

Fairness and Its Impact on Delay in 802.11 Networks

Gilles Berger-Sabbatel[†], Andrzej Duda[†], Olivier Gaudoin[‡], Martin Heusse[†], Franck Rousseau[†]

[†] LSR-IMAG Laboratory [‡]LMC-IMAG Laboratory
Grenoble, France

{Gilles.Berger-Sabbatel, Andrzej.Duda, Olivier.Gaudoin, Martin.Heusse, Franck.Rousseau}@imag.fr

Abstract— We analyze in this paper the fairness of the 802.11 DCF (Distributed Coordination Function) access method. Short-term fairness is an important property of a MAC layer for obtaining short delays. We show that contrary to the common wisdom, a 802.11 cell with two hosts does not exhibit short-term unfairness. Many papers considered 802.11 as short-term unfair by referring to a study of the Wavelan CSMA/CA access method [1]. The confusion comes from the extrapolation of its results to 802.11. Actually, these two access methods are very different: the Wavelan CSMA/CA access method performs exponential backoff when the channel is sensed busy, whereas 802.11 does the same only after a collision.

We propose a new fairness index: the number of inter-transmissions that other hosts may perform between two transmissions of a given host. By means of this index we analyze the fairness of 802.11 for the case of two hosts and derive the probability distribution of the number of inter-transmissions. Our results show that even on the short term time scale the 802.11 DCF access method provides fairness to competing hosts. When compared with Slotted ALOHA, a multiple access randomized protocol with good fairness properties, 802.11 presents even better fairness. To validate the model, we compare the analytical results with experimental histograms obtained via simulations and measurements.

I. INTRODUCTION

A MAC layer for a shared link medium should achieve several objectives: high efficiency, low latency, and fairness. This last objective can be characterized in two different manners:

- **Long-term fairness.** It is fairness observed over long time periods (corresponding for instance to the transmission of thousand packets). A MAC layer can be considered as long-term fair if the probability of successful access to the channel observed on a long term converges to $1/N$ for N competing hosts.
- **Short-term fairness.** This property is much stronger: the access to the channel should be fair over short time periods (lasting for an order of magnitude of 10 ms or approximately ten packets in 802.11b). A MAC layer can be long-term fair, but short-term unfair: one host may capture the channel over short time intervals.

Short-term fairness is extremely important for attaining the low latency objective. If a MAC layer presents short-term fairness, each host can expect to access the channel during short intervals, which in turn results in short delays. Previous studies

have shown that a MAC layer presenting long-term fairness and short-term unfairness generate serious performance impairments at upper layers, for example performance of TCP traffic may be severely degraded, because of delayed data and ACK segments, which result in retransmissions or decreasing congestion windows. The effect of short-term unfairness is even more penalizing for time sensitive flows that require short delays and small jitter.

The problem of fairness has been largely analyzed since the first multiple access methods for shared wired media such as Ethernet (802.3) and received a considerable attention for wireless LANs such as the first Wavelan products [1] and the IEEE 802.11 standard. The common wisdom concerning the fairness of 802.11 in the infrastructure mode states that its CSMA/CA access method provides long-term fairness and short-term unfairness (cf. Section II). Based on this knowledge, 802.11 is usually considered as unsuitable for time-sensitive traffic.

In this paper we analyze the fairness of the 802.11 DCF (Distributed Coordination Function) access method in the infrastructure mode [2] and show that even on the short term time scale its access method provides good fairness in the case of two hosts. We explain that the common wisdom about the short-term unfairness comes from the confusion of the access method used in the first Wavelan wireless cards and its modified version used in 802.11. The Wavelan access method was based on CSMA/CA principle in which a host *sensing the channel busy* goes into backoff mode: it selects a backoff time uniformly distributed over the range $\{1, 2, \dots, CW\} \times SLOT$, initially $CW = 32$ and doubles each time the carrier continues to be sensed. When two hosts for example compete for the channel, the host transmitting a frame has a higher probability of gaining the channel for the next transmission, because the other host senses the channel busy and doubles the retry interval. The DCF access method of 802.11 provides a major modification to this scheme: a host goes into backoff mode only when a *collision is detected* and not when it senses the channel busy. Hence, in 802.11, channel capture by one host may only appear after a collision.

We propose a new fairness index: the *number of inter-transmissions* that other hosts may perform between two transmissions of a given host. We derive the probability distribution of the number of inter-transmissions in the case of

two hosts. To validate the model, we compare the analytical results with experimental histograms obtained via simulations and measurements. We use Slotted ALOHA, a multiple access randomized protocol having good fairness properties, as a kind of a fairness gauge: actually, 802.11 demonstrates better fairness properties than Slotted ALOHA.

We also confront our results with a previously defined fairness method: the *sliding window method* that observes the patterns of transmissions and computes the *average Jain fairness index* in a window of an increasing size. We also analyze the fairness of 802.11 by using the sliding window method and the average Jain fairness index. Applied to the measurements and simulations, the index shows that the fairness of 802.11 is pretty good even on the short term time scale. All this evidence shows that the common wisdom concerning the fairness of 802.11 comes from the extrapolation of the results obtained for a different access method, namely the Wavelan CSMA/CA. We also provide experimental evidence showing that short delays can be obtained in 802.11 even if the channel is saturated. This nice property of 802.11 results from its good short-term fairness.

The rest of the paper is organized as follows. We start with the review of the existing work on fairness in wireless local area networks (Section II). Then, we define the notion of fairness (Section III) and analyze the 802.11 DCF access method—we derive the probability distribution of the number of inter-transmissions (Section IV). We compare the analytical results with experimental histograms obtained via simulations and measurements (Section V). Finally, we present some conclusions (Section VI).

II. RELATED WORK

The fairness of 802.11 when all hosts have equal opportunity of using a shared common channel has been largely analyzed in the literature. Koksal et al. analyzed the short-term unfairness of the Wavelan CSMA/CA medium access protocol [1]. They proposed two approaches for evaluating fairness: one based on the sliding window method with the Jain fairness index and the Kullback-Leibler distance, and the other one that uses renewal reward theory based on Markov chain modeling. The authors used Slotted ALOHA as an example of an access method with better fairness, but with much higher collision probability. This paper clearly identifies the short-term unfairness problem of an access method in which hosts perform exponential backoff whenever the channel is sensed busy.

Since this paper, many authors have stated that 802.11 suffers from short-term unfairness and referenced it as the paper that proves the short-term unfairness of 802.11 [3]–[5]. However, they have not realized that the access method of 802.11 has changed with respect to that of the Wavelan cards: in 802.11 standard [2] exponential backoff is only applied after a collision. This misleading common wisdom has emerged from the confusion of these two different access methods.

The confusion of the access methods in Wavelan and 802.11 dies hard: recently, some authors have described the 802.11

access method as based on the same principle as in the Wavelan cards, i.e. exponential backoff applied when the channel is sensed busy [6].

III. FAIRNESS

Our goal is to study the intrinsic fairness properties of the 802.11 DCF access method, so we concentrate on the homogeneous case in which all hosts benefit from similar transmission conditions: no host is disadvantaged by its signal quality, traffic pattern, or spatial position. This means that we do not take into account the problem of hidden or exposed terminals and we do not consider the RTS/CTS extension. In particular, we do not deal with the problems of unfairness due to different spatial host positions [6], [7]. Once we got insight into the intrinsic fairness of the 802.11 MAC layer, we can investigate the influence of other factors such as different spatial positions or traffic patterns.

In general, the fairness of a MAC layer can be defined in a similar way to Fair Queueing: assume N hosts and let $W_i(t_1, t_2)$, $i \in \{0, 1, 2, \dots, N\}$, be the amount of bandwidth allocated to host i in time interval $[t_1, t_2]$. The fair allocation requires that $W_i(t_1, t_2) = W_j(t_1, t_2)$, $i, j \in \{0, 1, 2, \dots, N\}$, regardless of how small the interval $[t_1, t_2]$ is.

We consider the case of greedy hosts (they always have a frame to send) that send frames of equal size. In this case, it is sufficient to only take into account the number of transmissions: the fair allocation needs to guarantee that over any time interval, each host transmits the same number of frames.

To evaluate fairness we will use two methods. The first one uses the *number of inter-transmissions* that other hosts may perform between two transmissions of a given host and the second one computes the *average Jain fairness index* in a window of an increasing size.

A. Number of inter-transmissions

Consider the case of $N = 2$: two hosts A and B share a common channel. To characterize fairness we take the point of view of host B and investigate K , the number of inter-transmissions that host A may perform between two transmissions of host B :

- $K = 0$ means that after a successful transmission of B , the next transmission will be done once again by B ,
- $K = 1$ means that A will transmit once and then the next transmission will be done by host B ,
- $K = 2$ means that A will transmit twice and then the next transmission will be done by host B , and so on.

Consider the following example pattern of transmissions: BBAAABABAAB—random variable K takes the following values: 0, 3, 1, 2.

In a deterministic channel sharing system such as TDMA, the distribution of K will simply be $P(K = k) = 0$ for $k = 0$ and $P(K = k) = 1$ for $k = 1$, meaning that both hosts perfectly alternate transmissions. The mean number of inter-transmissions in TDMA is $E(K) = 1$.

An example of a randomized access protocol that presents good fairness properties is Slotted ALOHA—it has been previously used for fairness comparisons [1]. Time in Slotted ALOHA is divided into slots, each access is independent from the previous one and when a collision occurs, a transmitting host waits a random number of slots distributed geometrically. If we ignore collisions, Slotted ALOHA with two hosts can be modeled as a simple Markov chain with two states. In this case the number of inter-transmissions is geometrically distributed with the parameter $1/2$ (this expression only holds for two hosts):

$$P(K = k) = \frac{1}{2^{k+1}}, \quad k \in \{0, 1, 2, \dots\} \quad (1)$$

Note that for Slotted ALOHA $P(K = 0) = 1/2$, so that each host has equal probability of accessing the channel at any time and the mean number of inter-transmissions is $E(K) = 1$, which is the same as for TDMA.

We can generalize the number of inter-transmissions to a larger number of hosts: we choose one host and observe how many times other hosts transmit frames before another transmission by the chosen host. Consider for example the following sequence of transmissions by five hosts: BAACEDCAB. The number of inter-transmissions observed from the point of view of B is 7.

Observing each outcome of random variable K gives us information on short-term fairness, whereas its distribution and moments convey indication about both short-term and long-term fairness. We can notice that large values of K mean lower fairness, because other hosts may capture the channel for several successive transmissions. Similarly, too small values of K also indicate lower fairness, as the chosen host captures the channel in an unfair way.

More precisely, the distribution of inter-transmissions $P(K = k)$ enables us to quantify fairness by means of:

- *capture probability*: $P(K = 0)$ characterizes the chances of a host to capture the channel. If $P(K = 0) = 1/N$, then all hosts have equal probability of accessing the channel.
- *mean number of inter-transmissions*: $E(K) = 0$ means that one host monopolizes the channel and $E(K) = N - 1$ means that on the average each host performs one transmission at a time, the situation that can be considered as fair ($N - 1$ is the number of inter-transmissions in TDMA). Values $E(K) < N - 1$ indicate a shorter tail of the distribution and better fairness, whereas $E(K) > N - 1$ indicates increased unfairness.
- *100th percentile*: characterizes the tail of the distribution, it is the largest l for which $\sum_{k=0}^l P(K = k) < q$, $0 < q < 1$. For instance the 95th percentile tells us that in 95% of cases the number of inter-transmissions will be less than the 95th percentile. Putting it another way, in only 5% of cases a host should wait more than the 95th percentile transmissions before the next access to the channel.

We can notice that the number of inter-transmissions is directly related to delays perceived by a host competing with other hosts for the channel access: when a host experiences large values of K , it also suffers from large delays, because it has to wait for the channel access while other hosts transmit several frames.

B. Sliding window method with the Jain fairness index

The sliding window method considers the patterns of transmissions and computes the average Jain fairness index in a window of an increasing size [1]. It is defined as follows: let γ_i be the fraction of transmissions performed by host i during window w ; the fairness index is the following:

$$F_J(w) = \frac{(\sum_{i=1}^N \gamma_i)^2}{N \sum_{i=1}^N \gamma_i^2} \quad (2)$$

Perfect fairness is achieved for $F_J(w) = 1$ and perfect unfairness for $F_J(w) = 1/N$.

The definition of window w also should take into account N , the number of competing hosts. We propose to normalize the window size with respect to the number of hosts and compute the Jain index for the window sizes which are multiples of N , because only in this case computing the Jain index makes sense. We call m such that $w = m \times N$, $m = 0, 1, 2, \dots$, a *normalized window size*. The Jain index will be computed as $F_J(m)$.

Both indices have the nice property of being able to capture the short-term as well as the long-term fairness.

IV. ANALYSIS OF THE FAIRNESS IN IEEE 802.11

We review below the mandatory DCF access methods of IEEE 802.11 [2] (we do not discuss the optional Point Coordination Function (PCF)) and the Wavelan CSMA/CA.

A. IEEE 802.11 DCF access method

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 Interframe 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 Interframe Space). If the 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 the DIFS period augmented with a random interval of time.

If there are multiple hosts attempting to transmit, the channel may be sensed busy and in this case hosts enter the collision avoidance phase: a host waits for a random interval distributed uniformly over $\{0, 1, 2, \dots, CW - 1\} \times SLOT$ (discrete distribution). The congestion window CW varies between $CW_{\min} = 32$ and $CW_{\max} = 1024$, the value of $SLOT$ is $20 \mu s$ (these parameters are for 802.11b). The host that chooses the smallest interval starts transmitting and the others freeze their intervals until the transmission is over. When hosts choose the same value of the random interval, they

will try to transmit at the same slot, which results in a collision detected by the missing ACK frame (only the transmitting hosts may detect a collision). Each time a host happens to collide, it executes the exponential backoff algorithm – it doubles CW up to CW_{\max} . CW is set back to CW_{\min} by each host after their first successful transmission.

B. Number of inter-transmissions

Consider the 802.11 access method with two greedy hosts A and B . We analyze the number of inter-transmissions in this case.

The order in which the hosts transmit their frames depends on the random values chosen in the contention window: we assume that a , the value chosen by A and b , the value chosen by B are independent random variables distributed uniformly over the interval $\{0, 1, 2, \dots, CW - 1\} \times SLOT$ (discrete distribution). When host A chooses value a_1 , such that $a_1 < b$, it performs one transmission. Then it chooses next value a_2 and if $a_1 + a_2 < b$, it transmits once again. In general A succeeds k times while $\sum_{i=1}^k a_i < b$ and B takes over as soon as $\sum_{i=1}^{k+1} a_i \geq b$.

When both hosts choose the same value, they will try to transmit at the same slot and collide. The collision will be detected by both hosts that apply the exponential backoff by doubling the contention window. The same stochastic process will start again, but with the doubled value of CW .

We derive below the distribution of K , the number of inter-transmissions as seen by host B by making the following assumption:

- we approximate the discrete uniform distribution by a continuous one. This means that collisions do not occur (the probability that both hosts choose the same value is infinitesimally small).

The distribution of K can be formally written as follows:

$$P(K = k) = P\left(\sum_{i=1}^k a_i < b \text{ and } \sum_{i=1}^{k+1} a_i \geq b\right),$$

where a_1, \dots, a_{k+1} , and b are independent random variables with the same uniform distribution over the interval $[0, CW - 1] \times SLOT$. This distribution is continuous and can be used for approximating the discrete one for large values of CW : the discrete distribution converges to the continuous one when $CW \rightarrow \infty$.

We divide each variable by $(CW - 1) \times SLOT$ to obtain:

$$P(K = k) = P\left(\sum_{i=1}^k u_i < v \text{ and } \sum_{i=1}^{k+1} u_i \geq v\right),$$

where u_1, \dots, u_{k+1} , and v are independent random variables with the same uniform distribution over $[0, 1]$. The distribution of K can be rewritten as:

$$P(K = k) = P(0 < V - W \leq U),$$

where U, V, W are independent random variables: $U = u_{k+1}$ and V are uniformly distributed over $[0, 1]$ with the probability

density function (pdf) $f_U(u) = \mathbf{1}_{[0,1]}(u)$. The distribution of $W = \sum_{i=1}^k u_i$ is known as the Irwin-Hall distribution and its pdf is the following [8]:

$$f_W(w) = \frac{1}{(k-1)!} \sum_{j=0}^k (-1)^j C_k^j [\sup(0, w-j)]^{k-1} \mathbf{1}_{[0,k]}(w).$$

Thus the distribution of K can be expressed as:

$$P(K = k) = \int \int \int_{0 < v-w \leq u} f_V(v) f_W(w) f_U(u) du dv dw.$$

As w is less than 1, $\sup(0, w-j) = 0$ for $j \geq 1$ and the integral becomes:

$$P(K = k) = \frac{1}{(k-1)!} \int_{v=0}^1 \int_{w=0}^v \int_{u=v-w}^1 w^{k-1} du dv dw.$$

Finally, we obtain the following result:

$$P(K = k) = \frac{k+1}{(k+2)!}, \quad k \in \{0, 1, 2, \dots\}. \quad (3)$$

This elegant formula has several interesting properties:

- *Access probability.* We can notice that $P(K = 0) = 1/2$, which shows that each host has the same probability of accessing the channel. $P(K = 0)$ is the same as in Slotted ALOHA (Eq. 1).
- *Distribution tail.* Compared to the distribution of Slotted ALOHA the distribution of 802.11 has a shorter tail, because of the factorial function (see also figures in the next section). This means that Slotted ALOHA is more unfair than 802.11: it is more probable for one host in Slotted ALOHA to perform several successive transmissions than in 802.11.
- *Mean.* The mean of K is given by

$$E(K) = \sum_{k=0}^{+\infty} \frac{k(k+1)}{(k+2)!}.$$

Let

$$g(x) = \sum_{k=0}^{+\infty} \frac{x^{k+1}}{(k+2)!} = \frac{e^x - 1}{x} - 1.$$

It is easy to see that $E(K) = g''(1)$, so the mean number of inter-transmissions is:

$$E(K) = e - 2 = 0.718\dots$$

This result also shows that the fairness of 802.11 is better than that of Slotted ALOHA: the mean number of inter-transmissions in 802.11 is lower than in Slotted ALOHA.

- *Independence of CW .* Surprisingly, the distribution of K is independent of the contention window size CW ! This means that changing its size, in particular by the exponential backoff algorithm, does not influence fairness.

At the beginning of this section, we have made the assumption whose impact on the accuracy of the approximation needs

to be evaluated: we have approximated the discrete uniform distribution by a continuous one, which means that we do not take into account collisions.

The impact of this assumption depends on the collision probability. Our previous models and measurements show that for $CW = 32$ and two hosts the collision probability is around 3% [9], which gives the order of magnitude of the expected precision. The difference between our approximation and the standard behavior of 802.11 can be shown up when a collision occurs: both hosts double their congestion windows, but then the host that succeeds the next transmission will lower its window back down to 32, which favors its following transmissions.

No simplifying assumptions were made in our simulations and the next section will show that the difference between the analytical results and simulation is fairly small, which confirms a minor impact of collisions on the approximation.

V. EXPERIMENTAL RESULTS

To validate our results and investigate furthermore fairness, we have developed a simulator and set up an experimental platform to measure the fairness indices and delays in a 802.11 cell.

The simulator implements the algorithm of the 802.11b for a chosen number of greedy hosts: all the parameters of the simulation such as CW_{min} , CW_{max} , $SLOT$ have the values defined in the 802.11b standard. We can modify the behavior of the algorithm if needed, for example set different values of CW or disable exponential backoff.

We have also set up an experimental platform to measure the fairness indices of 802.11. We use notebooks with two network interfaces, one wired 100 Mb/s Ethernet and a 802.11b wireless card. The notebooks run Linux RedHat 8.0 (kernel 2.4.20) with 802.11b cards based on the same chipset (Lucent Orinoco and Compaq WL 110). The notebooks use the wireless cards in the infrastructure mode—an access point is connected to the wired part of the network.

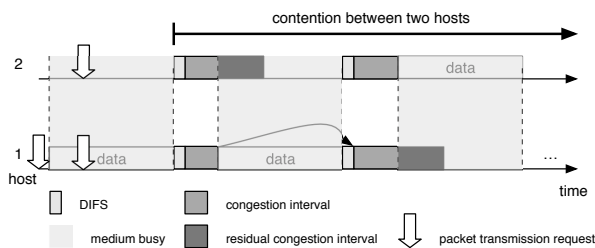


Fig. 1. Principle of measurements

The principle of measurements is the following: at the beginning of a session, one notebook starts transmitting a burst of several UDP packets to the access point. The other notebook tries to send a frame and observes how it inserts in the burst. In this way we can collect statistics for the histogram of the number of inter-transmissions. Our measuring tool synchronizes the notebooks at the beginning of the session

by sending a multicast packet on the wired network. Figure 1 illustrates the principle of measurements.

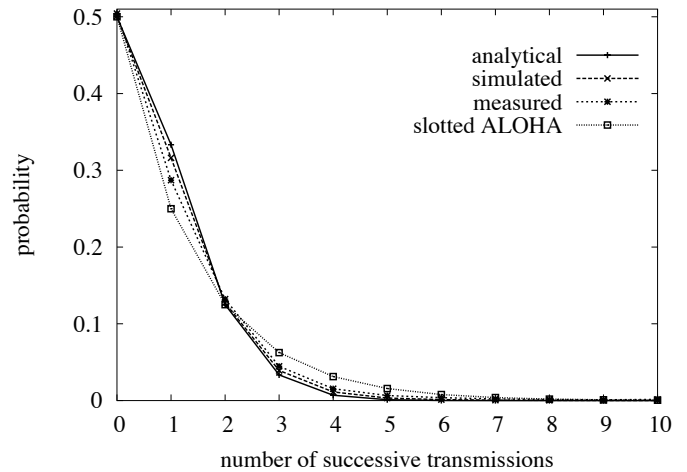


Fig. 2. Distributions of the number of inter-transmissions in 802.11b and Slotted ALOHA, $N = 2$

A. Number of inter-transmissions

Figure 2 presents the analytical, simulated, and measured distributions of the number of inter-transmissions in 802.11b with two competing hosts. We can see that all three distributions are close to each other, the analytical distribution slightly overestimating the other values for $K = 1$ and underestimating for $K > 2$. The figure also compares the distribution of the number of inter-transmissions in 802.11b with Slotted ALOHA. We can see that the distribution for Slotted ALOHA has a longer tail: the probability to have long runs of successive transmissions is higher.

TABLE I
MEAN NUMBER OF INTER-TRANSMISSIONS, $N = 2$

Case	$E(K)$
analytical Slotted ALOHA	1.0
analytical 802.11	0.718
simulated standard 802.11	0.768
simulated 802.11 (constant $CW = 32$)	0.747
simulated 802.11 (constant $CW = 1024$)	0.719
simulated 802.11 (constant $CW = 4096$)	0.718

Table I presents the mean number of inter-transmissions for two hosts. We can see that when we increase the congestion window, which means that the continuous uniform distribution better approximates the discrete one, the mean converges to the analytical value.

B. Sliding window method with the Jain fairness index

We have also collected traces in the 802.11b network with two hosts using `netperf` that generates two competing UDP flows. From the traces, we have computed the Jain fairness index over sliding windows of size $w = m \times N$, $m = 0, 1, 2, \dots$. For Slotted ALOHA, we have used simulation to generate traces.

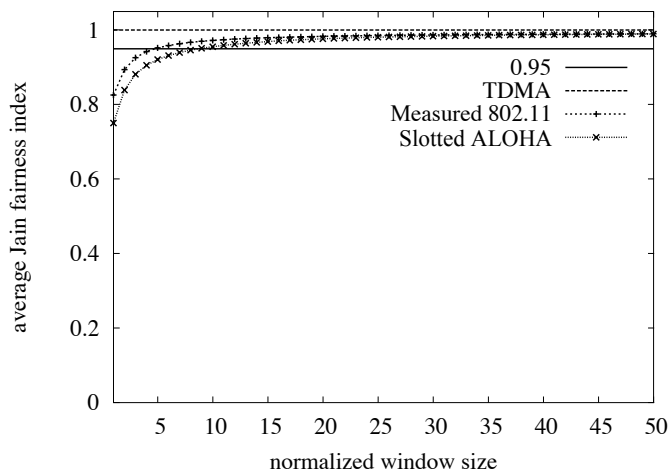


Fig. 3. Jain fairness index for normalized window size

Figure 3 shows the Jain fairness index measured in a 802.11 cell with two hosts in function of the normalized window size. The index of 1 (as in TDMA) represents perfect fairness. It can be seen that the threshold value of 0.95 is quickly attained for the normalized window size of 5. Recall that the same value was attained in a similar experiment for the Wavelan cards at the window size of 500, equivalent to 1000 in the original sliding window method [1]. This shows how different is the fairness of the 802.11 and the Wavelan cards. We can also compare 802.11 with Slotted Aloha—802.11 presents better short-term fairness.

C. Delay

We consider two hosts in a 802.11b cell that generate traffic of different classes designated according to the DiffServ model: one host generates high priority EF traffic and the other one lower priority AF traffic (we used this notation only to distinguish between different types of traffic—no special scheduling is implemented on hosts). To avoid interferences on the wireless channel, we measure the round trip time (RTT) in a configuration in which a host sends a packet over 802.11b and the reply returns via another interface (100 Mbit/s Ethernet).

In the experiment, one host sends EF traffic of a given packet rate whereas the other one tries to increase its AF traffic as much as possible starting from 256 Kbit/s to 10 Mbit/s in steps of 256 Kbit/s. The results are presented in function of the *offered load*, which is the sum of the EF and AF traffic in Kbit/s. The packet size corresponds to the UDP payload size.

Figure 4 illustrates the saturation of the network: the measured throughput attains some maximum value that strongly depends on the AF packet size. The cut-off point of the curves indicates the limiting offered load for which the cell enters saturation.

Figure 5 presents the RTT of the EF class transmitting at 128 Kbit/s, twice the rate of the previous figure. We can see that it remains small (under 6 ms) even if the cell is already




Fig. 4. Total throughput measured for the increasing offered load, constant 64 Kb/s EF traffic

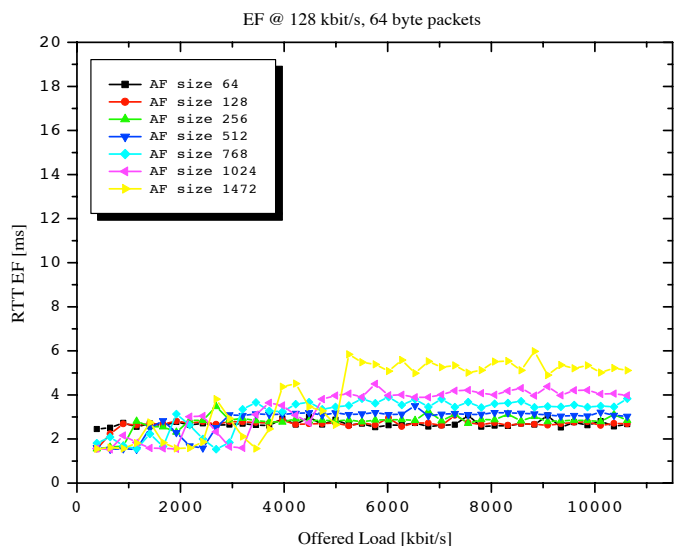


Fig. 5. RTT of the EF class for the increasing offered load, constant 128 Kb/s EF traffic.

saturated (the offered load increasing to 10 Mbit/s).

These results show that even if the channel is saturated by the AF traffic, it is still possible for the EF class to benefit from low delays provided that its packet rate remains lower than some limiting value. The reason of this nice effect is the good short-term fairness of the 802.11 access method: the channel access probability is equal for both hosts and even if the AF host tries to transmit as much as possible, the host generating the EF class benefits from the possibility of transmissions spaced by short inter-transmissions.

D. Case of several contending hosts

To complete our analysis, we present in this section some results for the case of $N > 2$ (for more details see [10]). As the analysis developed in Section IV-B applies only to two

contending hosts, we only report on experimental data.

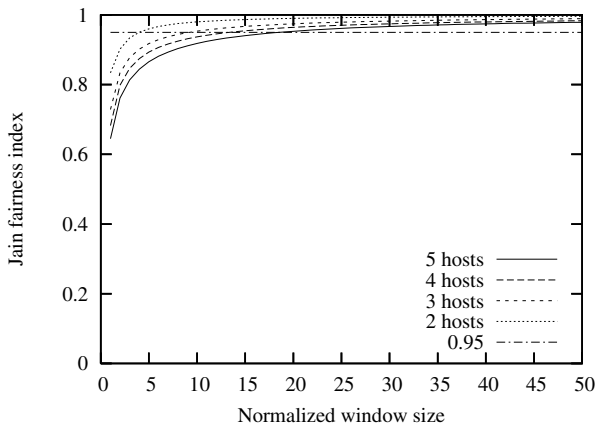


Fig. 6. Measured Jain fairness index of 802.11 for normalized window size

Figure 6 shows the measured normalized Jain fairness index for several hosts. The fairness index gets worse as the number of hosts increases, which is mainly due to the fact that after a collision, the first host that successfully transmits a frame is favored compared to the others. Nevertheless, the fairness index remains acceptable compared to the one measured for the first Wavelan cards.

TABLE II
WINDOW SIZE TO ACHIEVE 0.95 JAIN FAIRNESS INDEX

Number of hosts N	2	3	4
Wavelan	475	83	112
802.11	4	9	13

Table II compares the window size required to achieve the threshold of 0.95 for the Wavelan cards and 802.11. To obtain the normalized window size, we have divided the results reported by Koksals et al. [1] by the number of hosts (only the values of the Jain fairness index for $N = 2, 3$, and 4 hosts have been given in this paper).

VI. CONCLUSION

In this paper we have analyzed the short-term fairness of the 802.11 DCF access method, the property that enables obtaining low delays. The problem considered in this paper arose when we analyzed the performance of 802.11 and we tried to characterize the delay. Many papers considered 802.11 as short-term unfair by referring to the study of the Wavelan CSMA/CA access method [1]. When citing the results of this paper, several authors extrapolated from the Wavelan CSMA/CA to 802.11 without noticing that their respective access methods are very different. Our paper shows that contrary to this common wisdom, the 802.11 access method in a cell with two hosts does not exhibit short-term unfairness.

We have proposed a new fairness index: the number of inter-transmissions that other hosts may perform between two transmissions of a given host. We have derived its probability distribution for a 802.11 cell with two hosts. The distribution

has an elegant closed form that results from approximating a discrete uniform distribution by a continuous one. When compared with Slotted ALOHA considered as a randomized access protocol with good fairness properties, 802.11 shows even better fairness. Such a good behavior comes from the fact that hosts in 802.11 use their residual congestion intervals—when a host chooses a long interval, it will wait during one or several turns, but then it will eventually succeed because its congestion interval becomes smaller and smaller.

Unlike many papers, we confront the analytical results with measurements and simulation. Our experimental histograms validate the analytical results. We have also used the sliding window method to compute the average Jain fairness index over traces gathered via measurements and simulation. All this evidence shows that a 802.11 cell with two hosts presents much better fairness than the Wavelan wireless cards at the origin of the common wisdom on the short-term unfairness.

The presented results also apply to the 802.11a and 802.11g standards, because they have the same MAC access method as 802.11b, even if some parameters such as CW_{min} and $SLOT$ are modified.

Finally, we show that for an increasing number of hosts short-term fairness becomes a little bit worse, however it still remains acceptable. In fact, the collision rate for several hosts stays low (around 10 % for $N = 5$), so that the exponential backoff does not impact fairness in an important way. However, even if our analysis and experimental results show this good feature of 802.11, we believe that there is an open issue of designing a better MAC layer that optimizes the throughput and short-term fairness at the same time.

REFERENCES

- [1] C.E. Koksals, H. Kassab, and H. Balakrishnan, "An analysis of short-term fairness in wireless media access protocols," in *Proceedings of ACM SIGMETRICS*, 2000, extended version available at <http://nms.lcs.mit.edu/papers>.
- [2] ANSI/IEEE, "802.11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," 2000.
- [3] C. L. Barrett, M. V. Marathe, D. C. Engelhart, and A. Sivasubramanian, "Analyzing the short-term fairness of IEEE 802.11 in wireless multi-hop radio networks," in *10th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunications Systems (MASCOTS'02)*, Fort Worth, Texas, October 2002.
- [4] N. H. Vaidya, P. Bahl, and S. Gupta, "Distributed fair scheduling in a wireless LAN," in *Sixth Annual International Conference on Mobile Computing and Networking*, Boston, August 2000.
- [5] Y. Kwon, Y. Fang, and H. Latchman, "A novel mac protocol with fast collision resolution for wireless lans," in *Proceedings of IEEE INFOCOM 2003*, San Francisco, USA, March-April 2003.
- [6] Z. Fang, B. Bensaou, and Y. Wang, "Performance evaluation of a fair backoff algorithm for IEEE 802.11 DCFWMA," in *Proceeding of MOBIHOC'02*, Lausanne, Switzerland, June 2002.
- [7] V. Barghavan, A. Demers, S. Shenker, and L. Zhang, "MACAW: a media access protocol for wireless LAN's," in *Proceedings of ACM SIGCOMM*, August 1994.
- [8] N.L. Johnson, S. Kotz, and N. Balakrishnan, *Continuous univariate distributions*, Wiley, 1995.
- [9] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *Proceedings of IEEE INFOCOM 2003*, San Francisco, USA, March-April 2003.
- [10] G. Berger-Sabbatel, A. Duda, M. Heusse, and F. Rousseau, "Short-term fairness of 802.11 networks with several hosts," in *Proceedings of the Sixth IFIP IEEE International Conference on Mobile and Wireless Communication Networks*, Paris, France, October 2004.