SHORT-TERM FAIRNESS
OF 802.11 NETWORKS
WITH SEVERAL HOSTS

Gilles Berger-Sabbatel, Andrzej Duda, Martin Heusse,
and Franck Rousseau
LSR-IMAG Laboratory
Grenoble, France
{gberger, duda, heusse, rousseau}-@imag.fr

Abstract   Previously, we have analyzed the short-term fairness of the 802.11 DCF
(Distributed Coordination Function) access method in the case of a
network with two hosts. In this paper we extend the analysis to an
increased number of hosts. We use two fairness indices. The first one is
the number of inter-transmissions that other hosts may perform between
two transmissions of a given host. The second index is based on the
sliding window method that considers the patterns of transmissions and
computes the average Jain fairness index in a window of an increasing
size. Computed over traces gathered during measurement sessions the
indices show that the fairness of 802.11 is pretty good even on the short
term time scale. We also evaluate the impact of the short-term fairness
on performance by providing the measurements of the delay.

Keywords: Wireless LANs, Fairness of 802.11, CSMA/CA, Network measurements

1.   Introduction

Fairness is related to the ability of a MAC (Medium Access Control)
layer to equitably share a common channel among N contending hosts.
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. A stronger property is *short-term fairness*:
over short time periods, the access to the channel should also be fair. A
MAC layer may 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 obtaining low latency: 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.

In another paper [4] we have analyzed the fairness of the 802.11 DCF (Distributed Coordination Function) access method in the case of a network with two hosts. We have shown that contrary to the common wisdom, 802.11 does not exhibit short-term unfairness. The belief in the short-term unfairness of 802.11 comes from an analysis of the Wavelan CSMA/CA (Carrier Sense Multiple Access/Collision Avoidance) medium access protocol [6]. Its results have been extrapolated to 802.11 without realizing that the access method of 802.11 has changed with respect to that of the Wavelan cards: in the 802.11 standard [1] the exponential backoff is only applied after a collision and not when a host finds the channel busy.

This paper extends the analysis of the fairness in 802.11 networks to an increased number of hosts. Having more than two hosts in a cell changes considerably the access conditions to the radio channel—fairness may be degraded because of an increasing number of collisions. In this case the exponential backoff applied after a collision may have a negative impact on the fairness: a host doubles its maximum congestion window after a collision and has a higher probability of choosing a large interval. During this period other hosts may benefit from transmission opportunity.

We report on the measurements of two fairness indices in a 802.11 cell. The first one, proposed and extensively studied previously [4], is the *number of inter-transmissions* that other hosts may perform between two transmissions of a given host. The second index is based on a largely used method for evaluating fairness: the *sliding window method* that considers the patterns of transmissions and computes the *average Jain fairness index* in a window of an increasing size. Computed over traces gathered during measurement sessions the indices show that the fairness of 802.11 is pretty good even on the short term time scale.

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 2). Then, we define the notion of fairness (Section 3) and present the results of measurements (Section 4). Finally, we present some conclusions (Section 5).

2.   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 [6]. 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 a 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 cited it as the paper that proves the short-term unfairness of 802.11 [3][9][7]. 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 [1] 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 [5].

The belief in the short-term unfairness of 802.11 also comes from some properties of wired CSMA networks, for example several authors have studied the *Ethernet capture effect*, in which a station transmits consecutive packets exclusively for a prolonged period despite of other stations contending for access [8]. However, the 802.11 DCF method eliminates this problem by adopting a CSMA/CA strategy instead of CSMA/CD.

3.   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 [5, 2]. 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 $W_i(t_1, t_2), i \in \{0, 1, 2, ...N\}$, be the amount of bandwidth allocated to host i in time interval $[t_1, t_2]$. The fair allocation requires that $W_i(t_1, t_2) = W_j(t_1, t_2), i, j \in \{0, 1, 2, ...N\}$, 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.

3.1   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 [6]. 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):

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

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

- *100qth percentile:* characterizes the tail of the distribution, it is the smallest 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.

3.2   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 [6]. 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:

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

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 $w = m \times N, m = 0, 1, 2, ...$, 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.

4.   Experimental results

To evaluate the fairness in a 802.11 cell with several hosts, we have set up an experimental platform to measure the fairness indices and delays. We use notebooks running Linux RedHat 8.0 (kernel 2.4.20) with 802.11b cards based on the same chipset (Lucent Orinoco and Compaq WL 110). The results concerning fairness and delays also apply to 802.11a and 802.11g, because they use the same DCF access method as 802.11b with some parameter values modified.

The wireless cards work in the infrastructure mode—an access point is connected to the wired part of the network. Notebooks are greedy in our experiments—they try to send fixed sized UDP packets to the access point as quickly as possible and they always have a packet to send. We gather traces on the destination host: the arrival instant and the source host. Then we compute the fairness indices and delay statistics from the traces.

There are two ways of organizing measurement sessions:

- *Synchronized hosts.* All hosts start at the same instant so that their congestion window counters are reset. To do this, we use a wired 100 Mb/s Ethernet interface in addition to a 802.11b wireless

card. Sending a multicast packet on the wired network allows us to synchronize the hosts at the beginning of a measurement session.

- *Non synchronized hosts.* Hosts start at different instants so when a chosen host performs a successful transmission, other hosts may have lower values of their congestion windows because they use residual time intervals. These potentially shorter intervals mean that they may perform more inter-transmissions than in the previous case. The resulting distribution is different and corresponds to the worst case with respect to the fairness.

The case of synchronized hosts is easier to analyze theoretically, because all hosts begin to operate in a known initial state. However, it does not correspond to realistic working conditions in which hosts operate independently in a non synchronized manner. For the case of a 802.11 cell with two hosts we have done measurements with synchronized hosts to validate the theoretical analysis of the number of inter-transmissions [4]. Such a set up is cumbersome for a larger number of hosts so for this paper we have organized measurement sessions with non synchronized hosts. The results may present worse fairness, but if they are satisfactory, the measurements of synchronized hosts should result in even better fairness.

4.1   Number of inter-transmissions

We briefly recall the results of the analysis for the case of $N = 2$ (the analysis assumes the case of synchronized hosts).

Figure 1 shows 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 approximate 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. Note that $P(K = 0) = 1/2$ which confirms that the access probability is equal for both hosts.

Table 1 presents $E(K)$, the mean number of inter-transmissions for two hosts. It is lower than 1 showing that the fairness of 802.11 is better than that of Slotted ALOHA because the tail of its distribution is shorter.

Figure 2 shows the measured distribution of the number of inter-transmissions for several competing hosts. In the measurement session hosts are not synchronized, so that the measured distribution is slightly different from Figure 1: $P(K = 0) < 1/N$, because the host we observe has less chances to capture the channel than the other hosts—after a
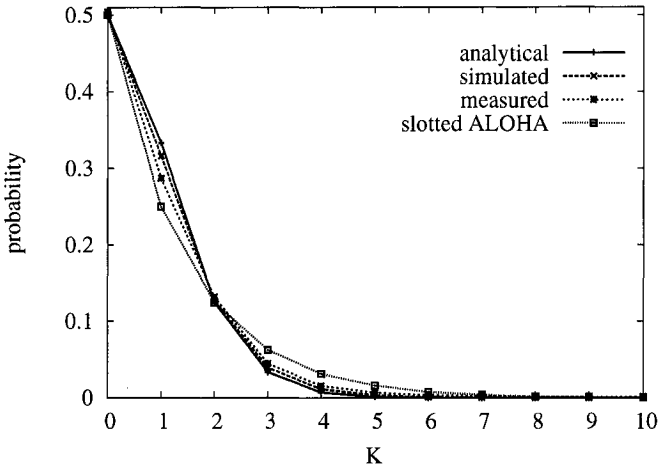
Figure 1. Distribution of the number of inter-transmissions in 802.11b and Slotted ALOHA, $N = 2$, synchronized hosts

Table 1. Mean number of inter-transmissions, $N = 2$

Case	$E(K)$
analytical Slotted ALOHA	1.0
analytical 802.11	0.718
simulated 802.11	0.768
measured 802.11	0.853

successful transmission by the chosen host, the other hosts benefit from smaller congestion window values corresponding to residual time intervals.

We can see from this figure that the distribution decays quickly in function of K. As for an increasing number of hosts, the number of collisions increases, one may expect that fairness degrades. However, we observe that the collisions do not have strong negative impact on the fairness: for $N = 5$ most of the values of K remain lower than 10.

The statistics presented in Table 2 confirm the good fairness properties of the access method: the average number of packets between two transmissions of a given host is equal to the number of hosts with which the host competes for the channel. Our measurements also show that the 95th percentile is fairly low, which confirms the fairness of the access

Figure 2. Measured distribution of the number of inter-transmissions in 802.11b, non synchronized hosts.

Table 2. Statistics of the measured distribution

Number of hosts N	mean	standard deviation	95th percentile
2	1.00	0.97	3
3	2.001	2.3	6
4	3.05	4	10
5	4.001	4.5	11

method. More precisely, when 5 (resp. 2) hosts contend for the channel, a host that has just transmitted a frame has probability of 5% to transmit a frame after the next 11 (resp. 3) successful transmissions.

4.2 Sliding window method with Jain fairness index

We have also computed the Jain fairness index over sliding windows of size $w = m \times N, m = 0, 1, 2, ...$ The index of 1 represents perfect fairness and we use the threshold value of 0.95 to characterize how quickly it is achieved.

Figure 3 shows the Jain fairness index measured in a 802.11 cell with several hosts in function of the normalized window size. It can be seen
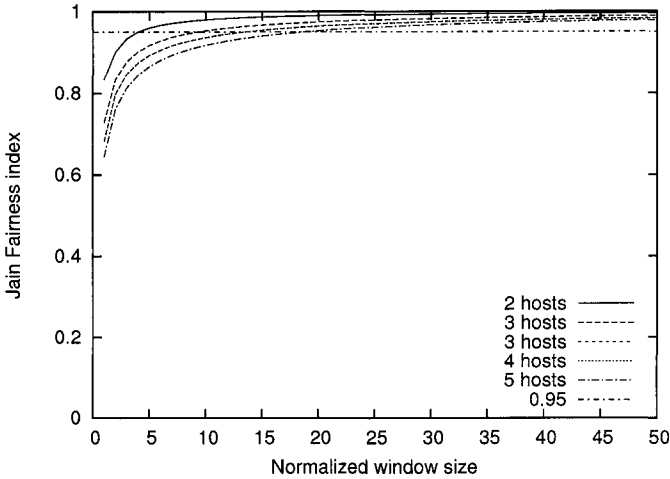
Figure 3.   Jain fairness index of 802.11 for normalized window size

that the threshold value of 0.95 is quickly attained for small values of the normalized window. Note that for an increasing number of hosts, the Jain index begins at a lower value, e.g. for $N = 5$ it starts with the value of 0.2.

Table 3.   Window size to achieve 0.95 fairness index

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

Table 3 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 Koksal et al. [6] 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).

4.3   Delay

To evaluate the impact of the short-term fairness on the delay we have measured delays perceived by one host when competing with other hosts. A delay is the interval between the ends of the last successful transmissions by a given hosts. It is composed of the waiting time while other hosts are transmitting and the actual transmission time of a frame.
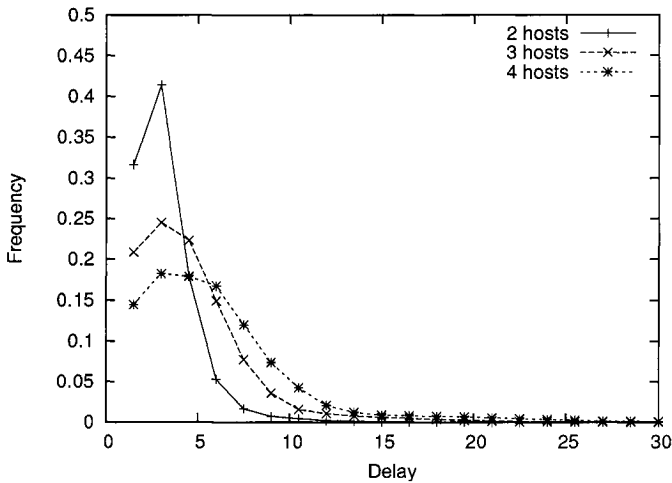
Figure 4. Measured distribution of the delays experienced by one host in milliseconds

Figure 4 presents the measured distribution of delays for an increasing number of hosts $N = 2, 3, 4$. We can see that the delay remains fairly low and the shape of the delay distribution is similar to that of the number of inter-transmissions (cf. Figure 2). This confirms the strong relationship between this fairness index and transmission latency.

5. Conclusion

In this paper we have analyzed the short-term fairness of the 802.11 DCF access method. The results of this paper complete our previous analysis of the 802.11 network with two hosts. They show that contrary to the common wisdom, 802.11 does not exhibit short-term unfairness. This fact has an important influence on transmission latency. As 802.11 presents short-term fairness, each host can access the channel during short intervals, which in turn results in short delays. The results also show that an increased number of collisions in a cell with several hosts does not have strong negative impact on the 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.

Although our measurements were only done for several hosts ($N \leq 5$), the results are also meaningful for a cell with much more hosts that gen-

erate real-life traffic: in our experiments the traffic sources were greedy so the load corresponds to a network with many hosts generating traffic from intermittent sources statistically multiplexed over the shared radio channel.

We work on extending our analysis to take into account other factors such as different spatial positions or traffic patterns. We also investigate improvements to the 802.11 DCF access method able to increase short term fairness and lower the delay.

References

[1] ANSI/IEEE. 802.11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, 2000.

[2] 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.

[3] C. L. Barrett, M. V. Marathe, D. C. Engelhart, and A. Sivasubramaniam. 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] G. Berger-Sabbatel, A. Duda, O. Gaudoin, M. Heusse, and F. Rousseau. Fairness and its Impact on Delay in 802.11 Networks. In *Proceedings of IEEE GLOBECOM 2004*, Dallas, USA, 2004.

[5] Z. Fang, B. Bensaou, and Y. Wang. Performance Evaluation of a Fair Backoff Algorithm for IEEE 802.11 DFWMAC. In *Proceeding of MOBIHOC'02*, Lausanne, Switzerland, June 2002.

[6] C.E. Koksal, 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.

[7] 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.

[8] K.K. Ramakrishnan and H. Yang. The Ethernet Capture Effect: Analysis and Solution. In *IEEE 19th Conference on Local Computer Networks*, 1994.

[9] 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.