

Performance Analysis of 802.11 WLANs under Sporadic Traffic*

M. Garetto and C.-F. Chiasserini

Dipartimento di Elettronica, Politecnico di Torino, Italy

Abstract. We analyze the performance of 802.11 WLANs that employ the Distributed Coordination Function (DCF). We consider contending stations within radio proximity, and investigate the case in which stations operate under non-saturated conditions. Our modelling technique can be used to study several important issues in 802.11 networks, such as the impact of bursty traffic and the system performance in a multirate environment. The accuracy of the analytical results is verified by simulation with *ns-2*.

1 Introduction

We present an analytical model of the 802.11 DCF that differs from previous work [1–6] in addressing all of the following issues: *(i)* it identifies the critical assumptions in the development of analytical models of 802.11 networks, and presents a fairly simple as well as accurate model of the DCF in presence of non-saturated traffic sources; *(ii)* it is general enough to account for different arrival processes and traffic patterns, in particular, it applies to the case of bursty traffic like that produced by the TCP protocol; *(iii)* it evaluates the system performance in a multirate environment; *(iv)* it applies to the case where a station seizing the channel is entitled to transmit a burst of packets, as specified in the IEEE 802.11e draft standard; *(v)* it evaluates several metrics of interest, such as the network throughput, the packet loss probability, the distribution of the MAC queue length at the wireless stations, and the average packet delay.

2 Modelling 802.11 WLANs

First, we briefly present the basic model describing the behavior of saturated sources; then, we describe two different approaches to deal with the more complicated case in which the stations transmission queue may become empty. Please refer to [7] for further details.

2.1 Modeling Saturated Sources

We follow the approach of [1]; in particular, we assume that the channel has three possible states: *(i)* busy channel due to a successful transmission; *(ii)* busy channel due to a collision; *(iii)* idle channel. We build a discrete Markov chain, embedded in the temporal evolution of the channel at the instants of a possible state change. By relying on the fundamental assumption that the state of a station is independent of that of the others, we can reduce to considering the behavior of a single, tagged station. Thus, we propose the simple Markov chain depicted in Figure 1. States labeled with b represent the station with backoff counter equal to 0, i.e., the case where the station actually transmits a frame in the current step. States labeled with B model the station while it decrements its

* This work was supported through the PATTERN project

backoff counter. States have an index in the range $\{0 \dots m\}$ representing the “backoff stage”, where m is the *maximum retry limit* for the frame to be transmitted. We denote the stationary distribution of the Markov chain by $\pi = \{\pi_s\}$, where s is a generic state of the model. The main differences between our representation and previous models for saturated sources are as follows: (i) In case of collision, a station with probability $\alpha_i = 2/W_i$ immediately retransmits the frame in the following step, being W_i the value of the updated contention window; (ii) All states belonging to backoff stage i and having backoff counter greater than one, have been collapsed in a single state B_i so as to reduce the number of states. The side-effect of this simplification is that the number of steps waited while decrementing the backoff counter is modeled as a geometrically distributed random variable, instead of a uniformly distributed variable. However, the transition probabilities $P(B_i, b_i)$'s have been chosen in such a way that the stationary probability of the collapsed states B_i are exactly the same that would result considering a uniformly distributed backoff, as done in [1]. By doing so, all performance metrics derived with our simplified model coincide with those obtained with a more precise model of the backoff counter based on a uniform distribution. Thanks to the particular structure of the Markov chain, the derivation of the stationary probabilities is straightforward. By following [1], we can derive: (i) the probability τ that a station transmits during a time step; (ii) the *conditional collision probability* p ; (iii) the channel state probabilities Π_s , Π_c and Π_σ , i.e., the probabilities that a generic discrete time step is occupied by a successful transmission, a collision, or an idle slot, respectively; (iv) the *aggregate packet throughput* T_P .

2.2 Modeling Non-saturated Sources

We assume that the MAC buffer at the wireless stations receive data packets from the upper layers according to some stationary, external arrival process with rate λ packets/s. For the sake of simplicity, we assume a Poisson arrival process at each station. The model can be easily extended to account for various degrees of traffic burstiness using batch arrivals. We denote by $\Lambda = \lambda n$ the total packet arrival rate at the stations queues. The MAC buffer at each station is assumed to be of finite size K . Packets that cannot be stored in the buffer are immediately discarded upon arrival.

As a first attempt, we incorporate in the Markov model presented above the information about the number of packets currently stored in the queue of the tagged station. The resulting model, named model “A”, comprises the states belonging to the set $\{b_{i,j}, B_{i,j}\}$. The two indexes $0 \leq i \leq m$ and $0 \leq j \leq K$ stand for the backoff stage and the number of packets currently stored in the buffer, respectively. We can derive the average probability τ that a station transmits in an arbitrary time step; then, by assuming that the state of a station is independent of that of the others, we can write the probabilities Π'_s , Π'_c and Π'_σ , that a time step is occupied by a successful transmission, a collision, or an idle slot, respectively, given that the tagged station does not transmit. From the solution of the model, obtained with a fixed point approximation, we obtain many significant performance measures, such as the aggregate throughput T_P , the entire distribution of the number of packets queued at a station, the average MAC queueing delay, the packet loss probability due to buffer overflow, as well as the discard probability due to the maximum retry limit. However, the obtained results (not shown here for the lack of space) show that the model overestimates the aggregate throughput, especially at the knee of the curve. The error that we encounter in the model is essentially

due to the assumption that the state of individual stations are independent. Our conclusion is that *an 802.11 network under non-saturated conditions cannot be correctly analyzed relying on the independence assumption among stations*, but it is necessary to model the joint behavior of the stations, and in particular the evolution of the number of backlogged stations.

Therefore, we extend the description of the state of the system by keeping track of the number of stations having non-empty transmission queue. We describe in detail the behavior of a tagged station using the same states as in *mod A*, but adding to each state of the tagged station an indication of the number k of stations (excluding the tagged one) having at least one packet in the queue. The state space of the resulting model, named model “B”, is the set $\{b_{i,j,k}, B_{i,j,k}\}$, where the newly introduced index k takes values in the range $[0 \cdots n - 1]$. The number k of competing stations during a time step may vary due to the following events: i) one or more of the stations having an empty buffer receive new data to send, thus increasing k ; ii) one of the stations having non-empty buffer successfully transmits a packet leaving an empty queue, thus decreasing k by one. Notice that these two events can occur simultaneously during the successful transmission of a packet.

The number of stations that join the competing set during a time step depends on the duration Δ of the step and on the current value of k . During an interval Δ , a station's queue receives at least one packet with probability $q = 1 - e^{-\lambda\Delta}$. The number of stations that join the set of competing stations is thus distributed according to a binomial distribution with parameters $(q, N - 1 - k)$.

The possible durations Δ of a time step are related to the probabilities that the channel is occupied by a successful transmission, a collision, or an idle slot. To compute such probabilities, we need to specify the probability $\tau(\tilde{C})$ that one of the stations belonging to the competing set transmits in a given time step. This set may include also the tagged station, thus $0 \leq \tilde{C} \leq n$. In a generic state of the model, we have $\tilde{C} = k$ if $j = 0$, whereas $\tilde{C} = k + 1$ if $j > 0$. The probability $\tau(\tilde{C})$, $\tilde{C} \geq 1$, is given by,

$$\tau(\tilde{C}) = \frac{\sum_{i=0}^m \sum_{j=1}^K \pi_{b_{i,j,\tilde{C}-1}}}{\pi_{b_{0,0,\tilde{C}}} + \pi_{B_{0,0,\tilde{C}}} + \sum_{i=0}^m \sum_{j=1}^K (\pi_{b_{i,j,\tilde{C}-1}} + \pi_{B_{i,j,\tilde{C}-1}})}$$

If $\tilde{C} = 0$, we obviously have: $\tau(0) = 0$. Note that, differently from the models previously described in the paper, the probability $\tau(\tilde{C})$ may vary from one state to another, as it depends on \tilde{C} . The tagged station transmits a packet in all states $b_{i,j,k}$, provided that $j \geq 1$. The *conditional collision probability* is given by: $p(\tilde{C}) = 1 - [1 - \tau(\tilde{C})]^{\tilde{C}-1}$; with probability $1 - p(\tilde{C})$, the transmission is successful. In states $b_{i,0,k}$ and $B_{i,j,k}$, the station does not transmit, and the probabilities $\Pi'_s(\tilde{C}, k)$, $\Pi'_c(\tilde{C}, k)$ and $\Pi'_o(\tilde{C}, k)$ that the channel is occupied by a successful transmission, a collision, or an idle slot, respectively, are functions of k and of the total number of competing stations \tilde{C} .

The additional parameter that we need to specify is the probability P_E that, upon successful transmission of a packet, a station other than the tagged one finds itself with an empty buffer, thus decreasing the number of competing stations. This quantity turns out to be the most critical to estimate by our model, as we do not maintain state information about the buffer occupancy at each station. An estimate of P_E can be derived

by considering this probability to be dependent on both the number \tilde{C} of competing stations and the backoff stage i . In particular, a good approximation is obtained assuming that, whenever the tagged station is at backoff stage i , the other competing stations are at a backoff stage h that differs from i at most by one, i.e., $|h - i| \leq 1$. Under this approximation, we obtain the following estimate of P_E :

$$P_E(\tilde{C}, i) = \frac{\sum_{h: |h-i| \leq 1} \pi b_{h,1,\tilde{C}-1}}{\sum_{h: |h-i| \leq 1} \sum_{j=1}^K \pi b_{h,j,\tilde{C}-1}} e^{-\lambda T_S} \quad (1)$$

Figure 2 shows the good agreement among the distributions of the number C of competing stations obtained with model “B” (referred to as *mod B*) and *ns-2*, for various values of Λ , thus proving the accuracy of our approach. The average buffer occupancy (and thus the average queueing delay) is also very well predicted by *mod B*, as shown in Figure 3.




Fig. 1. Markov chain model for saturated sources




Fig. 2. Distributions of C (competing stations) for $n = 10$ stations, $\Lambda = 550$ and 640 packets/s




Fig. 3. Average buffer occupancy vs. the total packet arrival rate

References

1. G. Bianchi, “Performance Analysis of the IEEE 802.11 Distributed Coordination Function,” *IEEE JSAC*, Vol. 18, No. 3, pp. 535–547, Mar. 2000.
2. C. H. Foh and M. Zukerman, “Performance Analysis of the IEEE 802.11 MAC Protocol,” *EWC 2002*, Florence, Italy, pp. 184–190, Feb. 2002.
3. E. M. M. Winands, T. J. J. Denteneer, J. A. C. Resing, and R. Rietman, “A Finite-Source Feedback Queueing Network as a Model for the IEEE 802.11 Distributed Coordination Function,” *5th EWC*, Barcelona, Spain, pp. 551–557, Feb. 2004.
4. R. Litjens, F. Roijers, J. L. van den Berg, R. J. Boucherie, and M. Fleuren, “Performance Analysis of wireless LANs: an Integrated Packet/Flow Level Approach,” *ITC Conference*, Berlin, Germany, Aug. 2003.
5. O. Tickoo and B. Sikdar, “Queueing Analysis and Delay Mitigation in IEEE 802.11 Random Access MAC based Wireless Networks,” *IEEE INFOCOM*, Hong Kong, Mar. 2004.
6. G. R. Cantieni, Q. Ni, C. Barakat, and T. Turletti, “Performance Analysis of Finite Load Sources in 802.11b Multirate Environments,” *INRIA Research Report RR-4881*.
7. M. Garetto and C. F. Chiasserini, “Performance Analysis of the 802.11 Distributed Coordination Function under Sporadic Traffic,” *Tech. Rep.*, Politecnico di Torino, July 2004.