ACKS: A Technique to Reduce the Impact of Legacy Stations in 802.11e EDCA WLANs

Luca Vollero, *Member, IEEE,* Albert Banchs, *Member, IEEE,* and Giulio Iannello

***Abstract*— The EDCA access mechanism of the upcoming 802.11e standard supports legacy DCF stations, but with substantially degraded performance. The reason being that DCF stations typically compete for access with overly small Contention Windows (CW's). In this letter we propose a new technique that, implemented at the Access Points (AP's), mitigates the impact of legacy stations on EDCA. The key idea of the technique is that, upon receiving a frame from a legacy station, the AP skips the ACK frame reply with a certain probability. When missing the ACK, the legacy station increases its CW and thus our technique allows us to have some control over the CW's of the legacy stations. We show by means of an example that this technique improves the overall performance of the WLAN.**

***Index Terms*— WLAN, 802.11e, 802.11, EDCA, DCF, ACKS.**

I. Introduction

ONE of the main problems of the EDCA (*Enhanced Distributed Channel Access*) mechanism of the upcoming IEEE 802.11e standard is that legacy IEEE 802.11 stations are not well supported. Although legacy stations using the DCF (*Distributed Coordination Funtion*) mechanism of 802.11 can operate in 802.11e WLANs under EDCA, the contention parameters with which these legacy stations compete cannot be controlled. This results in a degraded performance of the WLAN.

In this letter we propose a new technique, which we call the *ACK Skipping* (*ACKS*) technique, that mitigates the impact of legacy stations on an 802.11e WLAN under the EDCA mechanism working in Infrastructure Mode. The technique requires only a small modification in the 802.11e Access Point (AP), and leaves the legacy stations untouched. We show that this technique improves the performance of the WLAN.

II. 802.11 DCF and 802.11e EDCA

DCF and EDCA execute a very similar algorithm to transmit packets, which is described as follows. A station with a new packet to transmit monitors the channel activity. If the channel is idle for a period of time (named $DIFS$ in DCF and $AIFS$ in EDCA), the station transmits. Otherwise, it continues to monitor the channel until it is idle for a $DIFS/AIFS$, and, at this point, the backoff process starts by computing a random value uniformly distributed in the range $(0, CW - 1)$ and initializing the backoff time counter with this value. The CW parameter is called the Contention Window and its value depends on the number of failed transmissions. At the first transmission attempt, it is set equal to CW_{min}.

The backoff time counter is decremented once every slot time as long as the channel is sensed idle, "frozen" when a transmission is detected on the channel, and reactivated when the channel is sensed idle again for a $DIFS/AIFS$. As soon as the backoff time counter reaches zero, the station transmits. A collision occurs when two or more stations start transmission simultaneously. After each unsuccessful transmission, CW is doubled, up to a maximum value CW_{max}. If the number of failed attempts reaches a predetermined retry limit R, the packet is discarded.

From the above explanation, it can be seen that the behavior of a station depends on a number of parameters (namely $DIFS$ or $AIFS$, CW_{min} and CW_{max}). The main difference between DCF and EDCA is that, while in DCF the values of these parameters are fixed by the standard, in 802.11e EDCA these are open configurable parameters that can be set to different values for different Access Categories (AC's).

III. *ACKS*: The ACK Skipping Technique

As we have seen in the previous section, legacy DCF stations start the backoff process with a CW equal to CW_{min}. This initial CW is fixed by the standard to a small value, and it only doubles after each failed attempt. These small CW values raise the following problem in a WLAN under EDCA with 802.11e and legacy stations, in which typically 802.11e stations are expected to receive a high priority service and legacy stations a low priority one:

1) If high priority 802.11e stations are configured with smaller CW values than legacy stations, such that they receive a better service, the resulting overall efficiency of the WLAN is low, due to the fact that small CW values result in a high collision rate.
2) If high priority 802.11e stations are configured with larger CW values, in order to preserve the overall efficiency, low priority legacy stations receive a better service than desired, leaving 802.11e stations with a worse service.

Manuscript received August 27, 2004. The associate editor coordinating the review of this letter and approving it for publication was Prof. C.-F. Chiasserini. This work was funded in part by the Ministero dell'Istruzione, dell'Universitá e della Ricerca (MIUR) with the FIRB project WEB-MINDS, and by the European Community's Sixth Framework Programme with the IST FP6 Integrated Project DAIDALOS. The work was performed during a visit of L. Vollero at the Universidad Carlos III de Madrid funded by the IST FP6 Network of Excellence E-NEXT. Apart from this, the European Commission has no responsibility for the content of this paper.

L. Vollero is with the Università degli Studi di Napoli "Federico II", 80125 Naples, Italy (email: vollero@unina.it).

A. Banchs is with the Universidad Carlos III de Madrid, 28911 Madrid, Spain (email: banchs@it.uc3m.es).

G. Iannello is with the Università Campus Bio-Medico di Roma, 00155 Rome, Italy (email: g.iannello@unicampus.it).

Digital Object Identifier 10.1109/LCOMM.2005.04028.

It is obvious that none of the above two alternatives is desirable, as in both cases the service received by high priority stations is seriously degraded as a consequence of the impact of legacy stations. Instead, it would be desirable to increase the CW of legacy stations; in this way, high priority stations could receive a better service than low priority legacy stations without compromising the overall efficiency. The *ACK Skipping* (*ACKS*) technique achieves this goal without modifying the legacy stations.

ACKS is based on the following behavior of legacy stations: after sending a packet, a legacy station waits for an Acknowledgement (ACK) frame, and, if the frame is not received within an ACK timeout, it assumes a collision and increases its CW. The central idea is then the following: if the AP skips the ACK reply to legacy stations with a probability δ, these stations will 'see' a collision rate higher than the actual one, and will contend with larger CW's, resulting this in a smaller impact on the 802.11e stations[1].

The challenge with the *ACKS* technique is the configuration of the probability δ. This adds to the inherent challenge with 802.11e of configuring the EDCA contention parameters in order to provide the desired behavior.

In the rest of this paper, we use the case of throughput allocation as an example to illustrate the effectiveness of *ACKS*. However, we strongly believe that the applicability of the technique is not restricted to this case. On the contrary, we claim (following the arguments exposed in this section) that the technique can be used to improve the performance of a WLAN in any case in which the CW configuration of the legacy stations corresponds to a better service than the desired for this type of stations.

IV. Throughput Allocation Criterion

While there are many different criteria proposed in the literature for throughput allocation, *weighted max-min fairness* [2] is a widely accepted one; it is e.g. the one implemented by weighted fair queuing in wired links.

The weighted max-min fair allocation is the one that maximizes the minimum r_i/w_i in the system, r_i being the throughput allocated to entity i and w_i the entity's weight. Based on this criterion, we set our objective here to provide weighted max-min fairness in the WLAN, the WLAN stations being our *entities*, and the saturation throughput of a WLAN station its *allocated throughput*[2].

According to the weighted max-min fairness criterion, a configuration provides a better performance than another when it provides a greater $min(r_i/w_i)$. As a result, *ACKS* will be effective if and only if it is able to provide a greater $min(r_i/w_i)$ than any configuration of EDCA that does not use this technique.

[1] The idea of skipping the ACK reply with a certain probability was used in [1] for a different purpose, namely as a differentiation mechanism for 802.11 WLANs.

[2] The saturation throughput of a station in a WLAN is defined as the throughput that the station experiences if all stations always have packets to transmit. Note that this corresponds to the definition of *allocated throughput* in weighted max-min fairness, which assumes that all entities are using the throughput to which they are entitled.

V. *ACKS* Optimal Configuration

We next present an algorithm that finds the optimal configuration of the EDCA parameters and of the *ACKS* technique for weighted max-min fairness.

Let k be the number of AC's in the WLAN, w_i for $i \in \{1,\ldots,k\}$ the weight assigned to AC i, n_i the number of stations of the AC, w_{k+1} the weight assigned to the legacy stations and n_{k+1} the number of legacy stations. Our goal is to find the optimal configuration of all AC's (i.e. A_i, CW^i_{min} and CW^i_{max} for $i \in \{1,\ldots,k\}$), and of the probability δ with which ACK frames of the legacy stations are skipped.

Based on the arguments of [3], we set $A_i = 0$ and $CW^i_{min} = CW^i_{max}\ \forall i \in \{1,\ldots,k\}$, which leaves us with only one parameter per AC, CW_i. From [3], we have that the probability that a station of AC i transmits in a randomly chosen slot time under saturation conditions is

$$\tau_i = \frac{2}{CW_i + 1} \tag{1}$$

Similarly, from [4] we have that the transmission probability of a legacy station under saturation conditions is

$$\tau_{k+1} = \frac{2(1-2p)(1-p^{R+1})}{\begin{array}{c}W(1-(2p)^{m+1})(1-p) + (1-2p)(1-p^{R+1}) + \\ +W2^m p^{m+1}(1-2p)(1-p^{R-m})\end{array}} \tag{2}$$

where $W = CW_{min}$, m is such that $CW_{max} = 2^m CW_{min}$ (both set to the values given by the 802.11 standard) and p is the probability that a legacy station does not receive an ACK after transmitting a frame,

$$p = 1 - (1-\tau_{k+1})^{n_{k+1}-1} \prod_{i=1}^{k} (1-\tau_i)^{n_i}(1-\delta) \tag{3}$$

In the following, we present an algorithm that finds the optimal τ_i's. Once the optimal τ_i values are obtained, we set CW_i to the integer value that most closely approximates τ_i,

$$CW_i = round\ int\left(\frac{2}{\tau_i} - 1\right) \tag{4}$$

Given τ_j for $j \in \{1,\ldots,k+1\}$, the throughput received by a station of AC i can be computed as

$$r_i = \frac{\tau_i(1-\tau_i)^{n_i-1}\prod_{j\in\{1,\ldots,k+1\}\setminus i}(1-\tau_j)^{n_j}}{T_{slot}}\, l \tag{5}$$

where l is the average packet length and T_{slot} is the average duration of a slot time (the reader is referred to [3] for the computation of T_{slot} as a function of τ_j, $j \in \{1,\ldots,k+1\}$).

Similarly, the throughput received by a legacy station can be computed as

$$r_{k+1} = (1-\delta)\frac{\tau_{k+1}(1-\tau_{k+1})^{n_{k+1}-1}\prod_{i=1}^{k}(1-\tau_i)^{n_i}}{T_{slot}}\, l \tag{6}$$

Using a similar reasoning to [3], it can be proven[3] that the following condition holds in the optimal configuration,

$$\frac{r_i}{r_j} = \frac{w_i}{w_j}\ \forall i,j \in \{1,\ldots,k+1\} \tag{7}$$

[3] Specifically, it can be seen that, for any configuration that does not comply with the condition, there exists an alternative configuration that complies with the condition and provides a better performance.

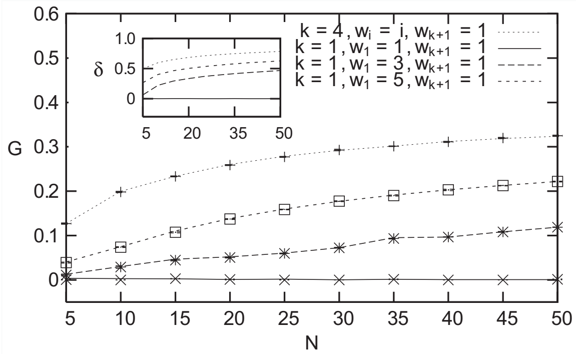

Fig. 1. Gain with the *ACKS* technique.

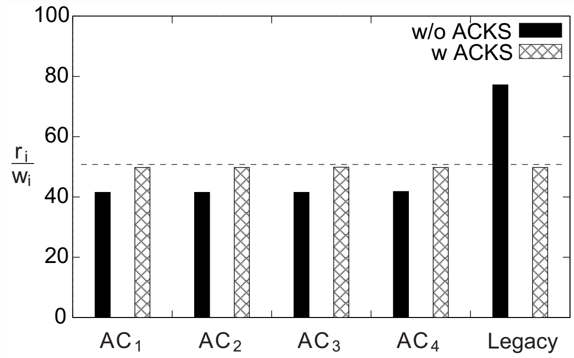

Fig. 2. Throughput distribution with and without *ACKS*.

Combining (5), (6) and (7) yields

$$\tau_j = \frac{w_j \tau_i}{w_j \tau_i + w_i(1-\tau_i)} \ \forall i, j \in \{1, \ldots, k\} \quad (8)$$

$$\tau_{k+1} = \frac{w_{k+1} \tau_i}{w_{k+1} \tau_i + w_i(1-\tau_i)(1-\delta)} \ \forall i \in \{1, \ldots, k\} \quad (9)$$

Let us imagine that the optimal configuration of τ_i for some AC i is known. Then, the optimal configuration of τ_j for all other AC's can be derived from (8), and τ_{k+1} can be expressed as a function of δ from (9),

$$\tau_{k+1}(\delta) = f_{(9)}(\delta) \quad (10)$$

τ_{k+1} can also be expressed as a function of δ from (2),

$$\tau'_{k+1}(\delta) = f_{(2)}(\tau_{k+1}, \delta) = f_{(2)}\left(f_{(9)}(\delta), \delta\right) \quad (11)$$

From the above, we have that δ can be obtained from solving the following non-linear equation,

$$\tau_{k+1}(\delta) = \tau'_{k+1}(\delta) \quad (12)$$

The left-hand side of the above equation is an increasing function that starts at $\tau_{k+1}(0)$ and grows up to 1. The right-hand side is a decreasing function that starts at $\tau'_{k+1}(0)$ and reduces down to $\tau'_{k+1}(1) < 1$. Therefore, if $\tau_{k+1}(0) \leq \tau'_{k+1}(0)$, the non-linear equation has only one solution, and otherwise it has no solution. Note that, when $\tau_{k+1}(0) > \tau'_{k+1}(0)$, a legacy station receives less than its share of throughput even with $\delta = 0$; therefore the optimal δ in this case is 0.

Let us take AC 1 as reference. With the above, given τ_1, we can derive the rest of the optimal parameters, namely the τ_i of the other AC's and δ. With these values, we can compute the throughputs of all stations, i.e. $r_i \ \forall i \in \{1, \ldots k+1\}$. The problem of finding the optimal configuration is thus reduced to finding the optimal τ_1 that maximizes $min(r_i/w_i)$. This can be performed using numerical techniques, which terminates the algorithm. Specifically, our algorithm uses the golden section search method to find the optimal τ_1. In each iteration, given a τ_1 value, the algorithm finds the corresponding δ by solving (12) numerically with the bisection method (unless $\tau_{k+1}(0) > \tau'_{k+1}(0)$, in which case we take $\delta = 0$).

We note that the computational cost of the proposed algorithm is very low. Indeed, we have assessed that the time it takes to run the algorithm in a Pentium 4 PC with 2.66 GHz of CPU speed is of a few tenths of ms only, which shows that the algorithm can be used at run-time.

VI. Performance Evaluation

We next assess the effectiveness of the *ACKS* technique; specifically, we measure the gain in performance resulting from using this technique. Based on the arguments given in Section IV, we measure the gain as follows

$$G = \frac{min(r'_i/w_i) - min(r_i/w_i)}{min(r_i/w_i)} \quad (13)$$

where the r'_i's are the throughputs obtained with *ACKS* and the r_i's are the throughputs obtained without *ACKS*. The configuration used in the former case is the one given by the algorithm of Section V. For the latter, we use a very similar algorithm with the only difference that we do not allow $\delta \neq 0$. For all experiments, the system parameters of the 802.11b standard and packet lengths of 1000 bytes have been taken.

Fig. 1 illustrates the gain obtained with *ACKS* for different configurations and numbers of stations ($n_i = N \ \forall i$). Analytical results are represented with lines and simulation results (average and 95% confidence intervals) with points and error bars. The figure also gives the optimal configuration of δ in each case. Fig. 2 depicts the r_i/w_i, obtained analytically, for each AC and for the legacy stations, in the case $k = 4$ with $N = 10$.

From the above results, we conclude that *ACKS* provides a substantial gain in performance, and that the technique is especially effective when 1) legacy stations are granted a low priority service, which is the typical case, and 2) the number of stations in the WLAN is large, which is the most critical case for performance.

References

[1] L. Vollero and G. Iannello, "Frame dropping: a QoS mechanism for multimedia communications in WiFi hot spots," in *Proc. International Workshop on Mobile and Wireless Networking (MWN'04)*, Aug. 2004, pp. 54–59.

[2] L. Massoulie and J. Roberts, "Bandwidth sharing: Objectives and algorithms," *IEEE/ACM Trans. Networking*, vol. 10, pp. 320–328, June 2002.

[3] A. Banchs, X. Pérez-Costa, and D. Qiao, "Providing throughput guarantees in IEEE 802.11e wireless LANs," in *Proc. 18th International Teletraffic Congress (ITC18)*, Sept. 2003, pp. 1001–1010.

[4] H. Wu *et al.*, "Performance of reliable transport protocol over IEEE 802.11 wireless LAN: Analysis and enhancement," in *Proc. IEEE INFOCOM'02*, June 2002, pp. 599–607.