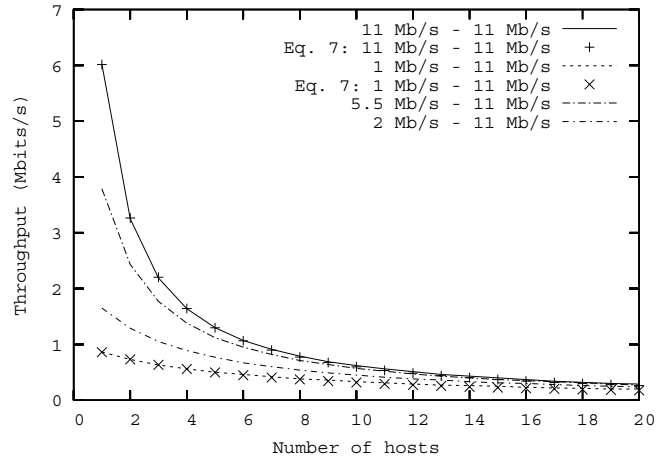

Fig. 3. Throughput experienced by a 802.11b host when all hosts except one transmit at 11Mb/s
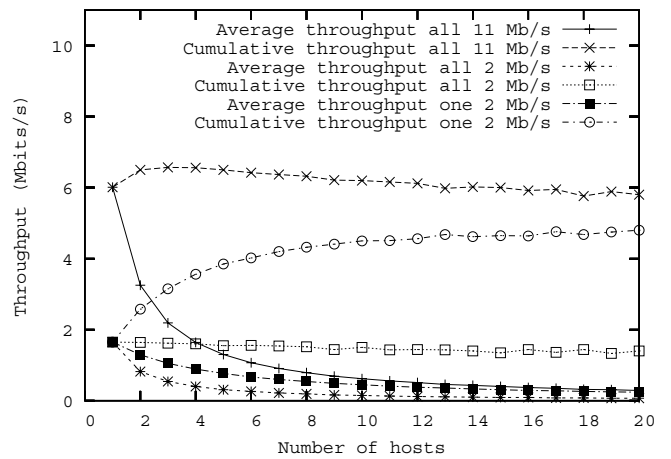

Fig. 4. Cumulative throughput in a cell with a single slow host

V. PERFORMANCE MEASUREMENTS

We have set up a platform to measure the throughput that hosts can obtain when sharing a 11 Mb/s 802.11b wireless cell. We have used four notebooks (*Marie, Milos, Kea*, and *Bali*) running Linux RedHat 7.3 (kernel 2.4.18) with 802.11b cards based on the same chipset (Lucent Orinoco and Compaq WL 110). The wired part of the network is connected by a Lucent Access Point. The notebooks use the Wvlan driver for the wireless cards. The cards do not use the RTS/CTS option that may optimize performance in case of the hidden terminal problem.

We measure the throughput using three tools:

- `netperf` generates TCP or UDP traffic to a target host running `netserver` and measures the throughput at the end of the session [4].
- `tcpperf` generates TCP traffic and measures the throughput obtained during each second.

- `udpperf` generates UDP traffic and measures the throughput obtained during each second.

The measurements are done using `netperf` and compared with the results of `tcpperf` and `udpperf`. We generate all traffic to a host on the wired part of the network.

A. Hosts with different rates, no mobility

In the first experiment, we place all notebooks near the Access Point to obtain good transmission conditions on the radio channel and we force one of them to work at a degraded bit rate. We measure the throughput for a varying number of hosts and different traffic conditions.

Table I compares the measured throughput with the expressions derived in Section III when hosts generate TCP streams. They compete with the Access Point that sends TCP ACKs on behalf of the destination ($N = 2, 3, 4, 5$). The table references figures that present the evolution of the throughput in time for a given number of hosts.

Table 10 compares the throughput obtained in the same setup with UDP streams. The hosts compete with each other for the radio channel ($N = 1, 2, 3, 4$).

Fig. 6. Measured throughput in Mb/s for two hosts, UDP traffic

Fig. 5. Measured throughput in Mb/s for two hosts, TCP traffic

B. Discussion

We can observe that the measured values correspond fairly well to the analytical expressions. The precision of the analytical expression is better for the UDP traffic than for that of TCP. In the case of TCP, the traffic pattern is much more complex than that of UDP, because the Access Point competes with other hosts when sending TCP ACKs. In addition to that, the traffic of TCP ACKs is governed by a complex dependence on the overall round trip time and the bottleneck link between the source and the destination. It can become correlated with the traffic of data segments – a TCP ACK is sent upon arrival of a data segment.

The influence of the estimate for the proportion of collisions on the overall throughput is limited, because it only accounts for a small factor in Eq. 7.
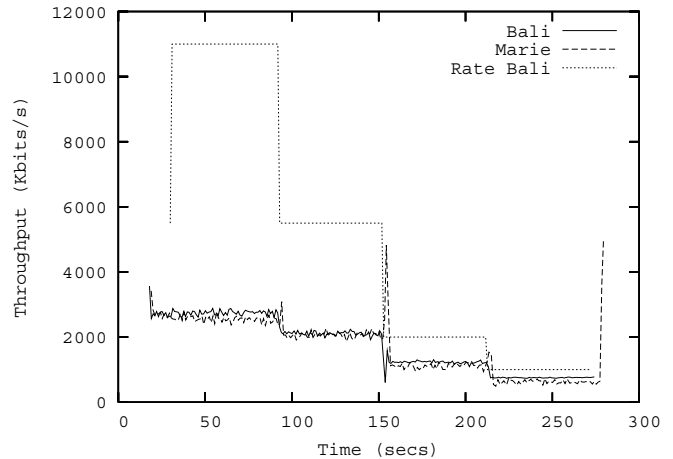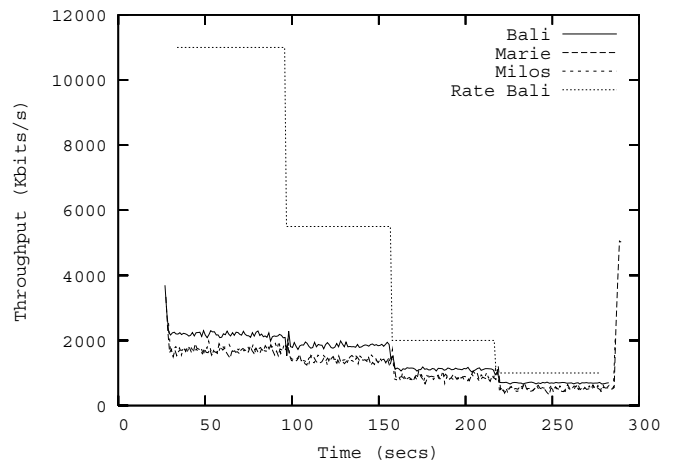

Fig. 7. Measured throughput in Mb/s for three hosts, TCP traffic

C. Hosts with different rates, real mobility

In the second experiment, *Bali* is a mobile host that moves around and its bit rate automatically adapts to varying transmission conditions. The other hosts are located near the Access Point and their transmission conditions are good.

Figure 11 shows the evolution of the throughput in time when two hosts send TCP traffic. *Marie* keeps its bit rate at 11 Mb/s and *Bali* changes its bit rate according to the quality of the radio channel (cf. Figure 12). They compete with the Access Point that sends ACKs on behalf of the destination ($N = 3$).

It can be observed that when transmission conditions are bad (period 300-380) the throughput of *Marie* increases. This can be explained by the fact that in adverse conditions the TCP source of *Bali* limits its sending rate so that *Marie* may benefit from the unused channel capacity. This situation becomes more explicit after instant 380 when *Bali* stops sending even if its bit rate is not zero while *Marie* gains almost all the

Bit rate of Bali	Bali	Marie	Milos	Kea	Eq. 10	observed P_c	Eq. 4	Figure No.
11	5.08	-	-	-	7.21	4%	3.1%	-
5.5	3.37	-	-	-	4.12	4%	3.1%	-
2	1.55	-	-	-	1.65	4%	3.1%	-
1	0.83	-	-	-	0.83	4%	3.1%	-
11	2.52	2.48	-	-	3.48	8%	6.2%	5
5.5	2.04	1.96	-	-	2.51	8%	6.2%	5
2	1.12	1.11	-	-	1.27	8%	6.2%	5
1	0.67	0.63	-	-	0.70	8%	6.2%	5
11	1.65	1.54	1.52	-	2.24	8%	9.1%	7
5.5	1.36	1.39	1.47	-	1.77	8%	9.1%	7
2	0.83	0.89	0.95	-	1.02	8%	9.1%	7
1	0.57	0.52	0.64	-	0.62	8%	9.1%	7
11	1.15	0.92	1.21	1.02	1.63	9%	12.0%	9
5.5	1.16	0.83	1.11	0.83	1.35	9%	12.0%	9
2	0.74	0.41	0.65	0.55	0.85	9%	12.0%	9
1	0.56	0.32	0.41	0.27	0.53	9%	12.0%	9

TABLE I

MEASURED THROUGHPUTS IN MB/S FOR A VARYING NUMBER OF HOST, TCP TRAFFIC

Bit rate of Bali	Bali	Marie	Milos	Kea	Eq. 7	observed P_c	Eq. 4	Figure No.
11	3.09	3.36	-	-	3.26	4%	3.1%	6
5.5	2.38	2.42	-	-	2.44	4%	3.1%	6
2	1.30	1.26	-	-	1.29	4%	3.1%	6
1	0.76	0.76	-	-	0.73	4%	3.1%	6
11	2.26	2.0	2.23	-	2.20	6%	6.2%	8
5.5	2.01	1.56	1.89	-	1.77	6%	6.2%	8
2	1.17	0.90	1.16	-	1.06	6%	6.2%	8
1	0.74	0.58	0.69	-	0.63	6%	6.2%	8
11	1.71	1.41	1.8	1.41	1.64	10%	9.1%	10
5.5	1.66	1.16	1.59	1.19	1.38	10%	9.1%	10
2	0.96	0.84	0.81	0.72	0.89	10%	9.1%	10
1	0.69	0.47	0.63	0.49	0.56	10%	9.1%	10

TABLE II

MEASURED THROUGHPUTS IN MB/S FOR A VARYING NUMBER OF HOSTS, UDP TRAFFIC

available throughput.

Figure 13 shows a similar graph when the two hosts send UDP traffic ($N = 2$). *Marie* keeps its bit rate at 11 Mb/s and *Bali* changes its bit rate according to the quality of the radio channel (cf. Figure 14).

The results are similar to the TCP traffic, but we can observe that now *Marie* does not gain as much throughput when the transmission conditions are bad (period 300-380) unless *Bali* stops sending data at instant 420.

VI. RELATED WORK

Many papers have studied the performance of 802.11 WLANs, however all of them assume that all hosts communicate using the same bit rate. The analysis is difficult so that many papers use simulation [5], [6]. Only a few papers analyze 802.11 analytically using Markov chains [2], [7]. This approach models the complex behavior of 802.11 better, but usually gives complex results or requires numerical resolution.

Short-term unfairness of CSMA-based medium access protocols has also been a topic of interest. In [3] the authors use experimental and analytical methods to evaluate the short-term fairness degree of the CSMA-CA as implemented in the WaveLAN network.

VII. CONCLUSIONS

We have analyzed the performance of the IEEE 802.11b wireless local area networks. This analysis shows that the throughput of the 802.11b WLAN is much smaller than the
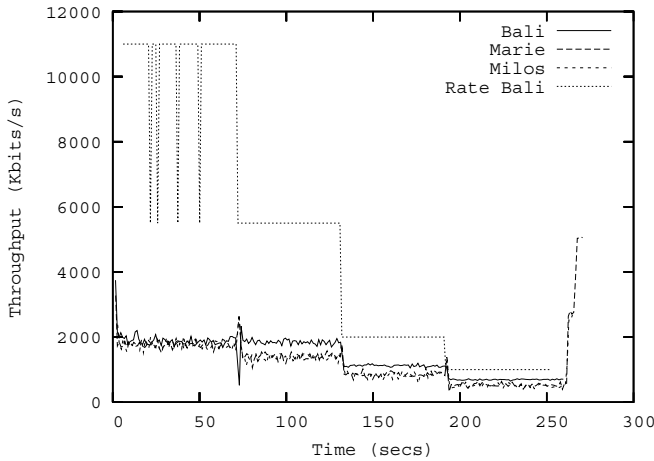
Fig. 8. Measured throughput in Mb/s for three hosts, UDP traffic
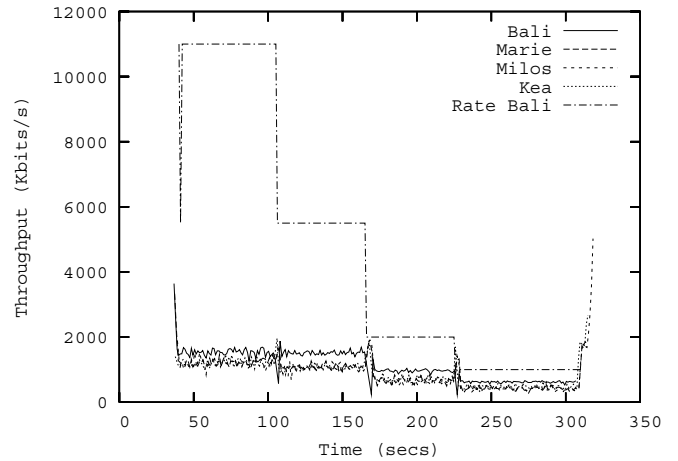

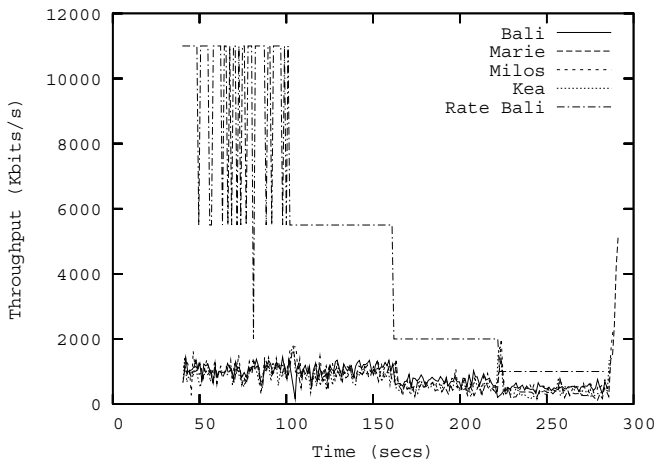Fig. 10. Measured throughput in Mb/s for four hosts, UDP traffic

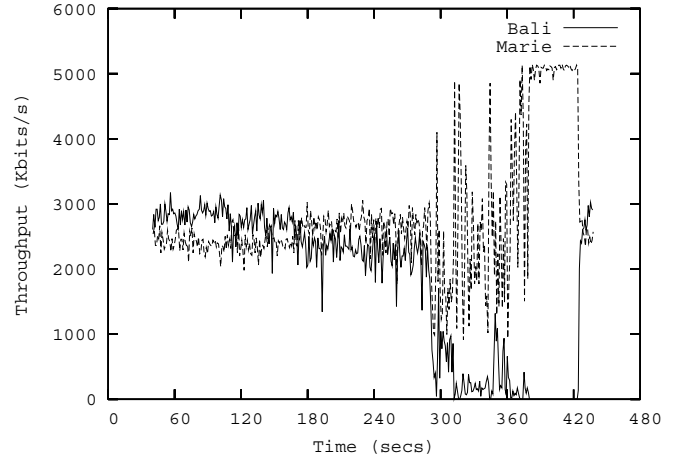Fig. 9. Measured throughput in Mb/s for four hosts, TCP traffic

Fig. 11. Measured throughput in Mb/s for two hosts (one mobile), TCP traffic

nominal bit rate. Furthermore, the proportion of the useful throughput strongly depends on the number of competing hosts.

When mobile hosts move, they may encounter bad transmission conditions and degrade the bit rate from 11 Mb/s to 5.5, 2, or 1 Mb/s. We have analyzed how a host with a lower bit rate influences the throughput of other hosts that share the same radio channel. Our analysis and performance measurements show that the slow host may considerably limit the throughput of other hosts roughly to the level of the lower rate.

However, this adverse performance effect should be alleviated by the observation that in real conditions and for TCP traffic, the host that degrades its bit rate will be anyway subject to important packet losses which in turn limit its sending rate. In this way, other hosts may benefit from the unused capacity.

Nevertheless, the performance anomaly analyzed in this paper should be taken into account when we consider the

deployment of 802.11 access points. The question becomes important in the case of hot spots that cover areas with an important number of hosts. When access points are located so that some mobile hosts are far away and use a smaller bit rate, this will result in the performance degradation perceived by all hosts.

VIII. Acknowledgments

References

[1] ANSI/IEEE, "802.11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," 2000.
[2] F. Cali, M. Conti, and E. Gregori, "IEEE 802.11 Wireless LAN: Capacity Analysis and Protocol Enhancement," in *INFOCOM*, 1998.
[3] C. Koksal et al., "An analysis of short-term fairness in wireless media access protocols," in *Proceedings of ACM SIGMETRICS*, 2000.
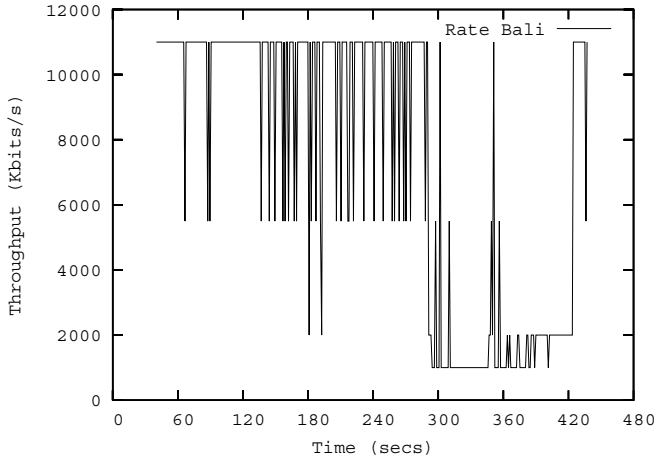
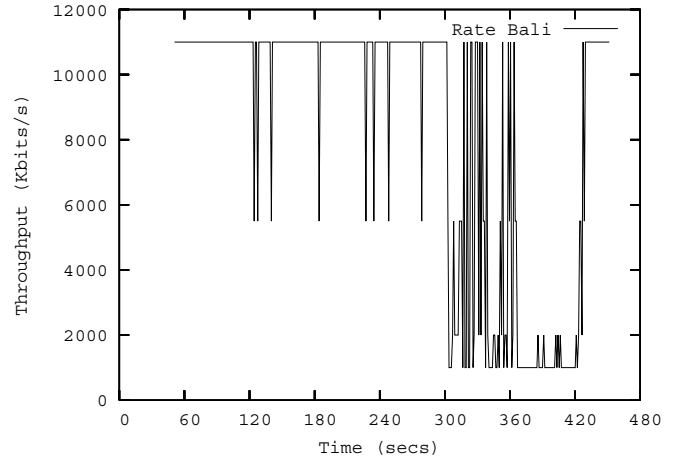Fig. 12.   Evolution of the Bali's bit rate in time
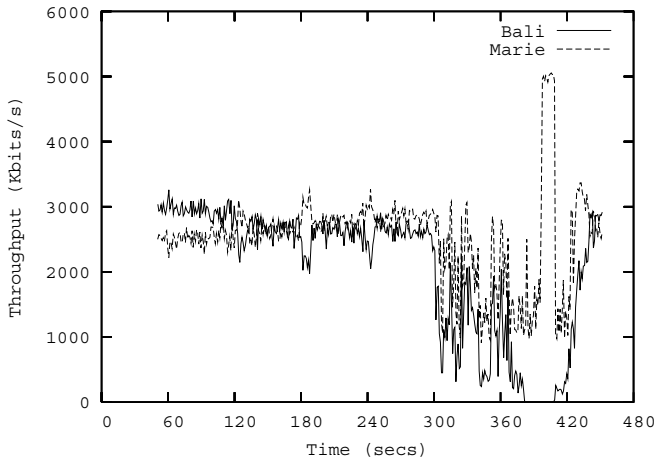

Fig. 14.   Evolution of the Bali's bit rate in time

Fig. 13.   Measured throughput in Mb/s for two hosts (one mobile), UDP traffic

[4] "Netperf: a Network Performance Benchmark," 15 Feb. 1995, hewlett-Packard, Information Networks Division. [Online]. Available: http://www.netperf.org/netperf/training/netperf.ps
[5] B. P. Crow et al., "Investigation of the IEEE 802.11 Medium Access Control (MAC)," in *Proceedings of INFOCOM 1997*, 1997.
[6] M. Natkaniec et al., "Analysis of Backoff mechanism in IEEE 802.11 standard," in *Proceedings of ISCC 2000*, 2000.
[7] G. Bianchi, "Performance Analysis of the IEEE 802.11 Distributed Coordination Function," *JSAC Wireless series*, vol. 18, no. 3, 2000.

APPENDIX

A. Estimate of t_{jam} in the UDP case

We assume the configuration described earlier with $N-1$ fast hosts and a slow one. The duration of a collision depends on the type of hosts involved in the collision. The probability P_{s} of having a packet sent at the lower rate r involved in the collision can be computed as the ratio between the number of host pairs that contain the slow host and the total number of pairs that can be formed in the set of all hosts:

$$P_{\mathrm{s}} = \frac{N-1}{\frac{N(N-1)}{2}} = \frac{2}{N}$$

. Thus, the average time of channel occupancy by the frames subject to a collision is

$$t_{\mathrm{jam}} = \frac{2}{N}\, T_{\mathrm{s}} + (1 - \frac{2}{N})\, T_{\mathrm{f}}.$$

B. Estimate of t_{jam} in the TCP case

In the case of TCP traffic, the Access Point sends TCP ACKs at a variable rate that depends on the rate of the host which is the destination of the ACK. For $N-1$ fast hosts and a slow one, the Access Point sends $(N-1)/2$ more packets than any other host.

Collisions may occur either between TCP data segments and TCP ACKs, or between two TCP data segments. Each frame may be sent at the lower rate r or the higher rate R. In the case of a collision between a TCP data segment and a TCP ACK, the collision lasts the time needed for the transmission of the data segment. The number of host pairs that can be formed without the Access Point is $(N-1)(N-2)/2$ and the number of hosts pairs that include it is $N-1$. As it sends $(N-1)/2$ more packets, it contributes to the collision count for $(N-1)^2/2$.

There are $N-2$ pairs that may include the host transmitting at the lower rate. Only one pair includes the Access Point and the slow host, but it contributes to the collision count for $(N-1)/2$.

Finally the probability P_{s} of having a packet sent at the lower rate r involved in the collision is:

$$P_{\mathrm{s}} = \frac{3\,N-5}{(N-1)(2\,N-3)}.$$

And the average time of channel occupancy by frames subject to a collision is:

$$t_{\mathrm{jam}} = P_{\mathrm{s}}\, T_{\mathrm{s}} + (1 - P_{\mathrm{s}})\, T_{\mathrm{f}}.$$