ACCEPTED FROM OPEN CALL

IEEE 802.11E QOS AND POWER SAVING FEATURES OVERVIEW AND ANALYSIS OF COMBINED PERFORMANCE

XAVIER PÉREZ-COSTA AND DANIEL CAMPS-MUR, NEC LABORATORIES EUROPE

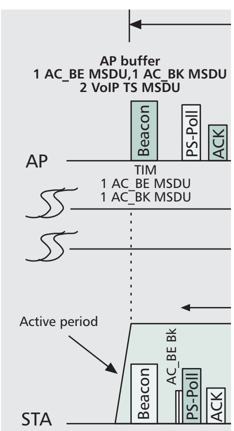

The authors provide an overview of the two power saving modes defined by IEEE 802.11e, unscheduled, and scheduled automatic power save delivery, followed by an evaluation of the performance of the different possible combinations of IEEE 802.11e QoS and power saving mechanisms

ABSTRACT

Wireless LAN has become the ubiquitous connectivity solution for computing in the home and hotspot environments. While the technology continues to grow within its core computer market, it is also steadily expanding into the much larger handset and consumer electronics markets. Given the increasingly diverse range of applications, additional requirements need to be met by products in order to ensure user satisfaction. Within these requirements two are commonly shared by most new devices: QoS support for prioritizing real-time services over non real-time and power saving functionality to achieve an operating time meeting users' expectations. In this article we first provide an overview of the two power saving modes defined by IEEE 802.11e, unscheduled, and scheduled automatic power save delivery, followed by an evaluation of the performance of the different possible combinations of IEEE 802.11e QoS and power saving mechanisms compared to the legacy 802.11 medium access protocol and power save mode. Our results show the level at which the 802.11e QoS and power saving mechanisms meet their design objectives, and provide quantitative results on the differences to be expected with respect to QoS, power saving, signaling load, and network capacity in a generic scenario.

INTRODUCTION

Wireless LAN (WLAN) [1] technology has become the de facto standard for wireless broadband Internet access in homes, offices, and public hotspot locations. In addition, during the last years several initiatives have been started around the world to offer wireless broadband access outdoors in city centers and other specific areas using this technology. The increasing availability of WLAN access points together with the high-speed data rate supported, its cheap cost per bit, and the fact that almost all laptops shipped today have a built-in WLAN interface have provided a constant boost to the popularity of the technology. As a result, there is an increasingly stronger trend toward the inclusion of WLAN capabilities in all kind of devices. On the computing side, notebooks, access points, and residential gateways are still growing at a solid pace. On the embedded side, mobile phone, consumer electronics, and computer peripheral manufacturers have all begun to integrate WLAN into their devices. Overall, market forecasts predict WLAN units to grow roughly 20 percent per year, reaching a billion units by 2012 [2]. Rising penetration of broadband access worldwide is certainly a catalyst for such growth.

The wide range of devices in the WLAN market target different customer needs but two main requirements are commonly shared by most new devices: quality of service (QoS) support for prioritizing real-time services over non realtime, and power saving functionality to achieve an operating time meeting users' expectations. A representative example of a WLAN product that needs to meet these two requirements is smartphones. Such mobile phones support applications like voice over IP (VoIP), video/music streaming, and web browsing, and therefore need to prioritize their traffic in a way that ensures a reasonable QoS. Additionally, given their small device size, which severely limits battery capacity, the QoS mechanisms need to be compatible with the power saving ones such that both requirements can be met at the same time. Other relevant examples of WLAN products where the QoS and power saving requirements are becoming increasingly important are game consoles and media players, home gateways, set-top boxes, displays, projectors, and wireless docking stations. It is worth noting that although some of these devices are plugged to the power grid, the steadily rising energy cost and pressure to reduce global CO_2 emissions has resulted in an emerging awareness of the energy consumption of all kinds of devices, pushing new models to become *greener*.

The QoS requirement can be addressed with the 802.11e [3] amendment of the IEEE 802.11 standard. Two channel access methods are defined: a contention-based channel access method called enhanced distributed channel access (EDCA) and a contention-free channel access referred to as HCF controlled channel access (HCCA). EDCA provides support for traffic prioritization based on four QoS classes in a way similar to differentiated services (DiffServ) [4], whereas HCCA provides support for parameterized QoS in a way similar to integrated services (IntServ) [5]. Access points (APs) and stations (STAs) including 802.11e functionality are denoted in the 802.11e standard as QoS APs and QoS STAs. For the sake of readability, in this article we refer to them simply as APs and STAs.

Traffic prioritization in EDCA is achieved through varying the amount of time a station senses the channel to be idle before backing off or transmitting (arbitration interframe space, AIFS) the length of the contention window (CW) to be used for the backoff process, and the duration a station may transmit after it acquires channel access (transmission opportunity, TXOP). This is realized by the introduction of four access categories (ACs) and multiple independent backoff entities. The four ACs are AC_VO, AC_VI, AC_BE, and AC_BK. The labels correspond to the applications for which they are intended: voice, video, best effort, and background.

HCCA is based on a polling mechanism. A station may have several traffic streams (TSs) characterized by a traffic specification (TSPEC). Relevant examples of TSPEC fields are mean data rate, maximum service interval and nominal service data unit size. If a traffic stream is accepted, the AP is required to honor the negotiated TSPEC by reserving a time interval, for that specific stream in the resource allocation schedule. This time interval is called an HCCA TXOP. Stations are informed about the reserved TXOP with a frame called CFPoll after the channel is sensed to be idle for a specific time period, which is shorter for HCCA than EDCA. A station receiving a CFPoll sends its frames to the AP after waiting for a short interframe space (SIFS). During a TXOP only the station that received the CFPoll can transmit frames.

With respect to the battery consumption requirement, the WLAN standards provide different solutions to extend the battery lifetime. IEEE 802.11 defines a power save mode (PSM) which allows WLAN devices to enter into a low power consumption state by buffering frames directed to these stations at the AP while they are saving energy. Once every beacon interval the access point sends a beacon indicating whether or not a certain station has any data buffered at the AP. Wireless stations wake up to listen to beacons at a fixed frequency (Listen Interval) and poll the AP to receive the buffered data by sending power save polls (PS-Polls). Whenever the AP sends data to a station, it indicates whether or not there are more data frames outstanding, using the More Data bit in the data frames, and a station goes to sleep only when it has retrieved all pending data. Although this mechanism significantly alleviates the power consumption problem, a dependency between the data frames medium access control (MAC) downlink delay (AP to station) and the listen interval is introduced. Consequently, some listen interval values can result in downlink delays that are unacceptable for certain QoS-sensitive applications (e.g., VoIP).

IEEE 802.11e defines an enhancement of the 802.11 power save mode, automatic power save delivery (APSD), which takes advantage of some of its QoS mechanisms and reduces the required signaling load. Two APSD modes are available: unscheduled APSD (U-APSD) and scheduled APSD (S-APSD). In the next section an overview of both U-APSD and S-APSD is provided.

Since the IEEE 802.11e QoS mechanisms have been extensively described and studied in the literature (e.g., [6, 7]), while the power saving ones have received much less attention (e.g., [8]), our contributions in this article consist of providing an overview of the two power saving modes defined by IEEE 802.11e, U-APSD and S-APSD, and evaluating the performance differences to be expected between the different possible combinations of the IEEE 802.11e QoS and Power Saving mechanisms. For comparison reasons we also include in the study the baseline 802.11 access mechanism distributed control function (DCF) and its power saving mode, which are currently the most commonly used. The objective of these contributions is to provide guidance to network operators when considering which option would better meet their users' needs.

In our previous work we studied in [9] the impact of using the 802.11 power save mode in combination with the 802.11e EDCA QoS mechanism, and designed algorithms for improving the QoS and power saving of 802.11e stations using 802.11 power save mode [10], U-APSD [11, 12], and S-APSD [13]. The article at hand extends our research in this area by summarizing in a comprehensive way our previous work plus additionally considering the configuration of EDCA in combination with SAPSD. To the best of the authors' knowledge there is no published related work analyzing the performance differences of all possible combinations of the 802.11e QoS and power saving mechanisms.

Throughout the article a basic knowledge of the 802.11e QoS mechanisms and 802.11 power save mode is assumed. For a detailed description about their functionality please see [6] for 802.11e HCCA and EDCA, and [7] for 802.11 power save mode.

APSD

APSD is the proposed 802.11e extension of the 802.11 power save mode. In line with the 802.11e QoS mechanisms, which define the *distributed* EDCA access to provide *prioritized* QoS guarantees and the *centralized* HCCA access to provide *parameterized* QoS guarantees, APSD defines a distributed power saving scheme, U-APSD, and a centralized one, S-APSD. Frames buffered at the AP can then be delivered to power saving stations by using either the EDCA access method if U-APSD is selected, or EDCA or HCCA if S-APSD is chosen. The period of time where a station is awake receiving frames delivered by the AP is defined in APSD as the service period (SP). An SP is started by a station or an AP

In line with the 802.11e QoS mechanisms, which define the distributed EDCA access to provide prioritized QoS guarantees and the centralized HCCA access to provide parameterized QoS guarantees, APSD defines a distributed power saving scheme, U-APSD, and a centralized one, S-APSD.

U-APSD uses data frames sent in the uplink as indications of power saving stations being awake. S-APSD APs schedule the instants where S-APSD sations should awake to receive the frames buffered at the AP.

depending on the considered APSD mechanism and is always finished by the reception at the station of a frame with the end of service period flag (EOSP) set.

In the following we provide an overview of the U-APSD and S-APSD functionality indicating the main enhancements introduced with respect to legacy 802.11 power save mode.

U-APSD

The main novel idea behind the U-APSD design is the usage of data frames sent in the uplink by stations (STA → AP) as indications (*triggers*) of the instants when power saving stations are awake. When such an indication is received at the AP from a power saving station, the AP takes advantage of it for delivering data frames buffered while the station was in sleep mode. Because of this specific functionality, this method is especially suited for bidirectional traffic streams even though it provides alternative methods for its usage in other cases.

An *unscheduled* SP begins when an AP receives a *trigger* frame, QoS data or QoS null,[1] from a station, and ends when the station receives a QoS Data or QoS Null frame with the EOSP set. For all frames except the final frame of the SP, the EOSP subfield of the QoS control field of the QoS data frame shall be set to 0 to indicate continuation of the SP. During an SP one or more data frames of delivery-enabled ACs might be delivered by the AP to a station up to the number of frames indicated in the *maximum service period length* (`Max_SP_Length`) value following the rules of an acquired transmission opportunity. `Max_SP_Length` is a field contained in the QoS Info field filled by the station at association that can indicate delivery of 2, 4, 6, or all buffered frames during a SP.

Four access categories (ACs) are defined for EDCA (AC_VO, AC_VI, AC_BE, and AC_BK), which correspond to the type of applications for which they are intended (i.e., voice, video, best effort, and background). Each AC of a station can be configured separately to be delivery/trigger-enabled. If one or more ACs of a station are *trigger-enabled*, when the AP receives a frame of subtype QoS data or QoS null of a trigger-enabled AC, an SP is started if one is not in progress. If a station has one or more ACs configured as delivery-enabled, when this station starts an SP the AP delivers the buffered frames corresponding to these ACs using EDCA. The configuration at the AP of the different ACs per station as delivery/trigger-enabled can be performed either at association time or through the usage of the traffic specification (TSPEC) element info field of the add traffic stream (ADDTS) frames.

In the case of a station that has no data frame to transmit in the uplink, QoS null frames[1] can be sent by the station to request the delivery of the frames buffered at the AP. This enables the usage of U-APSD by an AC of a station that does not generate uplink traffic often enough to meet the downlink QoS requirements of an application using this AC.

In order to guarantee backward compatibility of legacy stations that do not support APSD, the procedure of the AP to assemble the traffic indicator map (TIM) has been modified in such a way that if at least one of the ACs is non-delivery-enabled, it indicates the buffer status *only* of the non-delivery-enabled ACs. Note that in this case it means that the beacon will not indicate whether frames of delivery-enabled ACs are buffered. Only in the case that *all* ACs of a station are delivery-enabled does the TIM indicate the buffer status of delivery-enabled ACs.

Figure 1 provides an example of the operation of U-APSD. For further details about the U-APSD functionality, the reader is referred to [3].

S-APSD

The main idea behind the S-APSD design is the *scheduling* by the AP of the instants where each different station using SAPSD should awake to receive the frames buffered at the AP.

The usage by a station of the S-APSD delivery mechanism for a traffic stream where the access policy is HCCA, or for an AC where the access policy is EDCA, is configured by the transmission of an ADDTS request frame to the AP with the APSD and Schedule Subfields of the TS Info field element both set to 1. In case the AP can satisfy the requested service, it will indicate so in the schedule element of the response, which will include the *service start time* (SST) and the *service interval* (SI).

The AP is responsible for defining for each traffic stream or AC of a station using S-APSD, the SST, and the SI necessary for the periodical scheduling of the delivery of frames to the stations. If a station has set up S-APSD for a traffic stream or an AC, it shall automatically wake up at the scheduled time of each SP defined by $SST + N \times SI$, where $N \in \mathbb{Z}^+$.

The AP may update the service schedule at any time by sending a schedule element in a schedule frame which will be used once the acknowledgment to the schedule frame has been received. The new SST though, shall not exceed the beginning of the previous SP by more than a defined maximum service interval and shall not precede the beginning of the immediately previous SP by more than a defined minimum service interval. A station can also modify the S-APSD service schedule by modifying or deleting its existing traffic specification through ADDTS or DELTS messages.

As in the case of U-APSD, a station shall remain awake until it receives a frame with the EOSP subfield set to 1. If necessary, the AP may generate an extra QoS null frame with the EOSP set to 1. In Fig. 2 an example of the S-APSD operation is provided.

For further details about the S-APSD functionality, the reader is referred to [3].

Performance Evaluation and Discussion

In this section we evaluate the performance of the different possible combinations of IEEE 802.11e QoS and power saving modes as compared to the 802.11 legacy access and power save mode: HCCA+S-APSD, EDCA+S-APSD, EDCA+U-APSD, DCF+U-APSD, and

[1] *QoS nulls are the substitute in U-APSD of 802.11 power save mode PS-polls.*

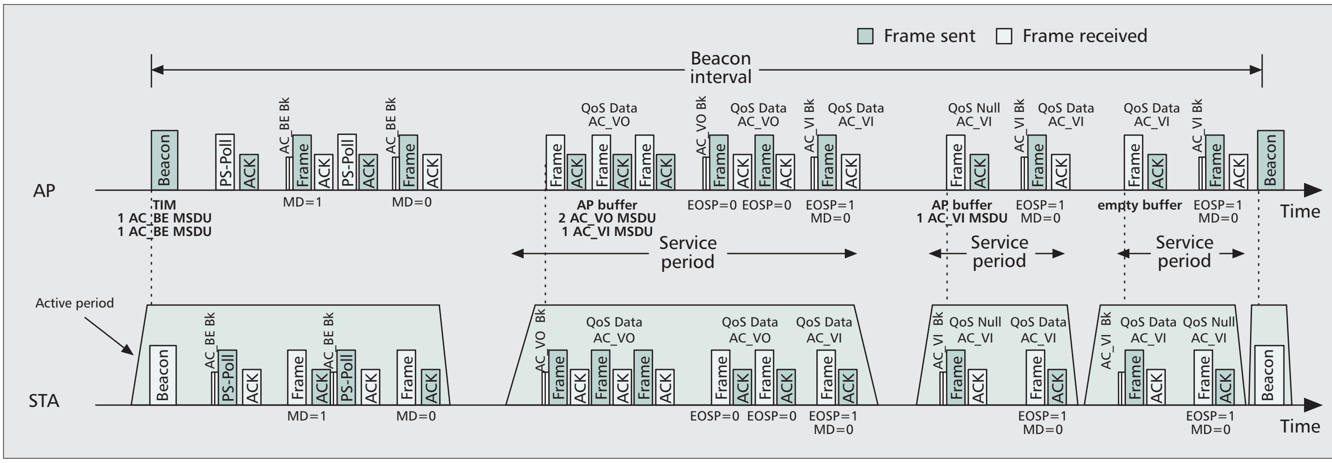

Figure 1. *Example of U-APSD operation. U-APSD configuration: AC VO and AC VI both trigger-and delivery-enabled. AC BE and AC BK neither delivery- nor trigger-enabled (i.e., use legacy 802.11 power save mode). 1) First active period of the station: usage of the legacy power save mode mechanism to retrieve frames of AC_BE and AC_BK categories. 2) Second active period: station wakes up because it has to send an AC VO QoS data frame in the uplink. The first AC_VO frame received by the AP acts as a trigger and starts an unscheduled SP. 3) Third active period: station wakes up because its U-APSD algorithm decides that an SP has to be started. In this case the STA sends a QoS null frame to trigger the unscheduled SP. 4) Fourth active period: station wakes up because it has an AC_VI data frame to be sent and starts an unscheduled SP.*

DCF+PSM. The objective of this study is to illustrate the quantitative differences that might be expected with respect to QoS, power saving, signaling load, and network capacity in a generic scenario. The analysis is performed via simulation. We extended the WLAN libraries provided by OPNET [14] to include 802.11 legacy power save mode, and the QoS and power saving mechanisms of 802.11e.

In the evaluation we study the impact of increasing the number of stations on the selected performance metrics. We define a basic *cluster* composed of four WLAN stations where each station is configured to send and receive traffic from their corresponding pair in the wired domain of its type of application; that is, one station sends and receives voice traffic, a second one receives a video stream, a third one does web browsing, and a fourth one does an FTP download. The number of stations increases in multiples of clusters starting from three clusters (12 stations) and ending with 63 clusters (252 stations) always keeping the relation of 1/4 stations of each application type. Within each group of stations of the same application type, one-third of the stations operate at 54 Mb/s, one-third at 24 Mb/s, and one-third at 12 Mb/s. All stations use the 802.11g PHY layer.

The configuration used for the different applications is detailed below:

- **Voice**: G.711 voice codec with silence suppression. Data rate: 64 kb/s. Frame length: 20 ms. Talk spurt exponential with mean 0.35 s and silence spurt exponential with mean 0.65 s.
- **Video**: MPEG-4 real traces [15]. Target rate: 230 kb/s. Peak: 5.5 Mb/s. Frame generation interval: 40 ms.
- **Web**: Page interarrival time exponentially distributed with mean 60 s. Page size 10 kbytes plus 20 to 80 objects of a size uniformly distributed between 5 and 10 kbytes.
- **FTP**: Download of a 20 Mbyte file.

Both for the web and FTP download applications we consider TCP New Reno and a round-trip time of 20 ms between the AP and the server storing the data. Web browsing uses HTTP 1.1.

Regarding the medium access and power saving mechanism used by each type of station, we consider the following five combinations:

- **HCCA+S-APSD** — Voice and video: HCCA+S-APSD. Service interval 40 ms and 80 ms, respectively. Web and FTP: EDCA+U-APSD. Trigger transmission based on beacon TIM indication.
- **EDCA+S-APSD** — Voice and video: EDCA+S-APSD. Service interval 40 ms and 80 ms, respectively. Web and FTP: EDCA+U-APSD. Trigger transmission based on beacon TIM indication.
- **EDCA+U-APSD** — Voice and video: EDCA+U-APSD. Periodic trigger interval 40 ms and 80 ms, respectively. Web and FTP: EDCA+U-APSD. Trigger transmission based on beacon TIM indication.
- **EDCA+PSM** — All applications: EDCA+PSM. Beacon period: 100 ms.
- **DCF+PSM** — All applications: DCF+PSM. Beacon period: 100 ms.

In order to decide when to schedule the resource allocations for HCCA and S-APSD, we use the algorithm described in [13], which results in a distribution in time as uniform as possible of the allocation of resources for the different flows. In the case of U-APSD, the periodic trigger generation is performed according to [11], which generates triggers at a configurable constant time interval if no data trigger has been sent within such an interval.

The EDCA QoS parameters used are chosen according to the 802.11e standard considering for aCWmin and aCWmax values of 127 and 1023, respectively. The TXOP length has been configured according to the 802.11g standard recommendation. In the DCF experiment, aCWmin and aCWmax are set as indicated in the 802.11g standard. Table 1 summarizes the

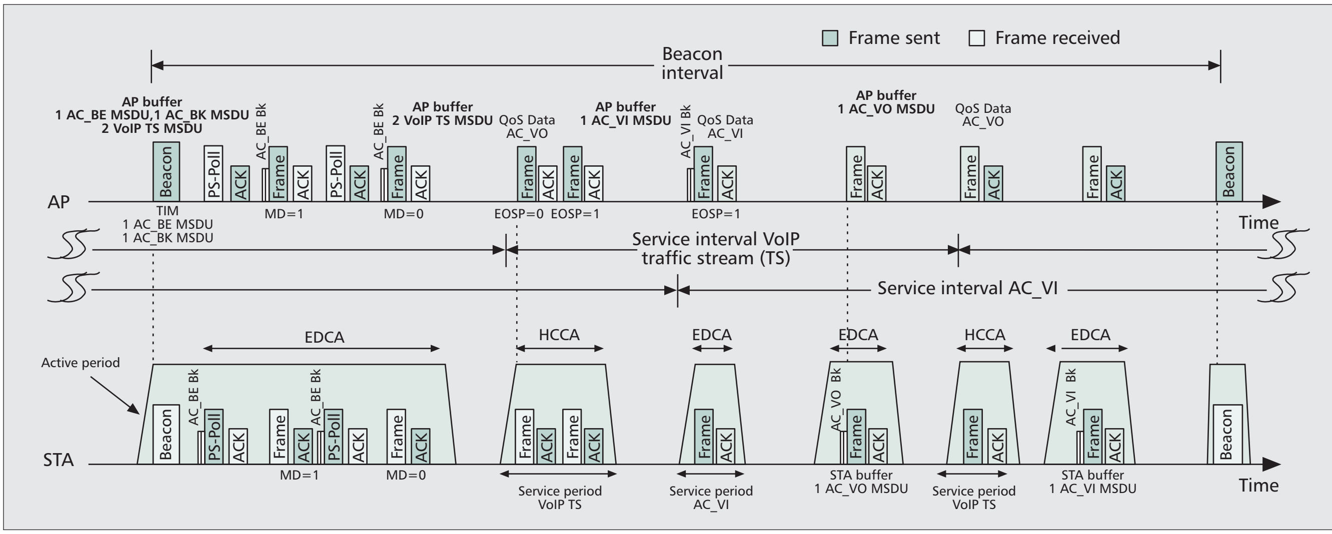

Figure 2. *S-APSD example of operation. S-APSD configuration: VoIP traffic stream in the downlink configured to use S-APSD with HCCA access mode, AC VI traffic in the downlink uses S-APSD with EDCA access mode, AC_BE and AC_BK configured to use legacy 802.11 power save mode. 1) First active period of the station: usage of the 802.11 power save mode mechanism to retrieve frames of AC_BE and AC_BK categories as indicated in the TIM. 2) Second active period: station wakes up at scheduled time for the VoIP traffic stream service period and receives the frames from the AP, which assigned itself a TXOP using the HCCA mechanism. 3) Third active period: station wakes up at scheduled time for AC_VI service period and receives the frame from the AP which obtained a TXOP using the EDCA mechanism. 4) Fourth active period: usage of the EDCA mechanism by the station to transmit an AC_VO frame in the uplink. 5) Fifth active period: station wakes up again at scheduled time for the VoIP traffic stream service period and receives one frame from the AP. 6) Sixth active period: usage of the EDCA mechanism by the station to transmit an AC_VI frame in the uplink.*

EDCA and DCF parameters used.

Power consumption values are obtained based on the data of a popular WLAN chipset [16] (Table 2). The average power consumption of STAs for each case is calculated by computing the percentage of time spent in each of its four possible states during the simulation duration, and then applying the power consumption at each state.

The length of the simulations performed is 300 s with a warmup phase of 30 s. Figures showing average values include the corresponding confidence intervals at 95 percent; graphs showing delay plot the 95 percent percentile of the delay obtained from all simulation samples. In the case of performance metrics grouped by application type, the values shown correspond to the expected performance experienced by a single station. These values are computed considering the results of all stations running the same kind of application.

EDCA	AIFS	CWmin	CWmax	TXOP length
AC_VO	2	31	63	3.264 ms
AC_VI	2	63	127	6.016 ms
AC_BE	3	127	1023	0
AC_BK	7	127	1023	0
DCF	2	15	1023	0

Table 1. *EDCA and DCF configuration.*

WLAN chipset [15]	Sleep	Idle	Rx	Tx
Power (mW)	20	390	1500	2000

Table 2. *Power consumption levels of a popular WLAN chipset.*

Real-Time Applications

In Figs. 3a and 3b we summarize the performance experienced by the Voice and Video applications in the considered scenario depending on the QoS and power saving configuration. For the delay and throughput metrics we focus in the downlink direction (AP → STA) since when power saving is used, this is the performance bottleneck of the system.

From the performance results, it can be observed that the configurations including the PSM functionality yield the worst performance in the three metrics considered. There are two main reasons for this. First, PSM synchronizes the requests by all stations of frames buffered at the AP to the period of time after the beacon frame reception, thus increasing the collision probability. Second, as explained earlier, PSM requires more signaling than the APSD mechanisms to retrieve the frames buffered at the AP, which increases again the collision probability in the channel.

Regarding throughput, EDCA+S-APSD is the best performing configuration for voice and HCCA+S-APSD for video. The reason for this result is that, unlike EDCA, where a higher priority is given to voice flows in front of video ones, our HCCA implementation [13] does not differentiate between real-time flows' duration, which in this case benefits the Video flows. Both configurations including S-APSD functionality outperform the alternative ones thanks to the fact that the AP can directly start the downlink transmission process when a scheduled service period

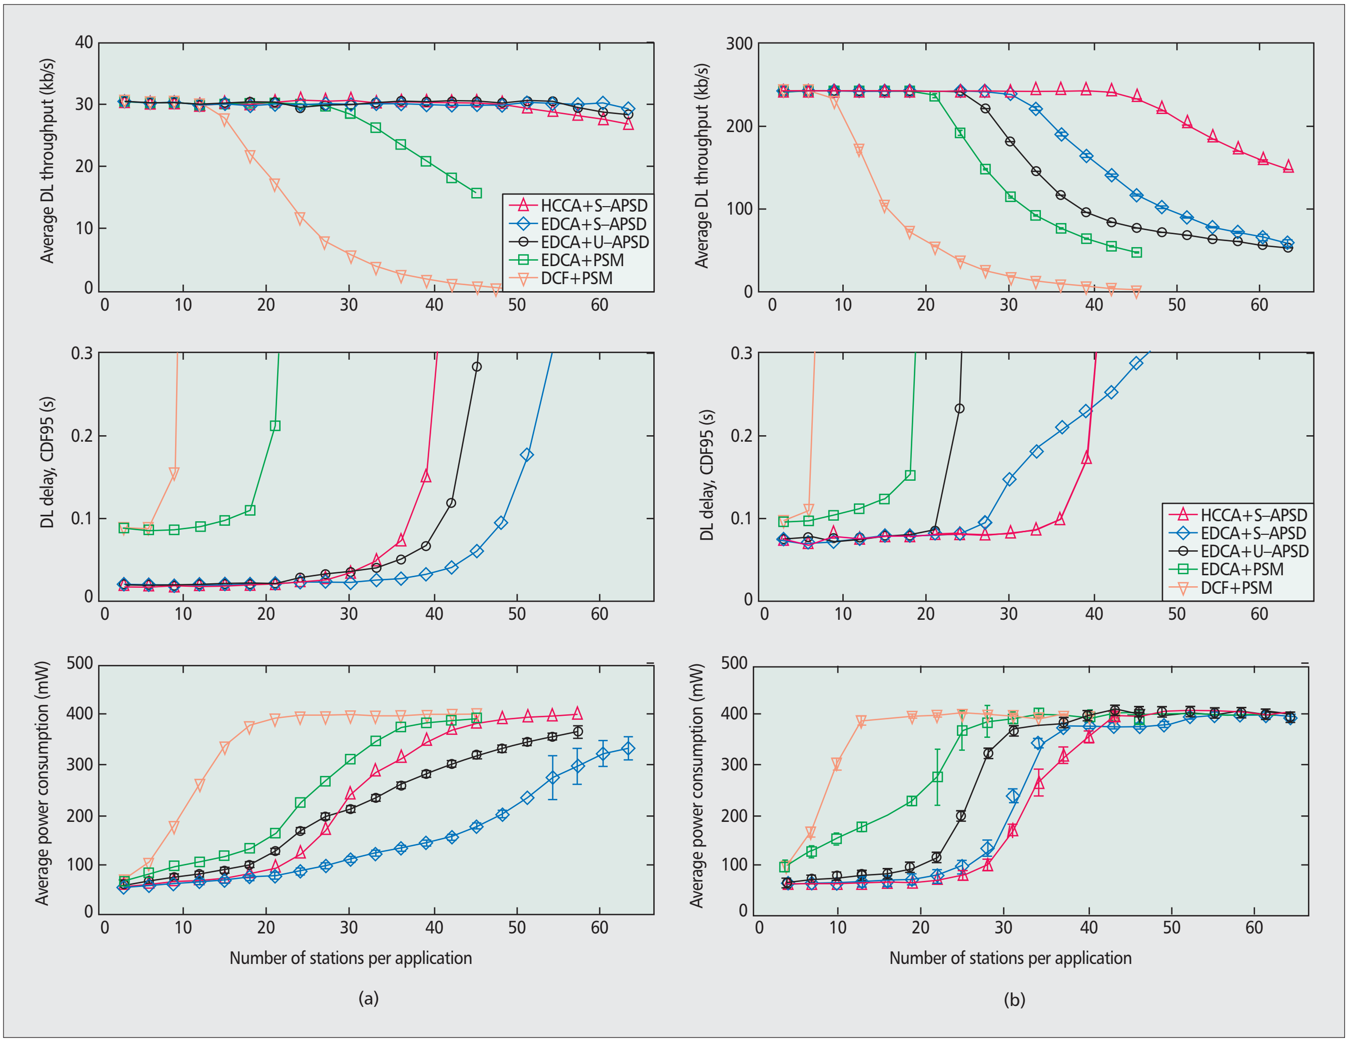

Figure 3a-b. *Performance experienced by single stations grouped by application type depending on the QoS and power saving configuration: a) voice over IP; b) video streaming.*

starts, without requiring any uplink signaling.

With respect to delay, configurations including S-APSD and U-APSD functionality successfully achieve their design objective of bounding the delay of voice and video to their preconfigured service interval. Configurations including PSM functionality, though, experience a delay bound close to 100 ms due to the dependency of the uplink signaling generation with the beacon interval.

Finally, the power saving performance of all approaches closely resembles the delay results since power consumption in a Wi-Fi network is directly related to the level of congestion in the network. Therefore, when the time to transmit a frame increases in an exponential way, the average power spent by the stations tends to be around 400 mW, slightly above the power spent by the wireless card in idle mode (390 mW), indicating that stations spend most of their time awake and waiting to transmit or receive data from the AP.

Non-Real-Time Applications

Figures 3c and 3d depict the performance obtained by the web and FTP applications for the different QoS and power saving schemes under study. Note that in all studied configurations, web and FTP stations are configured to use either U-APSD or PSM having the following common behavior: stations wait for the beacon to be notified about whether there are buffered frames at the AP for them and trigger the AP accordingly.

The first remarkable result for non-real-time applications is the very different maximum connection throughput achieved in the three stations per application case. While FTP stations reach a TCP connection throughput around 25 Mb/s *during* the duration of the FTP download, Web stations achieve a maximum TCP connection throughput below 600 kb/s. This phenomenon has been previously explained in the literature [17] and is due to the small size of the individual objects conforming a web page, resulting in web stations going back to sleep before the next window of TCP packets arrives at the AP. As a result, the effective round-trip time perceived by the TCP sender increases reducing in turn the achievable TCP throughput; see the web delay results in Fig. 3c. In the FTP download case, after accumulating a certain number of packets in the queue, the next

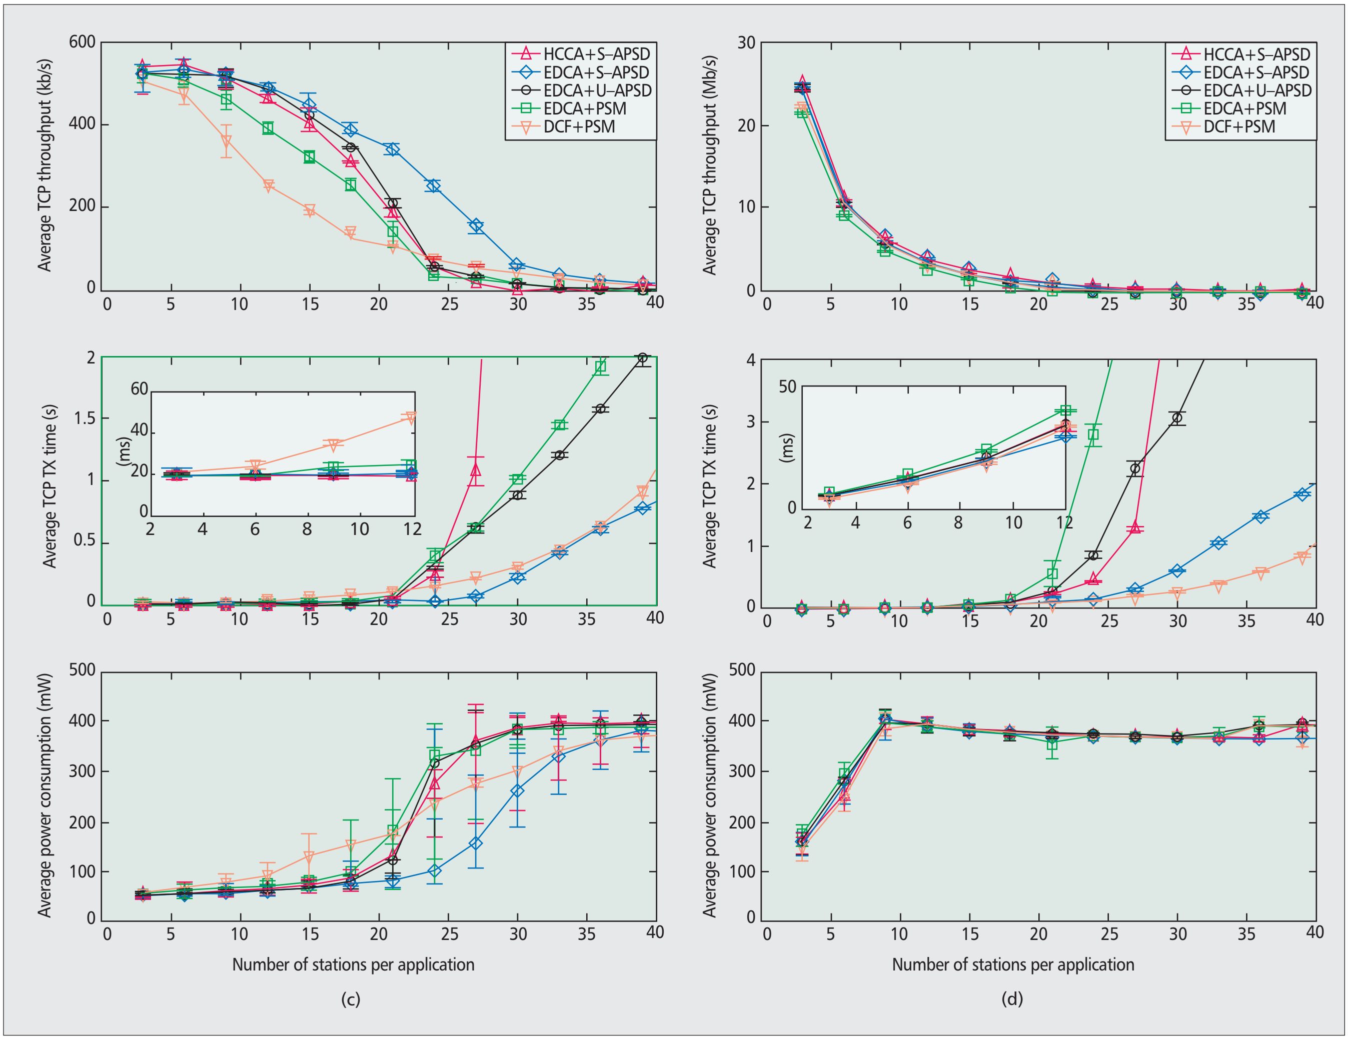

Figure 3c-d. *Performance experienced by single stations grouped by application type depending on the QoS and power saving configuration: c) web browsing; d) FTP download.*

window of TCP packets arrives before stations go back to sleep; thus, FTP stations remain awake and can make use of the whole available WLAN bandwidth; see the FTP delay results in Fig. 3d. This effect could be solved with the optional pipelining mechanism of HTTP 1.1, which is currently either not implemented or not enabled by default in most popular web browsers.

Regarding the performance differences between the considered QoS and power saving configurations, EDCA+S-APSD performs best when the congestion in the network is moderate, up to 24 stations per application, due to its reduced signaling load (as we see in the next section). However, under significant congestion conditions, the DCF+PSM configuration obtains better performance because DCF uses more aggressive medium access settings than EDCA for non-real-time applications.

Network Performance

Finally, in order to assess the overall network performance, in Fig. 4 we depict the total aggregated throughput and signaling load of each QoS and power saving configuration. Signaling load is computed as a percentage of the time the wireless channel is occupied by QoS or power saving related signaling (i.e., PS-polls and QoS nulls for PSM, U-APSD, and S-APSD, and CF-polls and QoS-nulls for HCCA+S-APSD).

The configuration achieving the highest total aggregated throughput when the network is congested is HCCA+S-APSD, from 36 stations per application on. The reason is that HCCA requires less channel capacity for medium access arbitration than configurations using the EDCA and DCF functionality, and thus is more robust in highly congested conditions.[2] However, the maximum throughput achieved is only 14 Mb/s, while in our scenario stations transmitting at 54 Mb/s, 24 Mb/s, and 12 Mb/s could potentially reach an average PHY rate of 30 Mb/s. Such low efficiency is due to the fact that under high congestion conditions, mostly voice and video frames are transmitted, resulting in a high MAC overhead. A higher efficiency could be achieved by making use of 802.11e and 802.11n overhead reduction mechanisms (e.g., Block Acknowledgments or piggybacking of signaling frames in data ones). It is worth noting that configurations including the PSM functionality present the worst total throughput results. As aforementioned, this is due to the higher

[2] *Note that the HCCA mechanism, in order to operate properly, requires that neighboring APs operate in non-overlapping channels and therefore can only be used when this condition holds. Currently, this issue is being considered by the IEEE 802.11 Working Group within the TGaa Task Group.*

collision probability and signaling load of PSM.

The signaling overhead of the different configurations under study is depicted in Fig. 4. Configurations including APSD functionality present a lower signaling load than PSM, especially when congestion increases, thanks to their maximum service period length feature. The larger the congestion, the more packets delivered within the same service period, up to the configured maximum service period length value, resulting in a significant reduction of the signaling load. Specifically, EDCA+S-APSD is the configuration introducing less signaling in this particular scenario, mainly due to the applications chosen, which generate most of their traffic in the downlink direction (the voice application being an exception). Configurations including S-APSD functionality do not require any signaling to initiate a downlink service period, once configured; thus, a significant reduction in signaling load is achieved. In addition, HCCA+SAPSD introduces more signaling, up to the 30 stations case, than all other configurations but DCF+PSM due to the VoIP application contribution. While HCCA+ SAPSD polls VoIP stations in uplink even if they are in a silence period, all other configurations do not generate additional signaling in such a case.

Summary and Conclusions

WLAN technology has become the de facto standard for wireless broadband Internet access in homes, offices, and public hotspot locations. Overall, market forecasts predict WLAN units to grow roughly 20 percent per year, reaching a billion units by 2012 [2]. The wide range of devices in the WLAN market target different customer needs, but two main requirements are commonly shared by most new devices: QoS support for prioritizing real-time services over non-real-time and power saving functionality to achieve an operating time meeting users' expectations.

Our contributions are as follows. First, we provide an overview of the two power saving modes defined by IEEE 802.11e, U-APSD and S-APSD. Second, we quantitatively evaluate in a generic scenario, the performance differences to be expected between the different possible combinations of the IEEE 802.11e QoS and power saving mechanisms, HCCA+S-APSD, EDCA+S-APSD, EDCA+U-APSD, and EDCA+PSM, with respect to MAC throughput and delay, power saving efficiency, and signaling load. As a reference, we also include in the study the combination of the baseline 802.11 access mechanism (DCF) and its corresponding power saving mode (PSM), which is currently the configuration most commonly used.

Through this study, an insight on the performance of the various possible combinations of the IEEE 802.11e QoS and power saving mechanisms and its causes has been provided. Additionally, we explain the reasoning behind the differences in the performance of the various possibilities to provide guidance to network operators when considering which option would better meet their users' needs.

The main conclusions that can be drawn from our study are:

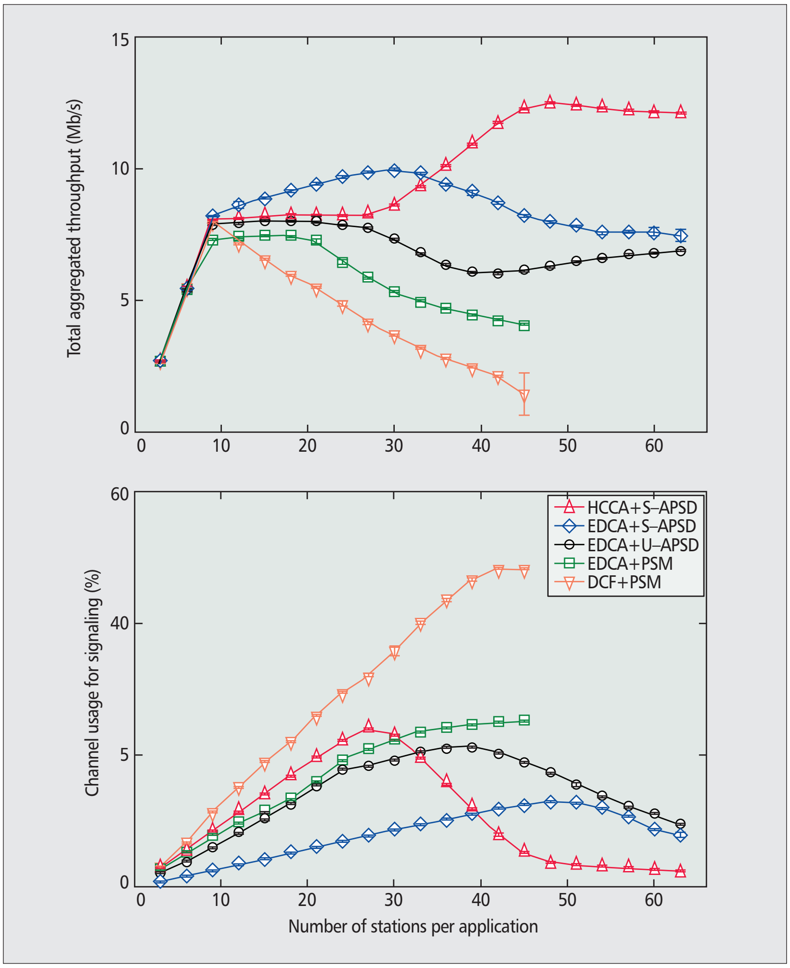

Figure 4. *Wireless LAN total throughput and signaling load.*

- Configurations including S-APSD functionality outperform their alternatives in almost all cases thanks to its lower required signaling load for downlink transmissions.
- Configurations including S-APSD and U-APSD functionality successfully achieve their design objective of bounding the delay to a configurable value if the network is not congested.
- The EDCA+S-APSD configuration results in non-real-time applications achieving a performance similar to or even better than the one they would expect if no QoS mechanisms were present because of the lower signaling load.
- EDCA+S-APSD is better suited for bidirectional applications with periods of no activity (e.g., VoIP), than HCCA+S-APSD since no signaling needs to be generated in the downlink during such periods.

Acknowledgments

The authors would like to thank all colleagues involved in NEC's 3G/WLAN Mobile Terminals development for their contributions to this research work.

References

[1] IEEE 802.11 WG, "IEEE Standard for Local and Metropolitan Networs Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Applications," IEEE 802.11-2007.

[2] Wi-Fi Alliance; http://www.wi-fi.com
[3] IEEE 802.11 WG, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications. Amendment 8: Medium Access Control (MAC) Quality of Service Enhancements," IEEE 802.11e, Nov. 2005.
[4] S. Blake *et al*., "An Architecture for Differentiated Services," RFC 2475, Dec. 1998.
[5] R. Braden, D. Clark, and S. Schenker, "Integrated Services in the Internet Architecture: An Overview," RFC 1663, June 1994.
[6] S. Mangold *et al*., "Analysis of IEEE 802.11e for QoS Support in Wireless LANs," *IEEE Wireless Commun*., Dec. 2003.
[7] J. Lee *et al*., "A Practical QoS Solution to Voice over IP in IEEE 802.11 Wireless LANs," *IEEE Commun. Mag*., Apr. 2009.
[8] J. Adams and G.-M. Muntean, "Power Save Adaptation Algorithm for Multimedia Streaming to Mobile Devices," *Proc. IEEE Int'l. Conf. Portable Info. Devices*, May 2007.
[9] X. Pérez-Costa, D. Camps-Mur, and T. Sashihara, "Analysis of the Integration of IEEE 802.11e Capabilities in Battery Limited Mobile Devices," *IEEE Wireless Commun*., Dec. 2005.
[10] X. Pérez-Costa and D. Camps-Mur, "APSM: Bounding the Downlink Delay for 802.11 Power Save Mode," *Proc. IEEE ICC*, May 2005.
[11] X. Pérez-Costa, D. Camps-Mur, and A. Vidal, "On the Distributed Power Saving Mechanisms of Wireless LANs: 802.11e U-APSD vs. 802.11 Power Save Mode," *Comp. Net*., vol. 51, no. 9, June 2007.
[12] D. Camps-Mur, X. Pérez-Costa, and S. Sallent "An Adaptive Solution for Wireless LAN Distributed Power Saving Modes," *Comp. Net*., vol. 53, no. 18, Dec. 2009.
[13] X. Pérez-Costa *et al*., "OAS-APSD: Design and Evaluation of an Overlapping Aware Scheduled Automatic Power Save Delivery Algorithm," *Proc. Euro. Wireless Conf.*, Apr. 2007.
[14] OPNET Simulator; http://www.opnet.com
[15] F. H. P. Fitzek and M. Reisslein, "MPEG-4 and H.263 Video Traces for Network Performance Evaluation," *IEEE Network*, Dec. 2001.
[16] Broadcom 4311AG 802.11a/b/g; http://www.broadcom.com
[17] R. Krashinsky and H. Balakrishnan, "Minimizing Energy for Wireless Web Access with Bounded Slowdown," *Proc. ACM MobiCom*, Apr. 2006.

Biographies

XAVIER PÉREZ COSTA (perez@neclab.eu) is a chief researcher at NEC Laboratories Europe in Heidelberg, Germany, where he is currently the responsible of managing several projects related to wireless networks. In the wireless LAN area, he leads a project contributing to 3G/WiFi mobile phones evolution and received NEC's R&D Award for his work on N900iL, NEC's first 3G/WiFi phone. In the WiMAX area he manages a team researching on NEC's WiMAX products future enhancements. In the future wireless Internet area, he is responsible for a team contributing to the EU-FP7 project Flexible Architecture for Virtualizable Wireless Future Internet Access (FLAVIA). He has participated in IEEE standardization working groups as well as their corresponding certification bodies, and has been included in the major contributor lists of IEEE 802.11e and 802.11v. He has served in the program committee of several conferences, including IEEE ICC, WCNC, and INFOCOM, and holds over 10 patents. He received both his M.Sc. and Ph.D. degrees in Telecommunications from the Polytechnic University of Catalonia (UPC) and was the recipient of the national award for the best Ph.D. thesis on multimedia convergence in telecommunications from the Official Association of Telecommunication Engineers (COIT).

DANIEL CAMPS MUR (camps@neclab.eu) is a researcher in the Wireless Access Group at NEC Network Laboratories in Heidelberg, Germany. He is currently pursuing a Ph.D. at the Telematics Department of UPC, where he obtained an M.Sc. degree in 2004. His current research interests include MAC QoS and power saving protocols for WLAN and WiMAX networks. His master thesis work received the mobile internet and 3G mobile solutions award from COIT. In the field of network simulations he also received OPNET's Significant Technical Challenge Solved Award.