Learning-BEB: Avoiding Collisions in WLAN

J. Barceló, B. Bellalta, C. Cano, M. Oliver
Dept. of Information and Communication Technologies,
Universitat Pompeu Fabra,
Passeig de Circumval.lació 8, 08003 Barcelona
Email: {jaume.barcelo,boris.bellalta,cristina.cano,miquel.oliver}@upf.edu

***Abstract*—Random access protocols have been the mechanism of choice for most WLANs, thanks to their simplicity and distributed nature. Nevertheless, these advantages come at the price of sub-optimal channel utilization because of empty slots and collisions. In previous random access protocols, the stations transmit on the channel without any clue of other stations' intentions to transmit. In this article we provide a framework to study the efficiency of channel access protocols. This framework is used to analyze the efficiency of the Binary Exponential Backoff mechanism and the maximum achievable efficiency that can be obtained from any completely random access protocol. Then we propose Learning-BEB (L-BEB).**

L-BEB is exactly the same as legacy BEB, with one exception: L-BEB chooses a deterministic backoff value after a successful transmission. We call this value the virtual frame size (V). This subtle modification significantly reduces the number of collisions. It can be observed that, as the system runs, the number of collisions is progressively reduced. Thus we conclude that the system learns. Further, if the number of contending stations is equal or lower than V and all stations consecutively successfully transmit, collisions disappear. This collision-free operation is maintained until a new station is activated and joins the contention.

L-BEB pushes the system performance beyond the upper bound inherent to completely-random access mechanisms. Moreover, L-BEB does not introduce any additional complexity to the algorithms currently in use in WLANs. All the claims in the paper are supported by extensive simulation results.

I. Introduction

The radio channel is a broadcast medium and nodes which are in each other interference range should take turns in transmitting. Simultaneous transmissions are called collisions. As a result of a collision, the messages being transmitted might be lost.

The Medium Access Control (MAC) is the function that arbitrates the access to the channel. In wireless networks, the MAC protocols play a key role in maximizing the channel utilization.

Carrier Sense Multiple Access (CSMA) and Time Division Multiple Access (TDMA) are two well known medium access mechanisms for shared medium communication systems. The former relies on stations sensing the medium before transmitting whereas the latter reserves time slices for each active station.

CSMA is simpler to operate. No tight time synchronization is required and the stations simply transmit when they have data ready to be transmitted. Thus they benefit from statistical multiplexing, supporting a larger number of bursty traffic sources. The problems arrive when two stations decide to transmit simultaneously. It may happen that the data from both transmissions is lost. Then a collision resolution mechanism must be activated. To avoid collisions, the stations distributedly execute backoff algorithms that randomly delay the transmission. Because of the random nature of the selection of transmission times, only a fraction of the time is devoted to successful transmissions, while the rest is wasted (either in the form of empty channel, or busy with collisions).

TDMA, on the other side, requires tight time synchronization among the participating stations. Additionally, a prior set-up is required to assign a time-slice (or slot) to the active stations. This set-up causes extra signaling overhead and often requires the presence of a central decision point. After the time slices are assigned, those slices are reserved for a given station. If that station has no data ready to transmit, the channel time is wasted. Conversely, if the station has a large amount of data to transmit, it can only transmit the fraction that fits in the reserved time slice. The rest of the data is buffered for later transmission. The great advantage of TDMA is that it avoids collisions and may achieve high channel efficiency.

Both CSMA and TDMA have advantages and disadvantages. The combination of the advantages of both mechanisms has been a long sought after goal. Sec. II describes related work in the subject. The prior art is characterized by its complexity, which have prevented widespread implementation of the ideas.

Sec. III briefly describes the CSMA mechanisms in IEEE 802.11 [1] which strongly rely on the Binary Exponential Backoff (BEB). This section also provides a framework to compute the performance of a backoff algorithm. It is demonstrated that an upper efficiency limit exists, under the assumption that the stations are unaware of other stations' intentions to transmit.

In Sec. IV we show that, by a simple modification of BEB, the channel access medium is converted from pure CSMA to a hybrid CSMA-TDMA. We call the new mechanism Learning-BEB (L-BEB) because it progressively learns from both successful and unsuccessful transmission attempts, in order to migrate to TDMA-like operation. L-BEB is even simpler than legacy BEB and does not require any additional signaling. In the worst case, the performance delivered by L-BEB is the same as the performance that is currently obtained from legacy BEB.

The argumentation of Sec. IV is supported by the simulation results in Sec. V. Specifically, it is shown that L-BEB outperforms legacy BEB for any number of stations by reducing collisions and increasing the number of successful slots. Further, cumulative collision plots are used to show how the system learns from previous transmission attempts and the number of collisions is reduced as the simulation progresses.

Finally, the paper is concluded in Sec. VI.

II. Related Work

The Aloha [2] protocol laid the foundations for many random access protocols to come. In random access protocols, the nodes optimistically send their packets. In Aloha, a node with data ready to send, sends it immediately. The nodes involved in a collision wait a random period of time before attempting retransmission. In CSMA [3], the nodes are smarter and listen before talking, thus reducing the chances of collisions.

Reservation-Aloha (R-Aloha), presented in [4] and further analyzed in [5], already proposed a combination of random access and TDMA. The time is divided in slots which are grouped in frames. The duration of the slots is fixed and the duration of frames is chosen to be longer than the propagation delay of the broadcast channel. When station X successfully transmits in slot Y of a frame, it implicitly reserves slot Y for the next frame. The reservation can be released either explicitly, using a special flag in the last packet transmission, or implicitly by not sending a packet in the reserved slot. R-Aloha presents several disadvantages when compared to protocols that are currently in use, such as IEEE 802.11. First, the fixed length of slots implies that a high fraction of the channel time is wasted due to empty slots (In IEEE 802.11, the empty slots are orders of magnitude shorter than busy slots). Second, R-Aloha requires time synchronization among terminals. And third, the number of slots in a frame effectively limits the maximum number of active terminals. As the frame becomes full, new entrants do not have any chance to transmit. If the frame size is variable, additional signaling is required to inform all the stations about the current frame size.

Packet Reservation Multiple Access (PRMA) [6] and Centralized-PRMA [7] further extended the idea of R-Aloha to support heterogeneous (real-time and bursty) traffic. However, the improvements came at the price of higher complexity and signaling requirements.

A CSMA-TDMA hybrid MAC protocol was also explored in [8]. It is called Probabilistic TDMA (PTDMA). As in TDMA, the time is divided in time slices called slots which are grouped in frames. A station can own a slot in the frame. If this is the case, a station can transmit in that slot with probability a. Otherwise, if the station does not own the slot, it can also transmit with probability b. a and b satisfy the following equation:

$$a + (n-1)b = 1 \tag{1}$$

where n is the number of senders. For low values of n, the behaviour of PTDMA is closer to CSMA. As the number of stations increase, the probability that a station transmits in a non-owned slot is reduced and the behaviour of PTDMA is biased towards TDMA.

In the context of wireless sensor networks, there have been recent research efforts in the field of CSMA-TDMA hybrids. Z-MAC [9] aims to join the advantages of CSMA and TDMA in a single protocol. From CSMA, it takes high channel utilization and low latency under low contention; as TDMA, it offers high channel utilization and a limited number of collisions under high contention. Differently from our proposal, it specifically addresses multi-hop networks. The downside of Z-MAC is its increased complexity, which include neighbor discovery, slots assignment, local frame exchange and global time synchronization.

As opposed to the related work described in this section, the protocol proposed in this paper is based on an extremely simple modification to the protocol currently in use. This modification can be even considered a simplification. Another key differentiation aspect is that our proposal supports different slot durations (as IEEE 802.11 does), allowing the empty frames to be shorter than transmission frames. This option dramatically boosts the performance by reducing the time that the channel remains empty. All the previous work cited above assumes fixed slot duration.

A separate line of research consists on squeezing the maximum efficiency out of BEB by tuning its operation parameters, without making any CSMA-TDMA hybridization attempts. This avenue of research has its origins in the finding that the optimal transmission probability in BEB is a function of the packet length (l) and the number of competing stations (n) [10].

It is natural to attempt to estimate the number of contending stations to optimize the performance of BEB. The fast and accurate estimation of n is not a trivial task and advanced filtering techniques are required. An extended Kalman filter is used in [11] while [12] further improves the estimation by means of a bayesian approach.

Nevertheless, even if perfect estimation of the number of contending stations is achieved, the obtained efficiency never surpasses the upper bound for BEB, which is further detailed in the next section.

Our proposal easily breaks the upper bound for BEB and neither requires the estimation of n nor the dynamic adjustment of the operation parameters.

III. Binary Exponential Backoff and Performance Analysis

This section introduces Binary Exponential Backoff (BEB) which is part of the popular suite of protocols IEEE 802.11. This protocol is an example of a CSMA algorithm in which the stations transmit without any previous knowledge about other stations' intentions to transmit. The second part of this section assesses the performance of BEB, and finds the theoretical efficiency upper bound for this sort of algorithms.

Throughout the analysis, a number of usual assumptions are adopted. These include the supposition that all the stations are in the transmission range of one another, *i.e.* there is no hidden terminal effect [13]. The time is divided in slots, and the stations are synchronized to those slots. Transmission attempts can occur only at the beginning of a slot. Additionally, an ideal channel is assumed and frame losses are caused only by collisions. To simplify the analysis, all the stations transmit using the same data rate. The frame length is also the same for all stations.

A. Binary Exponential Backoff

The Medium Access Control (MAC) mechanism used in IEEE 802.11 networks is called Distributed Coordination Function (DCF). Although the standard considers also a centralized alternative - the Point Coordination Function - it has been sparsely implemented.

DCF uses a truncated Binary Exponential Backoff strategy. When a station that has its MAC queue empty receives a packet from the upper layer, it is allowed to transmit the packet after sensing the channel empty [1]. Otherwise, when the MAC queue is not empty and a packet arrives to the head-of-line of the MAC queue after the previous packet is successfully transmitted, the station has to backoff.

The backoff consists on drawing a number from a Contention Window $[0, CW)$ and waiting for that number of slots before transmitting. For the first transmission attempt the minimum contention window is used (CW_{min}). If there is a collision, the contention window doubles ($CW = 2 \cdot CW_{min}$) and the station randomly chooses a new number and waits for that number of slots before re-attempting transmission. The CW doubles after each collision until it reaches a maximum value CW_{max}. After a successful transmission, the value of CW is reset to its minimum. Vanilla IEEE 802.11 takes the values 32 and 1024 for its minimum and maximum contention windows, respectively.

With the IEEE 802.11e [14] standard amendment for Quality of Service support, the values of CW_{min} and CW_{max} can vary. However, the essence of the BEB remains the same.

For our analysis we will consider traffic sources that are saturated, *i.e.* each active station has always a packet ready to transmit. Intuitively, if there is only one active station in the network, this station is expected to transmit one slot in every 16 slots. The reason is that the actual number of empty slots between transmissions will uniformly vary from 0 to 31.

It is apparent that an efficiency problem exists, since only one of every 16 slots is used while the rest remain empty. Nevertheless the problem is not as acute at it may seem at a first glance, because an empty slot is much shorter than a busy slot. Actually, the duration of an empty slot is $20\mu s$ in IEEE 802.11b while the duration of a successful slot is in the order of ms. The exact value of the latter depends on the length of the data contained in the packet.

As the number of stations increases, the number of empty slots decreases. At the same time, there are higher chances that two or more stations transmit on the same slot and that the packets are lost due to a collision. A slot containing a collision is even longer than a successful slot. Therefore it is critical to reduce the number of collisions.

[1]The channel has to be sensed idle for a DIFS (DCF Inter Frame Space) period of time.

The BEB reacts to collisions by doubling the contention window, thus diminishing the transmission rate of the stations. This reaction reduces the load on the network and should decrease the collision probability. Note, however, that it is necessary that there is one collision for the algorithm to realize that the network is highly loaded. Since the value of CW is reset to CW_{min} after a successful transmission, the station has to learn about the network congestion conditions for every packet, and every time there has to be a collision for the station to adjust its CW value. This is a relatively high price to pay for adjusting the CW to its optimal value.

Studies in [15] show that small contention windows are desirable when the number of contending stations is low, since a small contention window reduces the number of empty unused slots. Conversely, for a large number of stations, larger contention windows offer better performance because they reduce the collision probability. The framework provided by IEEE 802.11e can be used to dynamically tune the values of CW_{min} and CW_{max} to adapt to the number of contending stations. However, as explained in the previous section, this strategy requires previous estimation of the number of active stations n [12].

The qualitative analysis of BEB presented above describes the trade-off incurred in choosing the right CW. A quantitative analysis of BEB can be obtained using Markov Chains and the assumption that, regardless of the number of retransmissions, a packet collides with constant probability [16]. Using that model, it is possible to compute the probability that a given station attempts transmission in a given slot (τ). This probability can then be used to obtain the probability of an empty, a successful and a collision slot. With these values, the overall performance of BEB can be evaluated and compared to other mechanisms.

The backoff process pursues the random distribution of the transmission attempts among the slots. An important goal is to maximize the number of successful transmissions while minimizing the collision probability. It is also important to keep the number of empty slots relatively low. However, an empty slot is much more desirable than a collision since the duration of the empty slots is orders of magnitude lower than the duration of a collision.

B. Efficiency of CSMA Algorithms

In CSMA algorithms the stations autonomously decide whether to transmit or not. The transmission probability (τ) is the key parameter that determines the probability of empty, successful or collision slot (P_e, P_s and P_c [2] respectively). For a given number of contending stations n:

[2]The notation P_c is used in this paper to denote the probability that a slot is busy with collision. This is different to the conditional collision probability (p or p_c in many papers) which is the probability that a collision occurs conditioned to the event that a tagged station attempts transmission.

$$P_e = (1-\tau)^n, \quad (2)$$
$$P_s = n\tau(1-\tau)^{n-1}, \quad (3)$$
$$P_c = 1 - P_e - P_s. \quad (4)$$

The transmission probability τ for BEB can be derived from [16] and is:

$$\tau = \frac{2(1-2p_{cc})}{(1-2p_{cc})(CW_{min}-1)+p_{cc}CW_{min}(1-(2p_{cc})^m)},$$
$$p_{cc} = 1-(1-\tau)^{n-1}. \quad (5)$$

p_{cc} is the conditional collision probability; the probability that a collision occurs given that one tagged station is attempting transmission. CW_{min} is the minimum contention window and m the maximum backoff stage:

$$m = log_2\left[\frac{CW_{max}}{CW_{min}}\right]. \quad (6)$$

We define the efficiency (ϕ) as the fraction of time that the channel is used for successful transmissions. It is understood that the time that the channel remains empty or busy with collisions is wasted. The efficiency is a function of the probabilities described in (2) - (4) and the duration of an empty, successful and collision slot (T_e, T_s and T_c respectively).

$$\phi = \frac{T_s P_s}{T_e P_e + T_s P_s + T_c P_c}. \quad (7)$$

In (7) we can observe that the duration of empty, successful and collision slots also affect the observed efficiency. While T_e is constant and defined in the standard, T_s and T_c are a function of the length of the frames. The duration of successful and collision slots are similar, thus the duration of a collision can be approximated to the duration of a successful slot $T_c \approx T_s$. Using the approximation and substituting (2) - (4) into (7) we obtain:

$$\phi = \frac{n\tau(1-\tau)^{n-1}}{1-\frac{T_s-T_e}{T_s}(1-\tau)^n} \quad (8)$$

From (8) it can be observed that the efficiency increases when using large frames. Given a number of contending stations n and a successful slot duration T_s, the optimal transmission probability τ that maximizes efficiency satisfies:

$$\frac{d\phi}{d\tau} = \frac{(1-\tau)^{n-1}+(n-1)\tau(1-\tau)^{n-2}}{1-\frac{T_s-T_e}{T_s}(1-\tau)^n} - \frac{\frac{T_s-T_e}{T_s}n\tau(1-\tau)^{2(n-1)}}{(1-\frac{T_s-T_e}{T_s}(1-\tau)^n)^2} = 0 \quad (9)$$

In Figure 1, the efficiency using optimal values of τ (derived from (9)) is plotted. This values are compared to the ones that are actually obtained when using the values of τ provided by legacy BEB (which are derived from (5)).

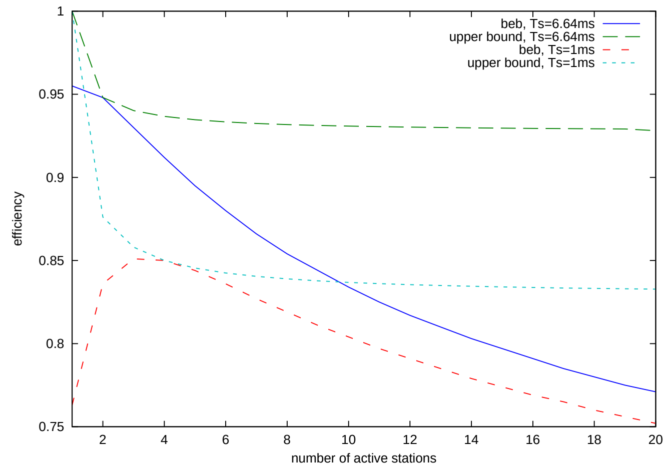

Fig. 1. This figure compares the efficiency of BEB to the efficiency that would be obtained if optimal transmission probability was used. The results are presented for different values of successful slot duration T_s.

The curves in Figure 1 for optimal transmission probability represent an upper bound for BEB, and for those protocols that simply tune the parameters of BEB in response to the number of competing terminals. To surpass that upper bound, it is not sufficient to adjust the size of the CW_{min}. Conversely, it is required that the stations gain some kind of knowledge about the other stations' future intentions to transmit. This can be achieved by setting the stations' backoff to a deterministic value after a successful transmission.

IV. Learning-BEB

In BEB, the stations randomly access the channel, without any attempt to collect any feedback from previous transmission attempts. This means that, if two saturating stations compete for the channel for a long time, the collision probability perceived by the stations remains constant. After a transmission attempt, a station samples a random backoff number from CW_{min} if the transmission attempt was successful. Otherwise, the current contention window is doubled before drawing the backoff number.

It is easy to modify the protocol to allow the stations to learn from previous transmission attempts and decrease the number of collisions. Consider the same example of two stations competing for the channel. In this case, the stations use a constant backoff value ($V = 16$) after a successful transmission. At the beginning, the two stations randomly transmit without any knowledge about the other station's intention to transmit. However, as soon as the two stations successfully consecutively transmit, each of the stations periodically transmit every $V = 16$ slots. Since the selection of the transmission slot is deterministic, the chances of suffering collisions disappear, and the stations will orderly transmit in a TDMA fashion.

Figure 2 shows a graphical example. It represents two time lines divided in slots. Even though the actual duration of empty, successful and collision slots is different, in the figure they are all represented equal for simplicity.

In the upper time line, the stations operate using legacy BEB. The two stations collide in their first transmission attempt. The stations double their contention window and draw backoff values, specifically 13 and 22. After 13 empty slots, *STA 0* transmits, and obtains a new backoff value equal to 23. After 8 empty slots, *STA 1* successfully transmits and draws a backoff value of 14. The result, in this example, is that both stations will collide in their following transmission attempt.

In the lower time line of Figure 2 the stations use Learning-BEB (L-BEB), which is the backoff mechanism proposed in this paper. The beginning of the example is similar: the two stations collide and draw different backoff numbers. At this point, the behaviour of the system will become completely deterministic. *STA 0* successfully transmits and thus sets its backoff counter to $V = 16$. After eight empty slots *STA 1* successfully transmit and sets its backoff counter to $V = 16$. 6 empty slots later, *STA 0* successfully transmits again. And after eight empty slots it is *STA 1*'s turn. Both stations continue to transmit in turns occupying slots 0 and 9 of a virtual 16-slot TDMA frame. The suppression of collisions should be warmly welcomed because implies more efficient channel utilization.

We introduce the concept of *virtual frame* to highlight the similarities with TDMA. The virtual frame consists on V slots. Throughout this article we consider $V = 16$ for similarity with legacy BEB. In legacy BEB a station uniformly draws a random number between 0 and 31 after a successful transmission. Thus we choose L-BEB to wait for 16 slots after a successful transmission. The value of V can be tuned to adjust the behaviour of L-BEB. Although we provide some insights about the implications of tuning V by the end of this section, an exhaustive study is considered out of the scope of this work.

In the example in Figure 2, the virtual frames appear as a dotted line. After a successful transmission, a station will retransmit in the same slot position in the next virtual frame. In Figure 2, if we number the frame's position from 0 to 15, *STA 0* and *STA 1* transmit in positions 0 and 9 respectively.

The frame is virtual because there is neither explicit signaling nor configuration to assign a slot to a station. Additionally, the virtual frame only applies to those stations that successfully transmit, because the rest operate as in legacy BEB by selecting random backoff numbers. Moreover, a station that deterministically selects its next transmission slot does not have any kind of reservation for that slot.

For exemplifying purposes we have considered the simplest case of two contending stations. Nevertheless, the same conclusions apply for an increasing number of stations up to V. There is an initial transitory phase in which collisions occur and the stations try to find their place in the virtual TDMA frame. A station that successfully transmits in a given slot of the virtual frame, will keep transmitting in the same slot until a collision occurs. If a collision occurs, the station should draw a random backoff number from a doubled contention window. Eventually, all the stations will sequencially successfully transmit. At this point, each station has found its slot in the virtual frame and collisions will vanish.

Obviously, the higher the number of contending stations, the longer it will be the transitory operation of the protocol, since it is more difficult that all stations choose a different slot. In a naive approximation, we consider the probability that n stations choose different slots from a V-slot virtual frame.

$$\prod_{i=1}^{n-1}\left(1-\frac{i}{V}\right) \quad ; \quad 1 < n \leq V \tag{10}$$

If the value of n is low, all the stations will probably choose a slot different from the others. Oppositely, when the value of n is larger, it is more probable that some of the stations successfully transmit while others collide. Continuing with the approximation of the virtual frame, if n_s ($n_s < n$) stations successfully transmitted in the previous virtual frame, the probability that the rest of stations choose a slot that does not result in collision is:

$$\prod_{i=n_s}^{n-1}\left(1-\frac{i}{V}\right) \quad ; \quad 1 \leq n_s < n \leq V \tag{11}$$

The intuition is that the higher the value of n_s, the closer we are to the TDMA operation of the system. In the next section, simulations will be used to find out how long it takes to reach the stationary condition, depending on the number of active stations.

Special attention deserves the case in which the number of contending stations n is greater thant the size of the virtual frame (V). It is not possible to fit more than V stations in a frame containing V slots. Thus, the system will not reach stability and the collisions will not completely disappear, no matter how long the system is running. Nevertheless, the system still performs as a CSMA/TDMA hybrid and, therefore, outperforms pure CSMA. There will be some stations that successfully transmit and deterministically choose their backoff, while the others collide and operate as CSMA stations.

A. Limitations of L-BEB

L-BEB shows its full potential after a short period of learning process. Ideally, after each station has found its place in the virtual frame, the system operates without collisions, until a perturbation moves the system back to the transitory phase. This perturbation could appear in the form of a new station entering the contention. It might happen that the new entrant successfully transmits in its first transmission attempt. If this is the case, no collision occurs and the systems continues its TDMA-like operation. Otherwise, when a collision occurs, there will be two stations selecting a random backoff algorithm before re-attempting transmission. This two stations might, in turn, generate new collisions initiating a chain reaction that brings the system to its transitory CSMA-like operation. Therefore, in a scenario with a high number of new entrants (in the order of multiple new incorporations per second), the

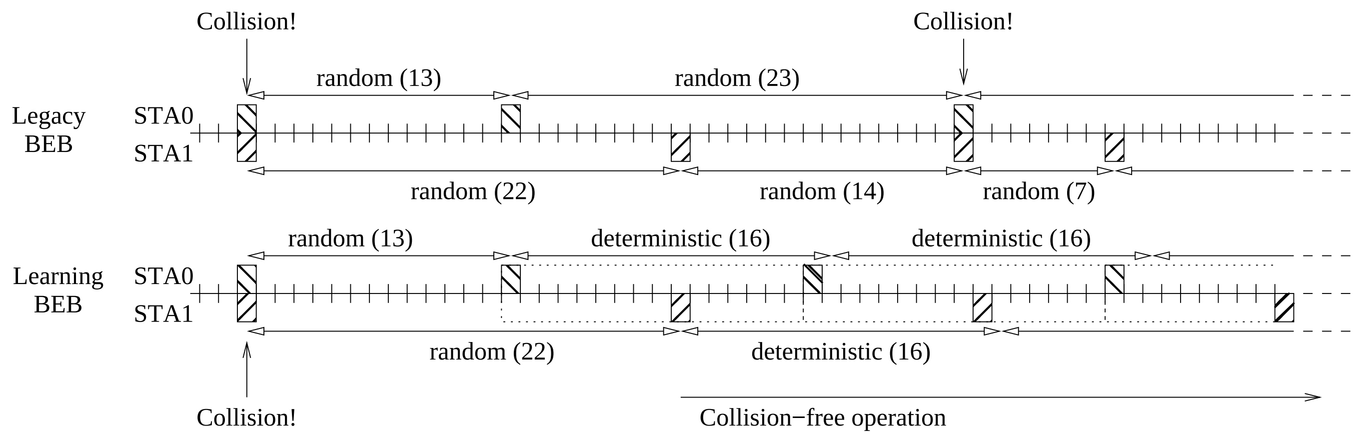

Fig. 2. The operation of BEB and L-BEB is compared. While the former computes a random backoff after a successful transmission, the latter always chooses 16. In L-BEB, after all the participating stations successfully consecutively transmit, the system enters in collision-free operation.

medium access mechanism will be closer to CSMA and the advantages of using L-BEB will not be so obvious.

The transitory operation of the protocol can be shortened by increasing V. It can be observed that a higher value of V leads to higher success probabilities in (10) and (11). Moreover, a higher value of V allows for more terminals to operate in a collision-free fashion since collision-free operation is only possible when $n \leq V$. However, increasing the virtual frame size V has the side effect of lowering the efficiency when the number of contending stations is low. One could argue that V should be chosen as a function of the number of contending stations. However, the estimation of the number of contending stations is not trivial. For this reason, we opt for a static configuration of V for the paper and leave the dynamic selection of V for further study.

In a realistic scenario, a packet might be lost due to bad channel conditions. A station losing a packet would not be able to differentiate whether the packet was lost due to a collision or because of poor channel conditions. In any case, the station will double the contention window and draw a random backoff number. This action will also endanger the stability the same way a new entrant does.

Finally, the argumentation in Figure 2 is valid only if all the stations share the same vision of the channel, i.e. if all the stations can listen to all the successful and collision slots. If there is a station that can not listen to another station transmission, the slot count would be different for different stations and the system performance would be the same as the one obtained in BEB. This last problem can be alleviated by using request-to-send and clear-to-send packets (RTS/CTS).

RTS/CTS signaling packets are transmitted before the actual data transmission and include a Network Allocation Vector (NAV) that describe the channel occupation intentions. RTS and CTS are sent by the sender and the receiver, respectively. Therefore, the channel occupation information reaches all the stations that can hear either the sender or the receiver. The RTS/CTS also limits the impact of collisions, since collisions can only occur in signaling (short) packets. Nevertheless, the RTS/CTS mechanism adds extra signaling overhead thus reducing the overall efficiency of the channel.

V. SIMULATION RESULTS

The goal of the simulations[3] is to show that a performance improvement can be obtained by substituting BEB for L-BEB. The performance is a function of the number of empty, successful and collision slots. It is desirable to maximize the number of successful slots while minimizing the number of collisions. Empty slots play a minor role in the performance evaluations, because they are much shorter than successful transmissions and collisions.

By counting the number of successful and collision slots, the performance of a backoff algorithm can be evaluated. Nevertheless, while the performance of BEB is maintained along the simulation, L-BEB learns and delivers better performance by the end of the simulation. However, we will postpone the analysis of the evolution of the performance of L-BEB. For the first round of simulations, the average number of empty, successful and collision slots in the first 1000 slots are studied.

The value of 1000 is arbitrary, choosing a higher value would highlight the advantage of L-BEB whereas a lower value would bring the curves in the plots closer. Please note that even though all the simulations contain the same number of slots, the simulated time of the different simulations is not equal. The reason is that the duration of empty, successful and collision slots is different. Assuming a successful transmission time $T_s = 6.64$ ms and that halve of the slots are empty, the duration of 1000 slots would be 3.33 seconds.

The number of active stations in the simulations ranges from 2 to 20. Figure 3 compares the number of collisions in the first 1000 slots when using BEB and L-BEB. Each simulation is repeated 100 times and both the average and the 95% confidence interval are computed. It can be observed that by employing L-BEB instead of BEB, the number of collisions is reduced for any number of stations from 2 to 20. Even when

[3]The simulations were performed in Octave. All the scripts are available upon request to the corresponding author.

the number of stations is greater than the size (in slots) of the virtual frame (V), L-BEB consistently achieves a lower number of collisions than BEB.

Figure 4 shows the number of successful slots. The first observation is that the number of successful slots is much higher when using L-BEB. This is a direct consequence of the lower number of collisions. Remember that, after a collision, the stations double their contention window and therefore reduce their transmission rate. L-BEB, reduces the number of collisions and allows the stations to keep a higher sending rate. Further, thanks to the CSMA-TDMA hybridization, L-BEB permits that the higher transmission rate does not translate to a higher number of collisions. This is true even when the number of contending stations is higher than the number of slots in the virtual frame.

The values in Figure 4 are those obtained in the first 1000 (transitory) slots. In steady state (collision-free) operation, the fraction of successful slots is n/V for $n \leq V$.

The original goal of the article was to increase the efficiency of the MAC access protocol by reducing the number of collisions. When the stations orderly and deterministically transmit, it is possible to outperform legacy BEB. It is also possible to cross the upper limit associated to random transmission. This is shown in Figure 5. The values of the axes replicate those of Figure 1 to ease comparison. The efficiency values presented in Figure 5 are those obtained during the first 1000 slots. As the system keeps learning, the efficiency further increases. The efficiency is a also a function of the frame length; for this reason, the efficiency is plotted for two values of T_s (the same values that were used in Figure 1).

Figure 6 evaluates the duration of the transitory in L-BEB. The transitory is characterized by collisions, while in steady-state conditions collisions theoretically disappear for $n \leq V$. To evaluate how the number of collisions fluctuate along the first 1000 slots, we plot the cumulative number of collisions. The cumulative number of collisions steadily grows at the beginning of simulations and becomes flat as the simulation advances and collisions disappear. The results presented in the plot are the average of the 100 simulations.

It can be observed that when the number of stations is 8, the steady-state condition is reached in about 200 slots. If the number of active stations is increased to 12, the steady-state condition is not reached within the simulation, since the curve does not become completely flat. Nevertheless, a reduction in the number of collision can be appreciated as the simulation progresses. Finally, for the case of $n = V$, the number of collisions is high even by the end of the simulation. Even though it is theoretically feasible to reach a steady-state condition without collisions for $n = V$, the probabilities are so small that it is not something that we can expect to happen in simulated or real scenarios.

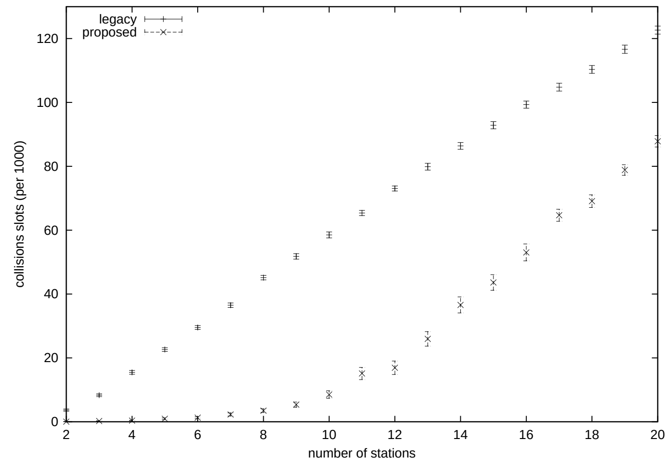

Fig. 3. The number of collision slots out of the total 1000 slots are plotted in this figure. The number of active stations range from 2 to 20. The results are presented with 95% confidence intervals

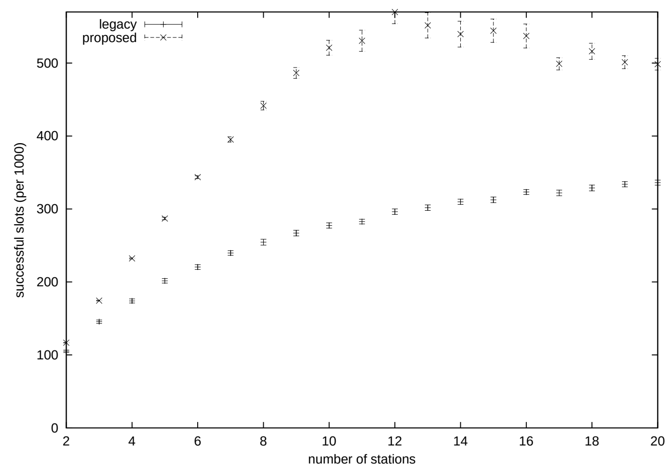

Fig. 4. The number of successful slots out of the total 1000 slots are plotted in this figure. The number of active stations range from 2 to 20. The results are presented with 95% confidence intervals

VI. Conclusion

This article addresses MAC protocols for wireless local area networks. In the extensively used Binary Exponential Backoff, the stations randomly select backoff (waiting) values to separate transmission attempts. Prior art struggled to optimize the parameters of BEB to improve its efficiency. Nevertheless, even if optimal transmission probability is used, the efficiency of the channel utilization is far from 100%. The explanation is that the stations blindly transmit, unaware of other stations's intentions to transmit.

We propose a framework to compute the efficiency of a MAC mechanism and we apply it to analyze BEB. We also derive the maximum efficiency that can be obtained from non-learning backoff schemes. Then we suggest a minor change to BEB: choosing a deterministic backoff value after a successful transmission, instead of a random one. By this simple modification, we allow that the stations learn from both collisions and successful transmission, thus reducing the chances of future collision. The system initially performs as legacy BEB, but after a few transmissions, the benefits of learning become clear and the collisions diminish. When the

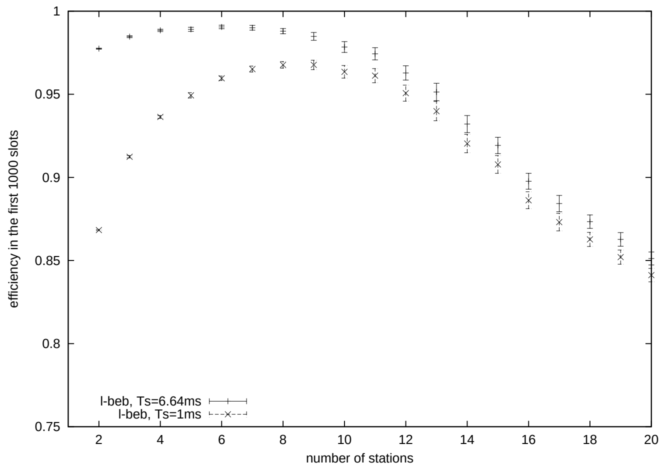

Fig. 5. The efficiency of L-BEB during the first 1000 slots as obtained from the simulations. The number of contending stations range from 2 to 20. The 95% confidence intervals are also plotted.

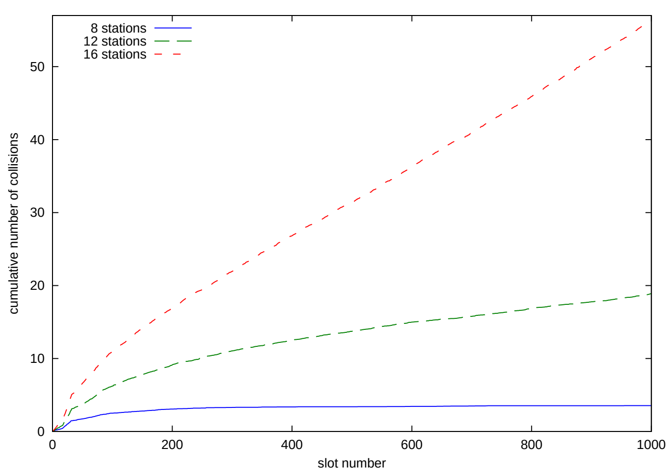

Fig. 6. The cumulative number of collisions in each of the first 1000 slots of the simulation of L-BEB. As the system reaches steady-state, collisions disappear and the curves become flat.

number of stations is lower than V, the collisions eventually vanish and the system operates in a collision-free fashion.

We call the new MAC algorithm Learning-BEB, since its performance improves over time, until a new station is activated. If this occurs, L-BEB has to learn again to adapt to the new scenario. Simulations have been used to show that L-BEB, in addition to reduce the number of collisions, also increases the number of successful transmissions. The combination of both effects positively impacts the efficiency, pushing it higher than the upper bound for non-learning algorithms.

The simulations are also used to understand the learning curve of L-BEB by analyzing the cumulative number of collisions over the first 1000 slots of the simulations. The conclusion is that, when the number of contending stations is low (lower than 8), the system quickly enters in collision-free operation. However, when the number of contending stations is higher, the learning pace is slower and the systems spends a long time in transitory operation. During the transitory, collisions still occur. Nevertheless, the number of collisions is lesser than in legacy BEB. If the number of contending stations is greater than the virtual frame size V, the system never reaches a steady-state (collision-free) condition. However, even in this extreme situation, L-BEB still outperforms BEB.

The main contribution of this paper is suggesting a minor change in the Binary Exponential Backoff that reduces the complexity and dramatically boosts efficiency.

Acknowledgment

The authors would like to thank the anonymous reviewers for their insightful comments.

References

[1] *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, ANSI/IEEE Std. Std 802.11, 1999 Edition (Revised 2003).

[2] N. Abramson, "The ALOHA System–Another Alternative for Computer Communications," *Cluster Computing*, vol. 5, pp. 187–201, 1970.

[3] L. Kleinrock and F. Tobagi, "Packet Switching in Radio Channels: Part I–Carrier Sense Multiple-Access Modes and Their Throughput-Delay Characteristics," *Communications, IEEE Transactions on [legacy, pre-1988]*, vol. 23, no. 12, pp. 1400–1416, 1975.

[4] W. Crowther, R. Rettberg, D. Walden, S. Ornstein, and F. Heart, "A system for broadcast communication: Reservation-ALOHA," *Proc. 6th Hawaii Int. Conf. Syst. Sci*, pp. 596–603, 1973.

[5] S. Tasaka, "Stability and Performance of the R-ALOHA Packet Broadcast System," *IEEE Transactions on Computers*, vol. C-32, pp. 717–726, Aug. 1983.

[6] D. Goodman, R. Valenzuela, K. Gayliard, and B. Ramamurthi, "Packet reservation multiple access for local wirelesscommunications," *Communications, IEEE Transactions on*, vol. 37, no. 8, pp. 885–890, 1989.

[7] G. Bianchi, F. Borgonovo, L. Fratta, L. Musumeci, and M. Zorzi, "C-PRMA: a centralized packet reservation multiple access for localwireless communications," *Vehicular Technology, IEEE Transactions on*, vol. 46, no. 2, pp. 422–436, 1997.

[8] A. Ephremides and O. Mowafi, "Analysis of a Hybrid Access Scheme for Buffered Users- Probabilistic Time Division." *IEEE Transactions on Software Engineering*, vol. 8, no. 1, pp. 52–60, 1982.

[9] I. Rhee, A. Warrier, M. Aia, and J. Min, "Z-MAC: a hybrid MAC for wireless sensor networks," *Proceedings of the 3rd international conference on Embedded networked sensor systems*, pp. 90–101, 2005.

[10] Y. Tay and K. Chua, "A Capacity Analysis for the IEEE 802.11 MAC Protocol," *Wireless Networks*, vol. 7, no. 2, pp. 159–171, 2001.

[11] G. Bianchi and I. Tinnirello, "Kalman filter estimation of the number of competing terminals in an IEEE 802.11 network," *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies. IEEE*, vol. 2.

[12] A. Lopez-Toledo, T. Vercauteren, and X. Wang, "Adaptive Optimization of IEEE 802.11 DCF Based on Bayesian Estimation of the Number of Competing Terminals," *IEEE TRANSACTIONS ON MOBILE COMPUTING*, vol. 5, no. 9, p. 1283, 2006.

[13] A. Zahedi and K. Pahlavan, "Natural hidden terminal and the performance of the wireless LANs," *Universal Personal Communications Record, 1997. Conference Record., 1997 IEEE 6th International Conference on*, vol. 2, 1997.

[14] "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications; Amendment: Medium Access Control(MAC) Quality of Service Enhancements," no. IEEE Std 802.11e, 2005.

[15] H. Anouar and C. Bonnet, "Optimal Constant-Window Backoff Scheme for IEEE 802.11 DCF in Single-Hop Wireless Networks Under Finite Load Conditions," *Wireless Personal Communications*, pp. 1–20.

[16] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *Selected Areas in Communications, IEEE Journal on*, vol. 18, no. 3, pp. 535–547, 2000.