# Applications and Challenges of the 802.11e EDCA Mechanism: An Experimental Study

**Albert Banchs, Arturo Azcorra, Carlos García, and Rubén Cuevas, University Carlos III of Madrid**

**Abstract**

In this article we conduct an experimental study of the enhanced distributed channel access mechanism of the upcoming 802.11e standard. The main focus of the study is on two applications of EDCA: *traffic engineering* and *service guarantees*. With traffic engineering we aim at distributing the bandwidth in the WLAN according to a given throughput allocation criterion. With service guarantees the objective is to ensure that the performance metrics (throughput and delay) experienced by a station meet a given quality criterion. We build a testbed with wireless cards that support a substantial subset of the EDCA functionality and analyze performance against different sets of requirements. Experimental results show that EDCA can effectively be used for traffic engineering purposes. The goal of providing service guarantees with EDCA is shown to be more challenging; we identify future research directions that need to be addressed in order to achieve this goal.

In recent years, much interest has been devoted to the design of wireless local area networks (WLANs) with quality of service (QoS) support. The Enhancements Task Group (TGe) was formed under the IEEE 802.11 project to recommend an international WLAN standard with QoS. This standard is called 802.11e and is being built as an extension of the basic WLAN 802.11 standard. While the standardization process of 802.11e is still an ongoing effort, the main features of the upcoming standard have already been agreed upon and are unlikely to change. These features are described in the latest version of the 802.11e standard draft [1].

The standard draft defines two different access mechanisms: the enhanced distributed channel access (EDCA) and hybrid coordination function (HCF) controlled channel access (HCCA). EDCA is a distributed scheme that extends the distributed coordination function (DCF) of 802.11, while HCCA is centralized. Recently, some commercial WLAN products that implement (to a certain extent) the EDCA mechanism have become available. The product roadmap for HCCA is less clear.

The focus of the present article is on an experimental analysis of the EDCA mechanism with a real-life testbed. To date, a considerable amount of work has been addressed to analyze the performance of this mechanism (e.g., [2]); however, to the authors' knowledge, all previous work in the literature was based on simulations and/or analysis; none investigates the performance of EDCA experimentally. It is well known from experimental studies with DCF [3, 4] that WLANs suffer from a number of nonideal effects that have nonnegligible impact on performance and are not accounted for in analytical and simulation studies. Therefore, experiences with real-life experiments are needed in order to complement previous theoretical studies of EDCA and assess its *real* performance.

The methodology used in this article to perform the experimental analysis of EDCA is the following. We consider two widely accepted applications of a QoS architecture for WLAN, *traffic engineering* and *service guarantees*, and evaluate how well the requirements of these applications are satisfied by our EDCA testbed. The article is structured as follows. We first review the EDCA mechanism of the 802.11e standard draft. Then we describe our testbed setup and discuss some initial experimental results that give an understanding of the various effects observed in the testbed. Next, we take up the issue of traffic engineering with EDCA and study the performance of two different traffic engineering criteria (*symmetric traffic allocation* and *olympic service model*). We then address the provisioning of service guarantees with EDCA and outline some possible future directions of research in this context. Finally, we summarize the results of our work.

## 802.11e EDCA

We now briefly summarize the EDCA mechanism as defined in the latest version of the 802.11e standard draft. For a complete description of the EDCA and HCCA mechanisms of 802.11e, the reader is referred to the standard draft [1].

EDCA controls the access to the wireless channel on the basis of the channel access functions (CAFs). A station may run up to four CAFs, with the frames generated by the station mapped to one of these CAFs. Each CAF executes an independent backoff process to determine the time of transmis-

*The work described in this article is based on results of IST FP6 Integrated Project DAIDALOS. DAIDALOS receives research funding from the European Community's Sixth Framework Program. Apart from this, the European Commission has no responsibility for the content of this article.*

  

sion of its frames. The backoff process is regulated by four configurable parameters: $CW_{min}$, $CW_{max}$, $AIFS$, and $TXOP$ *limit*.

A CAF $i$ with a new frame to transmit monitors the channel activity. If the channel is idle for a period of time equal to its arbitration interframe space ($AIFS_i$) parameter, the CAF transmits. Otherwise, if the channel is sensed busy (either immediately or during the $AIFS_i$ period), the CAF continues to monitor the channel until it is measured idle for an $AIFS_i$ time. At this point, the CAF starts the backoff process by initializing its backoff time counter to a random value uniformly distributed in the range $(0, CW_i - 1)$, where $CW_i$ is the contention window of CAF $i$ and depends on the number of failed transmissions. At the first transmission attempt, $CW_i$ is set equal to the minimum contention window parameter ($CW^i_{min}$).

As long as the channel is sensed idle the backoff time counter is decremented once every slot time, where the slot time duration (hereafter denoted $\sigma$) is a constant defined by the physical layer. When a transmission is detected on the channel, the backoff time counter is "frozen," and reactivated again when the channel is sensed idle for a certain period. This period is equal to $AIFS_i$ if the transmission is received with a correct CRC,[1] AND $EIFS - DIFS + AIFS_i$ otherwise, where the extended interframe space ($EIFS$) and DCF interframe space ($DIFS$) are physical layer constants. As soon as the backoff time counter reaches zero, the CAF transmits its frame in the next slot time. A collision occurs when two or more CAFs start transmission simultaneously. In the case of a single station running more than one CAF, if the backoff time counter of two or more CAFs of the station reach zero at the same time, a scheduler inside the station avoids an *internal collision* by granting the channel access to the highest-priority CAF.

An acknowledgment (ACK) frame is used to notify the transmitting CAF that the frame has been successfully received. The ACK is transmitted at the end of the frame after a period of time equal to the physical layer constant short interframe space (SIFS). If the ACK is not received within a specified ACK timeout, the CAF assumes that the transmitted frame was not received successfully and schedules a retransmission by reentering the backoff process. After each unsuccessful transmission, $CW_i$ is doubled, up to a maximum value given by the $CW^i_{max}$ parameter. After a successful transmission, the CAF is allowed to transmit several consecutive frames, the only restriction being that it cannot occupy the channel for a period of time longer than the transmission opportunity limit parameter (*TXOP limit*$_i$).

In WLANs, the erroneous reception of a frame may be caused by either:

- A collision (i.e., two stations transmit simultaneously and thus interfere with each other)
- The inherent noise and interference of the radio link

The 802.11 a/b/g standards tackle the latter by offering various modulation schemes (with different bit rates and robustness to errors) and using the one that adapts best to the characteristics of the radio channel. Specifically, 802.11b (which is the focus of our experiments) reduces the channel bit rate from 11 Mb/s (the nominal value) to either 5.5, 2, or 1 Mb/s depending on the number of detected errors. The first cause of erroneous reception, a collision between two frames, does not always produce an error. On some occasions, the frame received with the strongest signal survives the collision and is *captured* by the receiving station. This is the so-called *capture effect*.

## Experimental Setup

Although the 802.11e standard has not yet been finalized, there already are some commercial wireless cards that partially support EDCA. Our experiments are conducted with 802.11b wireless cards based on the Atheros AR5212 chipsets. The MADWIFI driver[2] provides full support for wireless adapters using Atheros chipsets on Linux platforms.

Atheros based cards implement a substantial subset of EDCA, but also suffer from some restrictions. Out of the 4 parameters defined in the standard draft ($CW_{min}$, $CW_{max}$, $AIFS$ and *TXOP limit*), only the first 3 are supported. In addition, only one CAF per station is implemented, instead of the 4 defined in the standard draft. We believe, however, that these restrictions do not represent a major limitation for our work. Indeed, the multiple CAFs per station feature is not needed for the applications we target, while the *TXOP limit* provides a level of differentiation similar to the $CW_{min}$ [5].

The configuration of the $CW_{min}$, $CW_{max}$ and AIFS parameters is performed with some functions that the MADWIFI driver provides for this purpose. Specifically, these functions allow configuring each wireless card with three integer values, $i$, $j$ and $k$, with which the $CW_{min}$, CWmax and AIFS parameters of the wireless card are set as follows:

$$CW_{min} = CW^{default}_{min}\, 2^i, \tag{1}$$

$$CW_{max} = CW^{default}_{max}\, 2^j, \text{ and} \tag{2}$$

$$AIFS = \mathrm{DIFS} + k\sigma, \tag{3}$$

where $CW^{default}_{min} = 32$ and $CW^{default}_{max} = 1024$ correspond to the default values of the DCF mechanism of the 802.11b standard. Note that the above restricts the possible $CW_{min}$ and $CW_{max}$ values to the integers power of 2.

To perform our experiments we built a testbed composed of four desktop PCs under Linux equipped with Atheros-based wireless cards. One PC was configured as the access point (AP) and the others as stations under the infrastructure mode; in the rest of the article we refer to them as stations 1, 2, and 3. Unless otherwise stated, experiments were performed with the four PCs located in the same room at a distance ranging from 2 to 4 m from each other. Except for the voice application used in the experiment of Fig. 6, traffic was generated with the *iperf*[3] tool sending (TCP or UDP) 1000-byte packets. The experimental results given are the average of five runs of 60 s duration each. In most cases, the difference between the maximum and minimum results of the five runs falls below 5 percent.

In order to gain understanding of the nonidealities of our testbed, we conducted a first set of experiments. First, by saturating the wireless channel with only one station sending UDP traffic, we determined that the throughput obtained is independent of the station's location. This shows that within the range of distances in our testbed all PCs work at the nominal bit rate of 802.11b (11 Mb/s), and error rates are negligible. By saturating the channel with two stations located at the

[1] *All stations check every received frame (regardless of its destination) for errors, using the cyclic redundancy check (CRC) field of the frame.*

[2] *Multiband Atheros Driver for WiFI (MADWIFI) consists of the following modules:* `hal`, `if_ath`, *and* `ath_pci`. *Our work was performed with* `hal` *v. 0.8.2,* `if_ath` *v. 0.7.0, and* `ath_pci` *v. 0.8.2. New versions of these modules with extended functionality are currently being developed. See http://sourceforge.net/projects/madwifi/*

[3] *http://dast.nlanr.net/Projects/Iperf/*

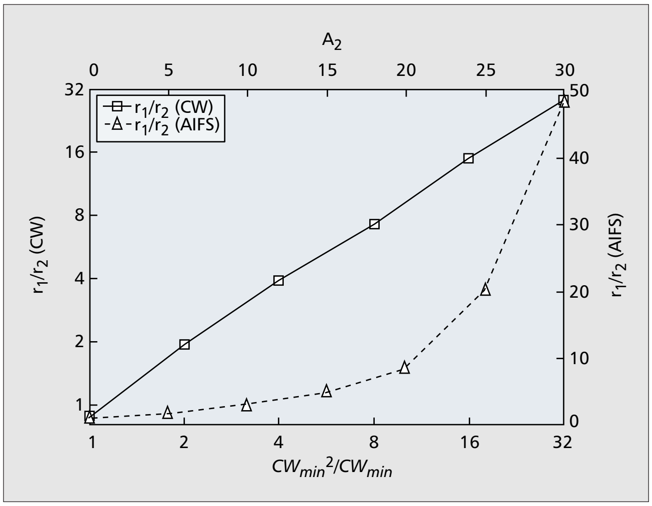


■ Figure 1. $CW_{min}$ *and* AIFS *parameters.*

same (about 2 m) and different (2 and 4 m) distances from the AP, respectively, and observing that the AP only detected errors in the former case, we found that the capture effect occurs within the range of distances of our testbed. Furthermore, by comparing the results of the latter experiment against the analytical model of [6] for the capture effect, we concluded that this is the main cause of the deviations observed in our experimental results.

## EDCA for Traffic Engineering

One of the possible applications of 802.11e EDCA is traffic engineering. With traffic engineering we aim at (by setting appropriately the EDCA parameters of the stations) distributing the throughput of the WLAN according to a traffic engineering criterion. We now study the effectiveness of EDCA for this purpose, taking as an example two different traffic engineering criteria: the *symmetric traffic allocation* and the *olympic service model*. We first propose some configuration guidelines of EDCA for traffic engineering. Next, we assess how well EDCA configured with these guidelines meets the performance goals of the two considered criteria under UDP traffic. Then we take one of the two criteria, the symmetric traffic allocation, and study its performance under TCP traffic. Finally, we study, also with the symmetric traffic allocation criterion, the impact of heterogeneous radio conditions on traffic engineering.

### Parameter Configuration for Traffic Engineering

Most traffic engineering criteria (in particular the ones we use as examples here) are based on *weights*. With these criteria, different stations get assigned different weights; the goal is to provide them with a throughput proportional to the weight.

In order to assess which of the EDCA parameters ($CW$ and $AIFS$) satisfies the above goal better, we conducted the following experiment. We saturated two stations with UDP traffic and studied the ratio of throughputs ($r_1/r_2$) when we increased, respectively, the $CW_{min}$ and $AIFS$ of the second station ($CW^2_{min}$ and $AIFS_2 = DIFS + A_2\sigma$), while leaving all the other parameters at their default value. The results, depicted in Fig. 1, show a linear relationship between $r_1/r_2$ and $CW^2_{min}/CW^1_{min}$, and an exponential one between $r_1/r_2$ and $AIFS_2$, which matches the analysis and simulations of [2].

From the above we conclude that the $CW_{min}$ parameter is the most appropriate for providing a station with throughput proportional to its weight. Specifically, results show that the throughput obtained is approximately inversely proportional to the $CW_{min}$. Following this, we configured the parameters of station $i$ according to the following formulae:[4]

$$AIFS_i = DIFS, \tag{4}$$

$$CW^i_{min} = CW^{default}_{min}\, round\_power2(w_{max}/w_i), \text{ and} \tag{5}$$

$$CW^i_{max} = CW^{default}_{max}\, round\_power2(w_{max}/w_i), \tag{6}$$

where $w_i$ is the weight assigned to station $i$, $w_{max}$ is the highest weight in the WLAN, and $round\_power2(x)$ is the integer power of 2 closest to $x$:

$$round\_power2(x) = \begin{cases} 2^{\lfloor \log_2(x) \rfloor}, & x - 2^{\lfloor \log_2(x) \rfloor} \le 2^{\lceil \log_2(x) \rceil} - x \\ 2^{\lceil \log_2(x) \rceil}, & \text{otherwise} \end{cases}. \tag{7}$$

Note that with the above configuration the station with the highest weight in the WLAN uses the default $CW$ values (i.e., $CW^{default}_{min}$ and $CW^{default}_{max}$), while the remaining stations use higher values. This ensures that efficiency is not degraded due to low $CW$ values yielding an unacceptably high collision rate. In addition, the number of steps between $CW_{min}$ and $CW_{max}$ is kept equal to the default value for all stations; this is done in order to preserve the good features of the exponential backoff increase algorithm. Finally, the $round\_power2$ function is introduced because of the restriction imposed by the wireless cards that $CW$s have to be a power of 2.

### Symmetric Traffic Allocation with UDP Traffic

Throughput in DCF is distributed "fairly" among all contending stations, in the sense that all stations with traffic to send are allocated the same throughput. While at first glance this criterion may seem reasonable, a closer look at the throughput distribution of a WLAN working in the infrastructure mode reveals that, as pointed out in [3], this distribution may be undesirable from a traffic engineering viewpoint.

In the infrastructure mode, stations only send and receive packets from the AP. Therefore, the relevant traffic engineering parameters in this case are the rate at which the stations can send packets (hereafter called the *upstream throughput*) and the rate at which the stations can receive (the *downstream throughput*). With the "fair" distribution of DCF, the AP, which behaves in the same way as any other station as far as the contention algorithm is concerned, is allocated the same throughput for sending packets as any other station. As a result, the sum of all downstream throughputs in the WLAN is equal to one single upstream throughput. In other words, if there are $N$ stations in the WLAN, each station is allocated $N$ times more throughput for upstream than for downstream traffic.

A more reasonable throughput allocation may be to allocate the same throughputs for upstream and downstream traffic (to which we refer as *symmetric traffic allocation*). A typical scenario in which this allocation would be desirable is a WLAN with all stations holding bidirectional voice over IP (VoIP) conversations. Note that to achieve this allocation we need to provide the AP with a throughput $N$ times larger than any other station, $N$ being the number of stations in the WLAN.[5]

---

[4] *While there are other proposals for configuration of the EDCA parameters (e.g., [7]), one of the advantages of the configuration proposed here is its simplicity.*

[5] *In a dynamic environment, the number of active stations in the WLAN can be obtained, for example, by caching the MAC address of the stations that have transmitted packets recently; this functionality is provided by the Wireless Tools (http://www.hpl.hp.com/personal/Jean_Tourrilhes/Linux/Tools.html).*

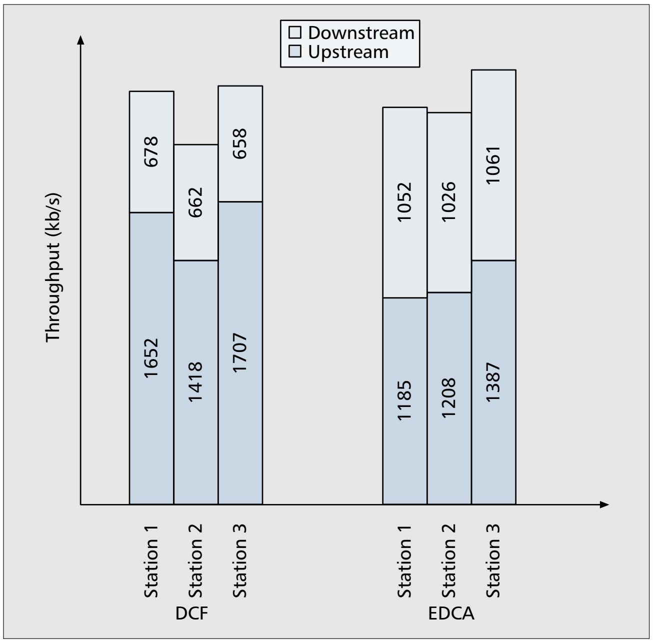


■ Figure 2. *Symmetric traffic allocation with UDP traffic.*

To study the effectiveness of EDCA in providing symmetric throughput allocation, we saturated the upstream and downstream directions of three stations with UDP traffic and examined the resulting throughput allocation for DCF and EDCA. Results for DCF were obtained using the default configuration. The EDCA configuration was derived from the formulae given in Eqs. 4–6 with weight $w_{AP} = 3$ for the AP and weight $w_{STA} = 1$ for the stations. Results are illustrated in Fig. 2. We observe that, as expected, DCF allocates approximately three times more throughput to upstream than to downstream, while, in contrast, our EDCA configuration behaves much better according to the symmetric traffic allocation objective.

It is worthwhile to observe that in both the DCF and EDCA results, the AP receives throughput approximately 20 percent higher than expected with the respective configurations. Note that if a station and the AP collide, the station's transmission will surely result in a failure, as the destination of the transmission (the AP) is not in listen mode. In contrast, there is some probability that the AP transmission results in success as a consequence of the capture effect. This asymmetry in the capture effect explains why in all our experiments the AP is always favored over the stations.

## Olympic Service Model with UDP Traffic

With the olympic service model (OSM) [8] users are classified in three categories (gold, silver, and bronze), and each category is mapped to a weight ($w_{gold} > w_{silver} > w_{bronze}$). Then a gold user is allocated throughput proportional to $w_{gold}$, a silver user throughput proportional to $w_{silver}$ and a bronze user throughput proportional to $w_{bronze}$. This allocation would be the desirable one for, say, a gold user holding a videoconference with high-rate video, a silver user holding a low-rate videoconference, and a bronze user holding an audioconference.

To study the effectiveness of EDCA to provide the OSM, we took $w_{gold} = 4$, $w_{silver} = 2$, $w_{bronze} = 1$, with one station for each category, and aimed to provide each station with balanced downstream and upstream throughputs, both proportional to the weight of the station's category. For this purpose we configured EDCA with the formulae given in Eqs. 4–6, each station with the weight of its category and the AP with a weight equal to $w_{AP} = w_{gold} + w_{silver} + w_{bronze}$ (i.e., the sum of the weights of all stations). In addition, we configured an internal Weighted Fair Queuing (WFQ) scheduler at the AP in order to distribute the downstream throughput among the competing flows proportional to their weights.

Figure 3 illustrates the throughput allocations resulting from the above configuration with the three stations saturating both the upstream and downstream directions with UDP traffic. Results confirm the effectiveness of EDCA to implement the OSM; upstream and downstream throughputs are approximately proportional to the weights. Downstream throughputs are about 30 percent larger than upstream; this is a consequence of the asymmetry in the capture effect discussed for the previous experiment as well as the error introduced by the rounding operation of Eq. 7.

## Traffic Engineering with TCP Traffic

Traffic engineering for TCP traffic becomes more challenging as a consequence of the well-known interactions with the congestion control algorithm of TCP and the presence of TCP ACKs[6] [3]. In order to gain insight into the impact of TCP on throughput distribution, we repeated the symmetric trafic allocation experiement of Fig. 2 with each station running $n$ long-lived TCP flows in both upstream and downstream directions. Experiments were performed with both the default output queue buffer size in all PCs ($B = 199$ packets) and (following the recommendations of [3]) an output queue buffer size of $B_{AP} = NB$ at the AP, $N$ being the number of stations in the WLAN (in our case $N = 3$).

Figure 4 shows the ratio of upstream and downstream throughputs resulting from the above scenario for DCF and EDCA and the two buffer configurations. According to our objective of a symmetric traffic allocation, this ratio should ideally be equal to 1. Results show substantial differences with the allocation obtained for UDP traffic in Fig. 2. For $n = 1$, TCP balances the throughput obtained by all flows independently of the EDCA parameters and buffer configuration. For

[6] *Note that throughput in EDCA is regulated by controlling the number of times a station accesses the channel, and from this perspective, a TPC ACK counts as much as any other packet independent of its size.*

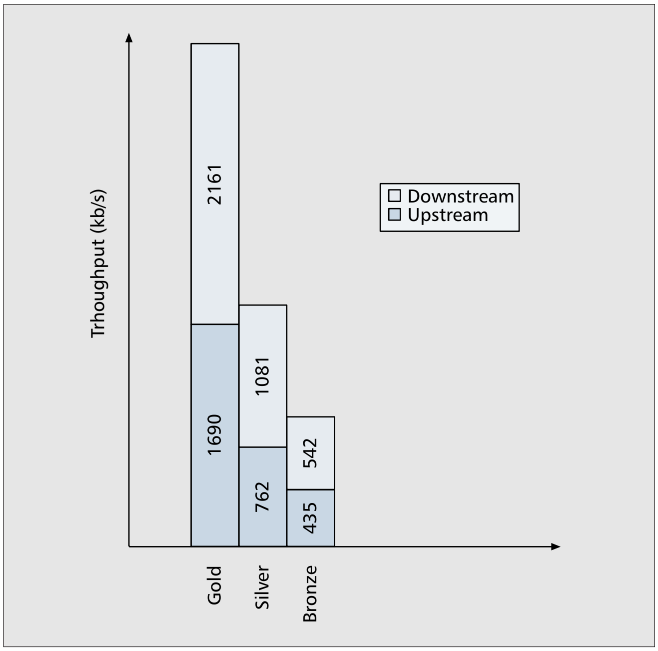


■ Figure 3. *Olympic service model with UDP traffic.*

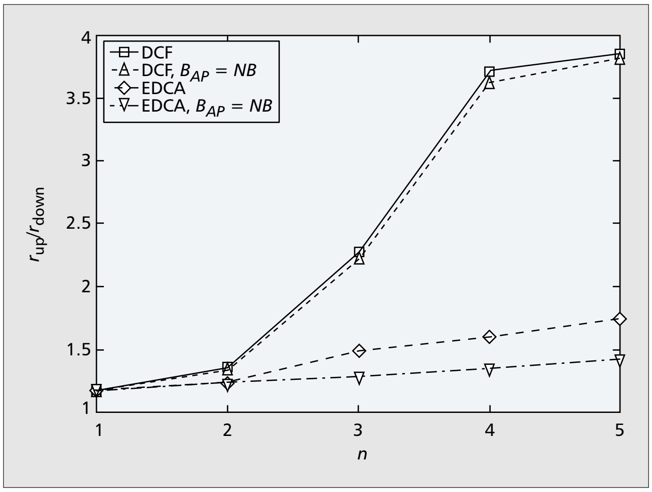


■ Figure 4. *Symmetric traffic allocation with TCP traffic.*

a larger number of flows, the distribution becomes unbalanced with DCF, while with EDCA (both when $B_{AP} = B$ and $B_{AP} = NB$) the ratios keep relatively close to 1. We conclude that the configuration of EDCA parameters is also useful with TCP traffic to keep the throughput distribution close to the desired allocation. We believe, however, that more extensive studies are needed to better understand the interactions between TCP and EDCA.

### Traffic Engineering Under Heterogeneous Radio Conditions

The fact that throughputs in WLAN depend on radio conditions poses an additional challenge to traffic engineering. Radio conditions may impact the throughput of a station in 802.11b WLAN in the following ways:

- Differences in the receive power of two stations may yield the capture effect.
- Some degradation of the radio link may yield occasional transmission errors.
- A severe degradation of the radio link yields frequent transmission errors, which triggers the selection of a new modulation scheme more robust to errors and with a lower transmission bit rate.

In order to illustrate the impact of radio conditions on throughput, we repeated the symmetric traffic allocation experiment of Fig. 2 with two stations sending UDP traffic upstream, the first station placed close (2 m) to the AP and the second one placed in the following locations:

- At a distance from the AP similar to the first station (2 m)
- At a longer distance (4 m) but still with a good radio link
- Outside the room, where the bit rate remains at 11 Mb/s
- At about 100 m, where the transmission bit rate decreases to 2 Mb/s

Results of the total throughput in the WLAN ($r$) and the ratio of throughputs ($r_2/r_1$) as a function of the location of the second station are given in Fig. 5. We observe that in position 1 both stations obtain approximately the same throughput. In position 2, the capture effect leads to the first station obtaining greater throughput (about 10 percent more) than the second station. In position 3 transmission errors occur, which slightly degrade the throughput of the second station and the total throughput; however, the impact is almost imperceptible. Finally, in position 4 the total throughput is drastically degraded as a result of the second station transmitting at a lower bit rate, but the ratio of throughputs does not change.

We conclude from the results that radio conditions have little impact on the ratio of throughputs; indeed, in our experiments the ratio is always close to 1, which is the ideal value according to the symmetric traffic allocation criterion. Another conclusion drawn from the results is that when a station decreases its transmission bit rate (position 4), although the ratios are preserved, the throughputs of all the stations in the WLAN are sharply degraded. Indeed, if a given station transmits at a lower bit rate, each time it accesses the channel it occupies it for a longer period; therefore, less channel time is available for other stations. This behavior is studied in detail in [4].

## EDCA for Service Guarantees

We now address the issue of providing service guarantees with EDCA and analyze the challenges that still have to be solved in order to reach this goal. To support the explanations provided here, we performed a number of experiments (Fig. 6) whose setup is described in the following paragraphs.

In experiment 1, station 1 and the AP held a bidirectional voice conversation using the GnomeMeeting[7] tool, which measures the round-trip delay. Stations 2 and 3 sent background data traffic of 1000-byte packets at a rate of $r_{background}$. In order to saturate the channel, stations were manually configured at a low transmission bit rate; specifically, the AP and stations 1 and 2 were configured at a bit rate of 1 Mb/s, and station 3 at a bit rate of 2 Mb/s. All stations had the default configuration (i.e. the one of DCF). We measured the voice traffic round-trip delay as a function of the background stations sending rate, and considered (as an example of quality criterion) that voice performance is acceptable as long as this delay falls below 50 ms (this is represented by the horizontal straight line in Fig. 6).

In experiment 2 we repeated the above setup but using EDCA. Following the results of Fig. 1, which show that the *AIFS* parameter is more appropriate for providing strict priorities, we configured real-time stations with the minimum AIFS allowed by the standard draft ($AIFS = DIFS$) and background stations with the maximum allowed AIFS ($AIFS = DIFS +$

[7] *http://www.gnomemeeting.org/*

[8] *In the recommendations given by the standard draft, real-time and background stations are also configured with different* $CW_{min}$ *and* $CW_{max}$ *parameters. In this article we restrict our analysis to only one parameter for simplicity.*

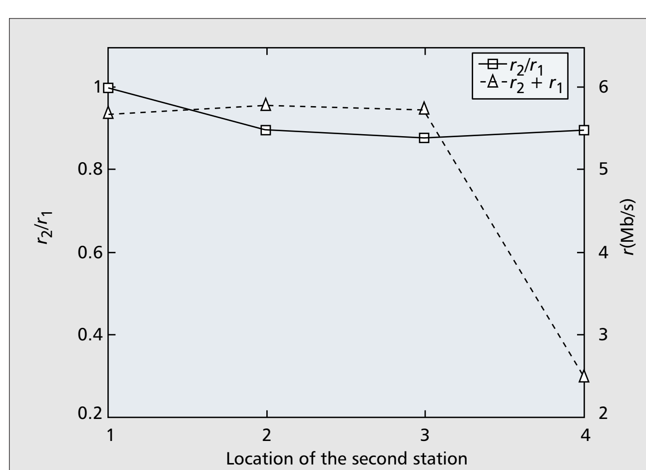


■ Figure 5. *Heterogeneous radio conditions.*

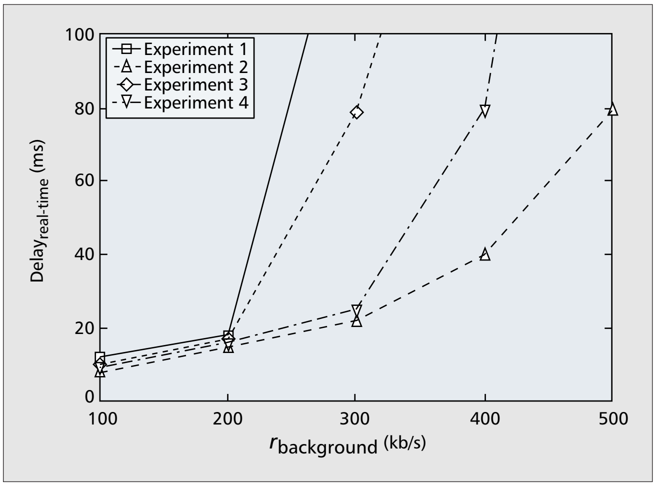


■ Figure 6. *Experiments with real-time traffic.*

13σ), and left the other parameters at their default values.[8]

In experiments 3 and 4 we repeated experiment 2 introducing the following variations. In experiment 3, the bit rate of station 3 was reduced to 1 Mb/s. In experiment 4, station 1 was moved about 2 m further from the AP than the other stations, such that the link quality was still good but not as good as the others.

### Analysis and Algorithms: Future Directions

The provisioning of service guarantees with EDCA requires the development of a theoretical basis that is not yet complete. Indeed, QoS architectures for wired links such as integrated services (IntServ) and differentiated services (DiffServ) rely on analyses (e.g., [9]) that allow performance predictions and can be used to guarantee QoS via an admission control algorithm: if accepting a new request would degrade the performance of admitted requests below the guaranteed metrics, the algorithm rejects the new request. The development of these analyses and algorithms for EDCA is one of the research challenges that remain yet to be addressed before this mechanism can be used for providing service guarantees.

Let us look at experiment 2 of the above set. Given the conditions of the experiment and the requirement that the round-trip delay for the voice applications cannot be greater than 50 ms, the maximum sending rate that can be admitted for the background stations is 400 kb/s. This value is higher than that obtained with DCF (experiment 1), where the maximum admissible sending rate is 200 kb/s. Nonetheless, an algorithm is still needed in EDCA in order to derive the maximum admissible sending rate and perform admission control accordingly.

In the last years considerable effort has been dedicated to analyze the theoretical performance of EDCA; however, there still exist a number of gaps that need to be filled before analysis can be used to provide service guarantees. In the following we outline future research directions to fill these gaps:

- **Analyses with realistic source models**. Present analyses of EDCA are restricted to the unrealistic case of saturated sources (i.e., they assume that all stations always have packets to transmit, e.g., [2]). Even for DCF (the basic mechanism that EDCA extends), although there are some analyses that consider nonsaturation conditions, no complete analyses with realistic source models exist (e.g., [5] is restricted to Poisson arrivals).
- **Analyses of the end-to-end delay distribution**. Real-time applications need not only a low average delay but also a low delay for most of their packets, which requires the computation of the end-to-end delay distribution. Existing EDCA analyses are restricted to average delay [2]. Even for DCF, analyses of the delay distribution are limited (e.g., the analysis of [10] is restricted to saturation conditions).
- **Algorithms to derive the optimal configuration**. The EDCA configuration proposed here as well as the one given by the standard draft recommendations is derived heuristically and guarantees no optimized performance. Instead, it would be desirable to use the optimal EDCA configuration (i.e., the configuration that, given a set of requests and requirements, provides the best possible performance). Existing algorithms to compute the optimal EDCA configuration [7] are restricted to saturation conditions and throughput, and do not serve the purpose of providing real-time applications with service guarantees.

### Inherent Uncertainties with Mobility: A New Concept for Admission Control

Most of the analyses referenced above assume that all stations transmit at the nominal bit rate (11 Mb/s) under ideal conditions (no transmission errors and no capture effect). These nonideal aspects, as shown by experiments 3 and 4, strongly impact performance and therefore must be taken into account to provide service guarantees.

Although there are some analyses that account for the above aspects (e.g., [11] provides an analysis for heterogeneous transmission bit rates and errors, and [6] provides an analysis for the capture effect), these analyses require the knowledge of some data (e.g., the receive power) that is not always easy to obtain. In fact, the receive power depends not only on the transmit power and distance, but also on some additional effects like fading that are not easy to derive.

Under static conditions (e.g., a company network based on WLAN where the stations are in fixed locations), it may be conceivable to obtain, from the location of each station and the measured quality of the corresponding radio link, the data necessary to evaluate the impact of the above aspects and use this data in the analyses and algorithms that provide service guarantees.

The goal of providing service guarantees, however, becomes more challenging when considering a mobile environment, in which radio conditions vary with time. The problem in this case arises from the fact that the conditions that yield a given admission control decision at one point may not hold after some time. This can be seen in the experiments of Fig. 6: under the conditions of experiment 2, background stations with a sending rate of 400 kb/s may be admitted. However, if the transmission bit rate of the second background station is reduced to 1 Mb/s (experiment 3) or the real-time station moves a bit further (experiment 4), the performance goals are no longer met, and the admission control decision needs to be reevaluated.

We conclude that under mobility conditions a new concept for admission control is required that constantly monitors the channel conditions and reevaluates admission control decisions. Such an algorithm is a subject of future research.

## Conclusions

In this article we have investigated the ability of the EDCA mechanism to satisfy, in a real environment, the requirements of two of the applications for which this protocol was designed: traffic engineering and service guarantees. The investigation was performed based on real-life experiments that have given us insight into a number of aspects not typically shown by simulations and analyses.

Results for traffic engineering show that this application is well supported by EDCA when only UDP traffic is present in the WLAN. For the two traffic engineering criteria studied (symmetric traffic allocation and olympic service model) the results obtained match, with reasonable accuracy, the desired allocations. Only some small deviations are observed because of the capture effect and the restriction that *CW* values must be set to integer powers of 2. Results with TCP are promising; although the interaction between the TCP congestion control and the EDCA mechanism causes some deviations from the desired allocations, the throughput distributions we have obtained are fairly good and much better than the ones obtained with DCF. Experiments under heterogeneous radio conditions reveal that transmission errors and the capture effect have a small influence, while the stations' bit rates have a much stronger impact on traffic engineering.

The other application we have analyzed, service guarantees, is much harder to satisfy. A substantial amount of theoretical work still needs to be developed before this application can be implemented with EDCA. This theoretical work includes analyses under nonsaturation conditions, analyses of the distribution of the end-to-end delay, and algorithms to determine the optimal EDCA configuration. In addition, our experiments show that the inherent uncertainties of a mobile environment represent a major challenge for service guarantees. In order to tackle this problem, we propose the design of a new concept for admission control that constantly monitors the WLAN conditions and reacts to changing conditions.

As a final conclusion, we believe that the experimental results reported in this article show that the EDCA mechanism has many useful applications that would not be possible to satisfy with the legacy DCF protocol, but there are still some challenges that need to be addressed before EDCA can meet all the design goals for which it was originally conceived.

## Acknowledgments

The authors would like to thank José Félix Kukielka, Christos Douligeris, and anonymous referees for their insightful comments on earlier versions of this article.

## References

[1] IEEE 802.11e/D13.0, Part 11, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications: Medium Access Control (MAC) Enhancements for Quality of Service (QoS)," draft supp. to IEEE 802.11 std., Jan. 2005.

[2] Z. Kong *et al.*, "Performance Analysis of IEEE 802.11e Contention-Based Channel Access," *IEEE JSAC*, vol. 22, no. 10, Dec. 2004, pp. 2095–2106.

[3] S. Pilosof *et al.*, "Understanding TCP fairness over Wireless LAN," *Proc. IEEE INFOCOM '03*, Mar. 2003, pp. 863–72.

[4] M. Heusse *et al.*, "Performance anomaly of 802.11b," *Proc. IEEE INFOCOM '03*, Mar. 2003, pp. 836–43.

[5] G. R. Cantieni *et al.*, "Performance Analysis under Finite Load and Improvements for Multirate 802.11," to appear, *Elsevier Comp. Commun.*

[6] R. Litjens *et al.*, "Performance Analysis of Wireless LANs: An Integrated Packet/Flow Level Approach," *Proc. 18th Int'l. Teletraffic Cong.*, Sept. 2003, pp. 931–40.

[7] A. Banchs, X. Perez, and D. Qiao, "Providing Throughput Guarantees in IEEE 802.11e Wireless LANs," *Proc. 18th Int'l. Teletraffic Cong.*, Sept. 2003, pp. 1001–10.

[8] C. Dovrolis and P. Ramanathan, "A Case for Relative Differentiated Services and the Proportional Differentiation Model," *IEEE Network*, vol. 13, no. 5, Sept. 1999, pp. 26–34.

[9] J. C. R. Bennett *et al.*, "Delay Jitter Bounds and Packet Scale Rate Guarantee for Expedited Forwarding," *IEEE/ACM Trans. Net.*, vol. 10, no. 4, Aug. 2002, pp. 529–40.

[10] A. Banchs, "Analysis of the Delay Distribution in 802.11 DCF: A Step Towards Delay Guarantees in WLANs," *LNCS*, vol. 3266, Sept. 2004, pp. 64–73.

[11] H. Pan, S. Sato, and K. Kobayashi, "On the Throughput of an IEEE 802.11a Wireless LAN System with Terminals under Heterogeneous Radio Conditions," *Proc. 18th Int'l. Teletraffic Cong.*, Sept. 2003, pp. 991–1000.

## Biographies

ALBERT BANCHS [M] (banchs@it.uc3m.es) received his M.Sc. and Ph.D. degrees in telecommunications from the Technical University of Catalonia in 1997 and 2002, respectively. His Ph.D. received the national award for best thesis on broadband networks granted by the Professional Association of Telecommunication Engineers. He worked for the International Computer Science Institute in 1997, for Telefonica I+D in 1998, and for NEC Network Laboratories from 1998 to 2003. Since 2003 he is with University Carlos III, Madrid. His current research interests include resource allocation, QoS, and performance evaluation of wireless and wired networks.

ARTURO AZCORRA [SM] (azcorra@it.uc3m.es) received an M.Sc. degree in telecommunications from the Technical University of Madrid in 1986, and a Ph.D. degree from the same university in 1989. His Ph.D. received the national award for best thesis, jointly granted by the Professional Association of Telecommunication Engineers and the National Association of Electronic Industries. In 1993 he obtained an M.B.A. from Instituto de Empresa. He was a lecturer at the Technical University of Madrid from 1987 to 1990, when he was promoted to associate professor. In 1998 he joined University Carlos III, Madrid, where he has stayed to the present as a full professor. Since 1998 he is also Deputy Vice-Provost for Academic Infrastructures at University Carlos III. He has been program committee member of several editions of IEEE PROMS, IDMS, QofIS, CONEXT, and IEEE INFOCOM. He is general co-chair for CONEXT '05 and TPC co-chair for INFOCOM '06. His list of papers published in national and international journals, books, and conferences is over 100 titles. Current research projects include broadband networks, multicast teleservices, active networks, and advanced IP networks.

CARLOS GARCÍA (cgg@it.uc3m.es) received an M.Sc. degree in telecommunications from University Carlos III, Madrid, in 2001. His M.Sc. thesis received the award for best thesis granted by the Nortel Networks Chair. He is currently a Ph.D. student and lecturer at University Carlos III. His research interests are QoS provision in 4G networks and IP over DWDM network deployment.

RUBÉN CUEVAS (rcuevas@inv.it.uc3m.es) is pursuing his M.Sc. degree in telecommunications at University Carlos III. In 2001 he received a scholarship from Telefonica I+D to work on VPN applications. Since 2002 he is involved in the development of a QoS management platform for wireless networks as part of his M.Sc. thesis.