SU-APSD: Static IEEE 802.11e Unscheduled Automatic Power Save Delivery

Xavier Pérez-Costa, Albert Vidal and Daniel Camps-Mur

NEC Network Laboratories, Kurfürsten-Anlage 36, Heidelberg, Germany
{perez, vidal, camps}@netlab.nec.de

***Abstract*: The integration of the wireless LAN technology in mobile devices as cellular phones or PDAs has become a user need due to its popularity in providing high speed wireless Internet access at a low cost. Such devices though should meet users' expectations with regard to QoS, e.g., guarantee a reasonable voice quality when VoIP is used, and power saving efficiency, e.g., standby and calling times should be similar to the ones of cellular phones. IEEE 802.11e defines QoS and power saving enhancements that should allow the wireless LAN technology address users' wishes in such specific devices. However, there are several questions that need to be answered in order to assess whether the additional complexity introduced by these mechanisms would be justified by a relevant performance improvement. With our work we contribute to answer these questions by i) identifying the likely 802.11e functionality to be implemented in mobile devices in the short-term (EDCA+U-APSD) ii) proposing a specific implementation of the U-APSD mechanism (SU-APSD) and iii) evaluating the performance improvement obtained as compared to 802.11 standard power save mode.**

1 Introduction

The 802.11 Task Group E has recently completed the design of QoS MAC enhancements which are specified in the 802.11e extension of the standard [1]. 802.11e defines the Hybrid Coordination Function (HCF) to support QoS. Two channel access methods are defined: a contention-based channel access method called the Enhanced Distributed Channel Access (EDCA) and a contention-free channel access referred to as HCF Controlled Channel Access (HCCA). Within a superframe two phases of operation are possible in 802.11e, contention period (CP) and contention-free period (CFP). HCCA can be used in both CP and CFP while EDCA can be used only during CP. A thorough overview of the 802.11e QoS enhancements can be found in [2].

Regarding the battery usage efficiency, the intrinsic nature of 802.11, which is based in a shared channel access (CSMA/CA), forces wireless stations to continuously listen to the channel to determine its current status. As a result, a handheld device connected through an 802.11 wireless LAN, will drain its battery after a few hours as opposed to current mobile devices that have a standby battery lifetime up to several days, e.g., cellular phones. Ideally, mobile devices including the wireless LAN technology should achieve a battery consumption similar to current handheld devices in order to meet users' expectations.

The IEEE 802.11 standard provides a power save mode that reduces the time required for a station to listen to the channel. Once every beacon interval, usually 100ms, the access point (AP) sends a beacon indicating whether or not a certain station has any data buffered at the AP. Wireless stations[1] wake up to listen to beacons at a fixed frequency 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 MAC downlink delay (AP to station) and the listen interval is introduced. Consequently, some listen interval values can result in downlink delays unacceptable for certain QoS-sensitive applications, e.g., VoIP. Further details about the power save mode operation can be found in [3].

In order to address some of the power saving issues that arise when a specific QoS is desired, the 802.11e standard includes an optional extension of the 802.11 power save mode defined as Automatic Power Save Delivery (APSD). If the AP supports APSD, 802.11e-capable stations can select the method for the delivery of the frames buffered at the AP between standard Power Save Mode and APSD. A main difference between the 802.11 standard power save mode and APSD is that with APSD a station is awake during a Service Period[2] (SP) instead of being awake from the transition to the awake state for receiving a Beacon until the return to the doze state after acknowledging receipt of the last frame buffered at the AP. Two types of SPs are possible under APSD: unscheduled and scheduled. Unscheduled SPs (U-APSD) are defined only for stations accessing the channel using EDCA while Scheduled SPs (S-APSD) are defined for both access mechanisms, EDCA and HCCA.

New mobile devices incorporating 802.11e functionality are more likely to include first the 'distributed' mechanisms of 802.11e, i.e., EDCA and U-APSD, than the centralized ones, i.e., HCCA and S-APSD. This can be seen for instance in the fact that the Wi-Fi^{TM} Alliance [4] has started first the certification of the Wi-Fi^{TM} multimedia extensions (WMM^{TM}) and the WMM Power Save^{TM}, which include EDCA and EDCA plus U-APSD functionalities respectively, while certifications for HCCA and S-APSD are being deferred. Based on that, we focus our study in the analysis of the performance enhancements provided by EDCA in combination with U-APSD since it is the most probable configuration to be implemented in the short-term in 802.11e-capable mobile devices.

To the best of the authors' knowledge there is no published related work regarding the performance im-

[1] In this paper we refer as AP and station to what in the 802.11e standard is denoted as QAP and non-AP QSTA respectively.

[2] The definition of a SP will be given in the next section.

provement obtained by the usage of the 802.11e QoS and power saving mechanisms as compared to standard power save mode. In the area of providing QoS in a wireless LAN a lot of research has been done during the last several years, see for example [2], [5]. With respect to the 802.11 power save mode, the infrastructure mode has been studied for instance in [6],[7] where the main focus was to improve the performance for web-like traffic. Regarding U-APSD, in [8] an early specification of U-APSD (802.11e draft 6.0) was studied where the U-APSD power saving performance was compared to the 802.11 power save mode one in a scenario where all users generate and receive constant bit rate voice traffic. In our previous work [9] we studied the effect of using the standard power save mode in conjunction with the 802.11e QoS mechanisms, and in [10] we proposed an adaptive algorithm to control the delay introduced by the 802.11 power save mode for application requiring QoS guarantees. The paper at hand, extends our previous results and the ones provided in [8] by considering the latest U-APSD specification (802.11e standard), including the 802.11 EDCA QoS mechanisms, using realistic traffic with non-deterministic behavior and evaluating the performance not only in power saving terms but also with regard to the QoS perceived by the users of the different classes, signaling load and available channel capacity.

The rest of the paper is structured as follows. In Section 2 an overview of the U-APSD functionality is given followed by a description of our proposed implementation in Section 3. An evaluation of the performance results obtained by using U-APSD as compared to 802.11 power save mode is provided in Section 4. Finally, Section 5 summarizes the results and concludes the paper.

2 Unscheduled Automatic Power Save Delivery (U-APSD)

Unscheduled Automatic Power Save Delivery (U-APSD) is the distributed APSD method defined in 802.11e to improve the QoS provided to stations accessing the channel using the EDCA mechanism as compared to legacy power save mode. The main 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 the power saving stations are awake and then take advantage of it for delivering any data frames that were buffered at the AP while the stations where in doze mode. Because of its specific functionality, this method is specially suited for bi-directional traffic streams even though it provides alternative methods for its usage in other cases. APSD has been designed such that it provides backwards compatibility to legacy terminals implementing standard power save mode only. Therefore, all new functionality was built on top of the already available standard power save mode one re-using as much as possible without altering the original power save mode specification. In the following we describe in detail the U-APSD functionality assuming a basic knowledge of the 802.11 standard power save mode and of the EDCA mechanism of 802.11e. Please see [2] and [9] for an overview of 802.11e and standard 802.11 power save mode respectively.

Functionality Description

As previously mentioned, the main difference between the power saving method defined in the 802.11 standard and APSD is that with APSD a station is awake during a Service Period (SP) instead of being awake from the transition to the awake state for receiving a Beacon until the return to the doze state after acknowledging receipt of the last frame buffered at the AP through PS-Polls.

An *unscheduled SP* begins when the AP receives a *trigger frame* from a station and ends when the station receives a QoS Data or QoS Null frame indicating the end of the service period (EOSP). QoS Null frames are the substitutes in U-APSD of PS-Polls in standard power save mode to allow a station to request the delivery of the frames buffered at the AP even if a station has no data frame to transmit in the uplink. This enables the usage of U-APSD by applications which do not generate uplink traffic often enough to meet the QoS application requirements.

Each AC at the stations can be configured separately to be delivery/trigger-enabled respectively. When a station has an AC configured as *delivery-enabled*, the AP is allowed to use EDCA to deliver traffic from the AC to a STA during an unscheduled SP triggered by a station. When a station AC is *trigger-enabled*, frames of subtype QoS Data and QoS Null from the station, that map to the AC, trigger an unscheduled SP if one is not in progress.

During a 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 following the rules of an acquired transmission opportunity. The maximum service period length is a field contained in the QoS Info field filled by the station at association. In each directed data or management frame associated with delivery-enabled ACs sent to a station, the More Data (MD) bit indicates that more frames are buffered for the delivery-enabled ACs. For all frames except for 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 the continuation of the SP.

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 ACs non delivery-enabled. Note that, in this case, it means that the Beacon will not indicate whether frames of ACs delivery-enabled are buffered. Only in the case that all ACs are delivery-enabled the TIM indicates the buffer status of delivery-enabled ACs.

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 element info field of the add traffic stream frames (ADDTS). The option of configuring U-APSD at association has the advantage of not having to generate

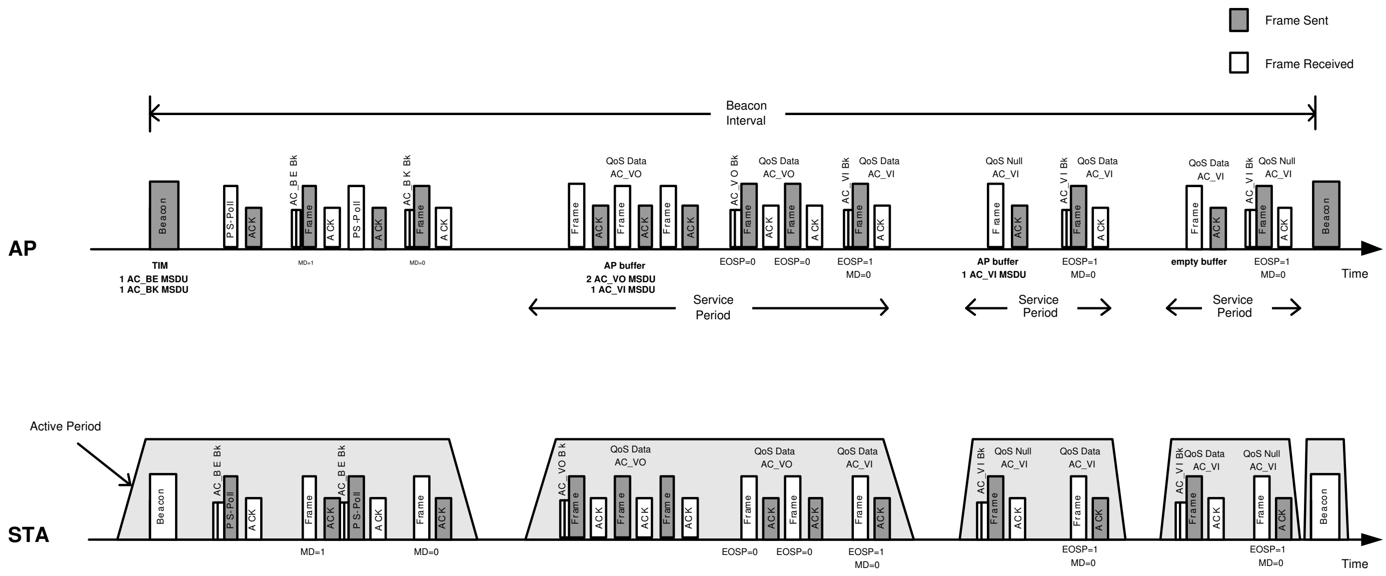

Figure 1: U-APSD example of operation. U-APSD configuration: AC_VO and AC_VI are both trigger- and delivery-enabled and AC_BE and AC_BK are neither delivery- nor trigger-enabled (i.e., use legacy 802.11 power save mode). 1) First active period of the station: usage of the legacy power save mode mechanism to retrieve frames of AC_BE and AC_BK categories. 2) Second active period: the 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: the station wakes up because the U-APSD algorithm decides that a SP has to be started. In this case the STA sends a QoS Null frame to trigger the unscheduled SP. 4) Finally, in the fourth active period, the station wakes up because it has an AC_VI data frame to be sent and starts an unscheduled SP.

ADDTS frames each time a new communication starts for each of the ACs for which we would like to use U-APSD, but because the TIM in the Beacon might not indicate whether frames are buffered at the AP, it requires to periodically send QoS Null frames if no other triggers are sent in order to learn the actual buffer status of the AP. On the other hand, if U-APSD is configured through ADDTSs each time a new application starts which requires its usage, then the stations can rely on the TIM information of the Beacons to be informed about new traffic at the AP, enter into U-APSD mode through ADDTS when a communication starts and revert back to legacy power save mode when it ends.

In Figure 1 an example illustrating the U-APSD procedure is provided.

Advantages of U-APSD with respect to 802.11 power save mode

Three main advantages are introduced with U-APSD with respect to 802.11 power save mode. These advantages will be evaluated in Section 4.

- The first advantage is the improvement of the QoS of the ACs with U-APSD enabled since, because of the possibility of generating triggers at any point in time, the delay introduced by U-APSD can be limited based on the application QoS needs instead of depending on the listen interval configuration of the 802.11 power save mode. Additionally, triggers get access to the channel according to the priority of the frames intended to be retrieved as compared to the PS-Polls which by default use AC_BE.
- The second advantage is the reduction of the overhead required to retrieve frames from the AP thanks to the usage of data frames as triggers. This is specially relevant for symmetric applications like VoIP, because then almost no QoS Null frames are required and therefore the significant PS-Poll overhead required by the standard 802.11 power save mode is significantly reduced.
- The third advantage is again an important reduction of the overhead required to retrieve frames from the AP depending on the maximum service period length (Max_SP_Length) per trigger. While for 802.11 power save mode a signaling frame (PS-Poll) is required for obtaining each single frame buffered at the AP this is not necessary for U-APSD which delivers up to Max_SP_Length frames per trigger.

3 Proposed U-APSD Implementation

The U-APSD specification provided in the IEEE 802.11e standard introduces different mechanisms to control the QoS and power saving provided to a station for each different AC. The specific implementation of these mechanisms to actually deliver the desired QoS is though, as usual, left open to allow differentiation between vendors. In the following we describe our proposed implementation of an U-APSD algorithm.

For the design of the U-APSD algorithm we assumed that the EDCA QoS parameters are properly set by a QoS algorithm at the AP to provide the required bandwidth and delay guarantees for the applications. Therefore, the objectives of our algorithm design are:

- Limit the MAC downlink delay introduced by U-APSD to a value acceptable for the applications needs

- Minimize the required signaling overhead introduced in the channel

In order to limit the MAC downlink delay (AP → station) introduced by U-APSD, the first problem that needs to be solved is to guarantee that triggers will be generated to meet the delay requirements even if the application running at the station does not generate enough data frames (triggers). Our proposed solution to this problem is to schedule at fixed time intervals the generation of QoS Null frames to bound the maximum delay introduced by U-APSD. The time intervals could be set based on some information provided to the MAC layer depending on the application or simply fixed based on an assumed maximum delay requirement per AC. Additionally, to avoid the introduction of unnecessary overhead in the channel due to the generation of QoS Null frames, we propose to re-schedule the QoS Null frame generation of U-APSD trigger-enabled ACs each time that a new data frame of such an AC is prepared to be sent. The re-scheduling of QoS Null frames due to a QoS Data frame of trigger-enabled AC is depicted in Figure 2. In this way, we guarantee that a QoS Null frame will be generated to trigger a SP only after reaching the maximum delay we would like to allow for U-APSD trigger-enabled ACs since the last transmission of a data frame of such an AC. Note that if more than one AC is U-APSD trigger-enabled for a station at the same time, we only need to schedule the transmission of QoS Null frames for the one with the minimum delay requirements because the requirements of other ones would be fulfilled by the most stringent one.

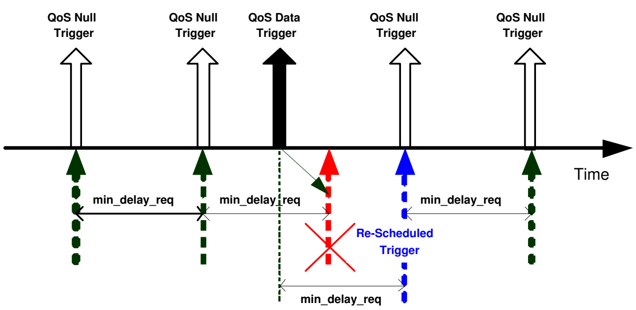

Figure 2: U-APSD Re-scheduling example of operation

Once a trigger is received at the AP, we start delivering first the frames of the highest priority delivery-enabled AC and continue until all buffered frames of delivery-enabled ACs for a certain station have been delivered or the maximum service period length has been reached. Since, due to the maximum service period length, it could happen that a SP concludes before all the frames of a certain delivery-enabled AC had been delivered to a station, we use in our algorithm the More Data information provided in the data frames to decide whether a new SP has to be started after receiving the EOSP indication.

The maximum service period length is another parameter that needs to be configured for U-APSD. In our proposed implementation we configure the maximum service period length (Max_SP_Length) such that it permits sending the maximum number of frames allowed by

---

Algorithm 1 U-APSD algorithm that soft bounds the MAC downlink delay to the minimum delay value configured for the trigger-enabled ACs of a station

---

Trigger interval selection
for All Trigger-Enabled ACs **do**
  if $delay\_req[AC_i] < min\_delay\_req$ **then**
    $min\_delay\_req = delay\_req[AC_i]$
  end if
end for

Start trigger generation
Generate the first trigger when the first AC of a station is configured as trigger-enabled
$trigger\_time = current\_time + min\_delay\_req$

Periodic trigger generation
if $trigger\_time = current\_time$ **then**
  $trigger\_time = trigger\_time + min\_delay\_req$
end if

Trigger re-scheduling
if A frame is sent in the uplink of a trigger-enabled AC **then**
  Cancel the pending Trigger and schedule a new one at $trigger\_time = current\_time + min\_delay\_req$
end if

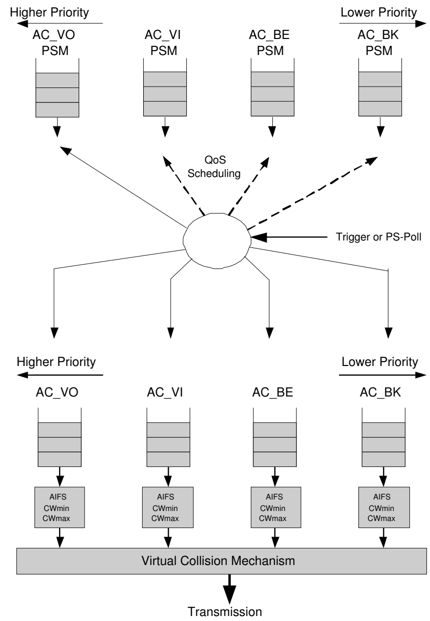

Figure 3: U-APSD queueing architecture

the 802.11e specification, i.e., all. We have chosen this option for two main reasons. First, because of power saving efficiency since the more frames sent in a row after getting access to the channel, the less contention for the channel required and therefore the SPs can complete faster. Second, because by retrieving the frames from the AP as soon as possible we maximize the chances that the traffic meets its delay requirements. Note that even if the Max_SP_Length has no limit, the maximum time that a station can keep control of the channel will continue to be bounded by the TXOP limit.

Finally, to avoid that high priority traffic gets dropped in congestion conditions at the AP due to the low retrieval rate of low priority one, we propose to have separate power save mode queues at the AP (physical or logical) for each AC to guarantee that no high priority traffic is dropped because of having the power save mode queue full of low priority traffic. Figure 3 depicts the resulting queueing architecture at the AP.

Based on the different considerations explained above we designed the static algorithm for U-APSD described in Algorithm 1.

4 Performance Evaluation & Discussion

In this section we evaluate the performance of our proposed algorithm for U-APSD as described in Section 3 with respect to standard 802.11 power save mode. The analysis has been performed via simulation. We extended the 802.11b libraries provided by OPNET [11] v10.0 to include the EDCA QoS mechanisms of 802.11e, our proposed U-APSD algorithm and 802.11 standard power save mode.

EDCA-TXOP durations are configured for all ACs to allow the transmission of one data frame after gaining access to the medium. The RTS/CTS mechanism has not been enabled to avoid its influence over the mechanisms being studied. The beacon interval used is 100ms and the listen interval configured for the power saving stations is 1. The maximum service period length (Max_SP_Length) configured allows to send all delivery-enabled frames for a certain station.

In the analysis, the performance of two different U-APSD configurations will be compared to the standard 802.11 power save mode (StdPSM) one. First, a configuration where only the ACs categories which are supposed to require QoS guarantees, i.e., AC_VO and AC_VI, are configured to use U-APSD while AC_BE and AC_BK use standard power save mode. This case will be referred as *UAPSD+StdPSM*. Second, we consider a configuration where *all* ACs are configured to use U-APSD. This case will be referred as *UAPSD*. The delay requirements set for the different ACs when being configured to use U-APSD are 40, 60, 100 and 100ms for AC_VO, AC_VI, AC_BE and AC_BK respectively. We have chosen the delay requirements for AC_VO and AC_VI 20ms above the frame generation rate of the traffic sources used (20 and 40 ms respectively) to model the fact that in general the trigger generation interval will not be matching perfectly the frame generation rate but will rather have a value that covers most of the common values. With respect to the delay requirements for AC_BE and AC_BK, we have chosen the same value of 100ms to compare with standard power save mode (beacon interval: 100ms) but with the enhanced functionality of U-APSD.

Since our focus is on the differences in the performance between the power saving mechanisms when being used together with EDCA we assume a fixed configuration of the 802.11e EDCA QoS parameters based on the 802.11e document recommendation [1]. The parameters used are detailed in Table 1.

The configuration of the applications used is detailed below:

	AIFS	CWmin	CWmax
AC_VO	2	31	63
AC_VI	2	63	127
AC_BE	3	127	1023
AC_BK	7	127	1023

Table 1: EDCA configuration for the different ACs

- AC_VO: G.711 Voice codec with silence suppression. Data rate: 64kbps. Frame length: 20ms. Talk spurt exponential with mean 0.35s and silence spurt exponential with mean 0.65s.
- AC_VI: MPEG-4 real traces of the movie 'Star Trek: First Contact' obtained from [12]. Target rate: 64kbps. Frame generation rate: 40ms.
- AC_BE: Web traffic. Page interarrival time exponentially distributed with mean 10s. Page size 10KB plus 1 to 5 images of a size uniformly distributed between 10KB and 100KB.
- AC_BK: E-mail. Send interarrival time exponentially distributed with mean 120s. Receive interarrival time exponentially distributed with mean 60s. Size exponentially distributed with mean 100KB.

The length of the simulations performed is 300 seconds with a warm-up phase of 30 seconds. The number of seeds used to obtain each value in the graphs has been chosen such that performance results crossings between ACs are due to the differences between the power saving mechanisms.

In the evaluation we study the impact of increasing the number of stations over the uplink/downlink MAC throughput and delay, the power saving efficiency and the resulting power saving costs. The experiment starts with a scenario of four wireless LAN stations where each station is configured to send and receive traffic from their corresponding pair in the wired domain of a different AC, i.e., one station sends and receives AC_VO traffic, a second one sends and receives AC_VI traffic and so on. Due to space restrictions we do not show in the following sections the uplink results (STA → AP) for delay and throughput since they are the less relevant ones.

Note that the two power saving mechanisms subject of study modify only the behavior of a station to obtain data frames from the AP and therefore the performance bottleneck is the downlink direction.

QoS Differentiation

Figure 4 illustrates the effect of configuring U-APSD for the AC_VO and AC_VI categories instead of standard 802.11 power save mode over the data throughput at the MAC layer. As expected, when there is no congestion in the wireless channel, similar results are obtained with both approaches in throughput terms since the signaling load required by the different mechanisms has no major impact.

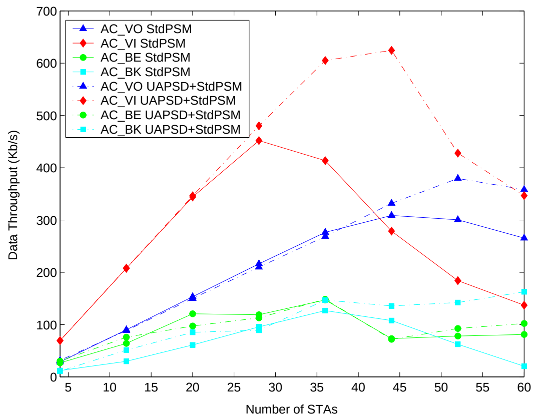

Figure 4: Impact of the number of stations on the MAC downlink throughput, StdPSM vs UAPSD+StdPSM

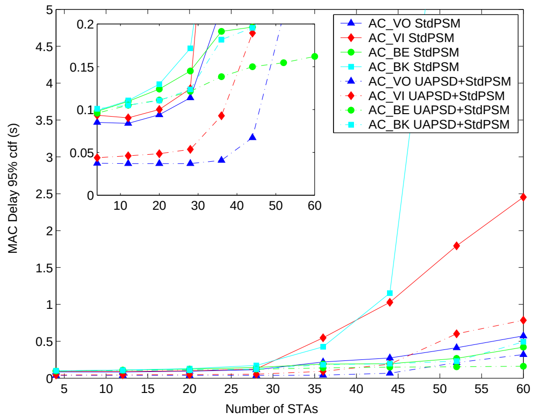

Figure 5: Impact of the number of stations on the MAC downlink delay, StdPSM vs UAPSD+StdPSM

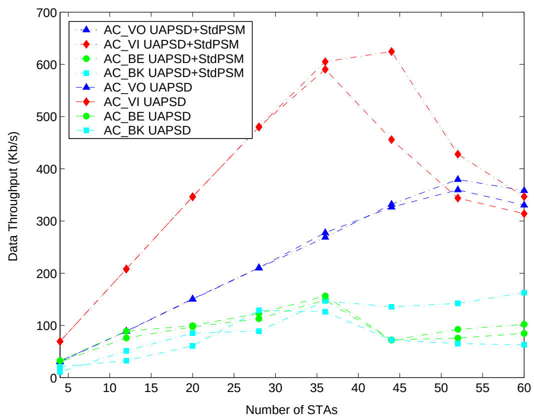

Figure 6: Impact of the number of stations on the MAC downlink throughput, UAPSD+StdPSM vs UAPSD

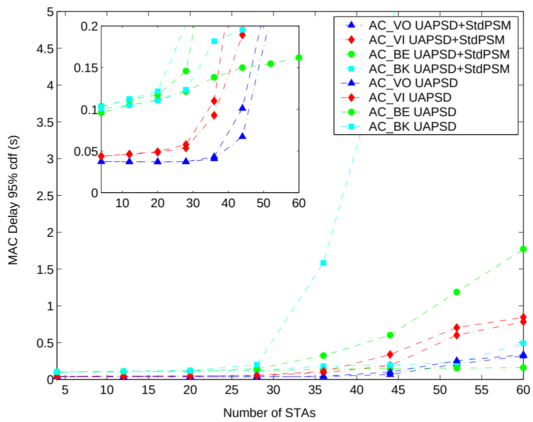

Figure 7: Impact of the number of stations on the MAC downlink delay, UAPSD+StdPSM vs UAPSD

However, when the channel starts to get congested, *U-APSD+StdPSM* (dot-dashed line) clearly outperforms *StdPSM* (solid line) specially for AC_VO and AC_VI due to the data frames acting as *triggers*, which reduce the signaling load, and the fact that the triggers are sent in a non synchronous way, reducing thus the probability of collision.

The corresponding results with respect to the MAC downlink delay are shown in Figure 5. The results perfectly match the throughput ones in congestion conditions, UAPSD+SdtPSM presents a significant performance improvement as compared to StdPSM. A clear differentiation though is observed in no congestion conditions that did not appear before. The reason for the significant lower delay for AC_VO and AC_VI in the UAPSD+StdPSM case is that triggers are generated every 40ms and 60ms respectively as compared to 802.11 power save mode where the Beacon is sent every 100ms. Additionally, the lower signaling load and probability of collision in the UAPSD+StdPSM case result also in an improvement for the AC_BE and AC_BK categories.

Figures 6 and 7 show the downlink metrics performance for the UAPSD+StdPSM (dot-dashed line) and the UAPSD (dashed line) cases. We can see here that the UAPSD+StdPSM configuration outperforms the UAPSD one where all the ACs use U-APSD. The reason for the better performance of the UAPSD+StdPSM configuration is that when AC_BE and AC_BK use U-APSD, triggers are sent periodically (every 100ms) to check whether there are frames buffered in the AP's power save mode buffer as compared to StdPSM where signaling frames are only generated when data frames are buffered at the AP.

Due to the small amount of traffic generated by the AC_BE and AC_BK applications, most of the times there are no frames buffered but the AP still has to send a QoS Null with the EOSP set to 1 to finish the SP. The transmission of these QoS Nulls adds useless additional load in the wireless channel, resulting in an overall worse system performance.

Power Saving Cost

In this section we analyze which are the costs of the two power saving schemes under study in terms of signaling load and wireless LAN channel capacity.

Figure 8 shows the total signaling load, this means ei-

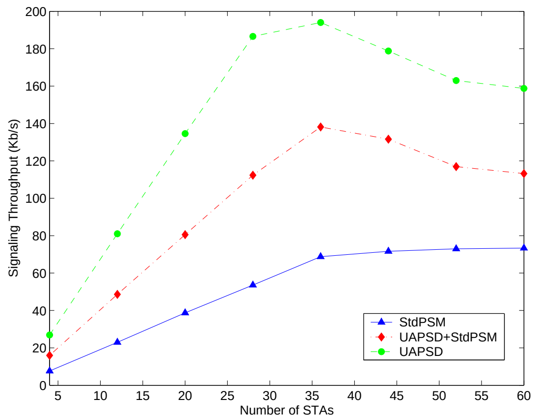

Figure 8: Impact of the number of stations on the signaling load

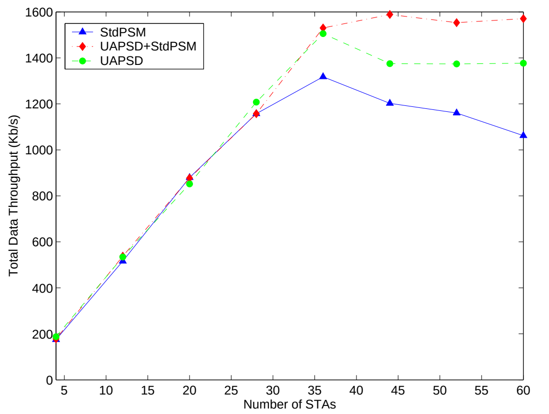

Figure 9: Impact of the number of stations on the WLAN channel

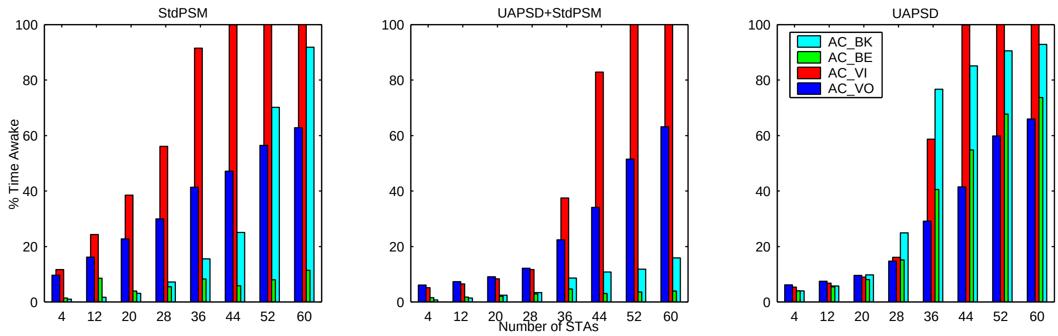

Figure 10: Impact of the number of stations on the Power Saving Efficiency

ther PS-Poll or QoS Null frames, introduced in the channel by the different mechanisms. The results show that the StdPSM case is the one that introduces the minimum signaling load. The reason is again that in 802.11 standard power save mode a PS-Poll is only sent when the station is notified that a frame is buffered in the AP, therefore useless signaling load is never introduced. In the UAPSD+StdPSM configuration a higher signaling load is observed mainly because of the AC_VO stations that keep on sending QoS Null frames within the silence periods. Note that sending a QoS Null if there is no frame buffered will generate another QoS Null response from the AP and this is additional signaling load. In the UAPSD configuration we have, in addition to the previously commented effect, the AC_BE and AC_BK stations sending often QoS Nulls when there is no traffic buffered for them at the AP.

The additional signaling load introduced in the U-APSD approaches is the price paid in order to fulfill the QoS requirements of the applications while being in power save mode. The effect of the Max_SP_Length helps reducing the signaling load by allowing the AP to transmit more than one data frame in response to a trigger frame but, in our experiments, this is not noticeable because the stations generate triggers too often to have many frames waiting in the AP's power save mode buffer.

Figure 9 shows that even if the signaling load is lower for the StdPSM configuration, the wireless channel capacity increases when using U-APSD approaches mainly because of the reduction in the downlink delay observed in the previous section. Here we can see again that the UAPSD+StdPSM configuration outperforms the UAPSD one when reaching saturation because of a lower signaling load.

Power Saving Efficiency

Figure 10 shows the percentage of time the stations are spending in active mode. The results show that the fact of using U-APSD affects in a different way the AC_VO and AC_VI stations as compared to the AC_BE and AC_BK stations, and the reason is the different generated traffic load of each AC. On the one hand AC_VO and AC_VI clearly benefit from using U-APSD, specially when the network is not yet congested. On the other hand, AC_BE and AC_BK benefit when AC_VO and AC_VI are configured to use U-APSD because there is less congestion after the Beacon frames but their performance gets much worse when they use U-APSD. The reason for the degraded performance of AC_BE and AC_BK when configured to use U-APSD comes from the fact that the Web and E-mail applications generate a low traffic load in both uplink and downlink. Because of that, if these ACs are configured to use StdPSM, they are almost all the time in the doze state and only awake when there is a burst of traffic coming from the downlink. However, when these ACs are configured to use U-APSD, they must send triggers periodically (every 100 ms) to check whether the buffer is empty or not.

Usually, when the AC_BE/AC_BK stations send a trigger, there will not be any frame in the AP's power save mode buffer and the AP will have to generate a QoS Null frame to finish the SP. This frame exchange is very expensive in terms of power saving due to the low priority that both the trigger sent by the station and the QoS Null sent by the AP have when trying to access the channel. This is noticeable specially in saturation conditions where the difference between the U-APSD case and the other ones, for the AC_BE and AC_BK stations, is very significant.

5 Summary & Conclusions

The upcoming mobile devices including wireless LAN access capabilities introduce new technological challenges that need to be addressed. We identified the combination of the 802.11e mechanisms EDCA and U-APSD as a key element toward the introduction of wireless LAN capabilities in the short-term in battery-limited mobile devices. The superposition of the U-APSD mechanism on top of the EDCA one though required to be studied to assess whether the additional complexity was justified by a significant QoS and power saving improvement as compared to simply using 802.11 power save mode.

The first requirement to proceed to the study of the U-APSD mechanism was to have a good understanding of its functionality and to design an algorithm in order to make use of it. We provided a thorough description of the U-APSD mechanism and its operation together with our proposed implementation of a static U-APSD algorithm. Based on that, we evaluated the performance of the combination of the EDCA QoS mechanisms with U-APSD or 802.11 power save mode via simulation. Through this analysis, an insight on the impact of U-APSD or 802.11 power save mode over the EDCA QoS functionality and its causes was acquired. Therefore, the results of this study are twofold. First, we provided quantitative results of the performance differences to be expected when U-APSD is used as compared to 802.11 power save mode with respect to QoS perceived by the users, power saving efficiency, signaling load and available channel capacity. Second, we provided the reasoning behind the performance differences pointing out the elements to be taken into account when deciding which power saving mechanism to use.

The main conclusions that can be drawn from our results are i) U-APSD significantly outperforms 802.11 power save mode in all considered performance metrics ii) the U-APSD performance metrics improvement results not only in better QoS and power saving but also in a considerable higher number of users/applications that can be accepted for the same channel capacity and iii) the usage of U-APSD for ACs which are not expected to have frames buffered frequently at the AP presents some particularities that have to be considered or otherwise use 802.11 power save mode.

6 Acknowledgments

The authors would like to thank their colleagues from the System Platforms Research Laboratories from NEC in Japan for inspiring the direction of this research work: Morimoto, Omori, Sashihara and Hirano.

REFERENCES

[1] 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 Standard 802.11e, November 2005.

[2] S.Mangold, S.Choi, G.R.Hiertz, O.Klein, and B.Walke. "Analysis of IEEE 802.11e for QoS Support in Wireless LANs". IEEE Wireless Communications Magazine, December 2003.

[3] IEEE 802.11 WG. "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications". IEEE Standard 802.11, June 1999.

[4] Wi-Fi Alliance. Http://www.wi-fi.com.

[5] A. Banchs and X. Perez-Costa. "An Assured Rate Service Extension for 802.11 Wireless LAN". In *Proceedings of IEEE Wireless Communications and Networking Conference (WCNC)*, Orlando, FL, March 2002.

[6] R.Krashinsky and H.Balakrishnan. "Minimizing energy for wireless web access with bounded slowdown". In *Proceedings of the eighth Annual International Conference on Mobile Computing and Networking (MOBICOM)*, September 2002.

[7] D.Qiao and K.G.Shin. "Smart Power Saving Mode for IEEE 802.11 Wireless LANs". In *Proceedings of IEEE INFOCOM*, March 2005.

[8] Y.Chen, N-Smavatkul, and S.Emeott. "Power Management for VoIP over IEEE 802.11 WLAN". In *Proceedings of IEEE WCNC*, March 2004.

[9] X.Pérez-Costa T.Sashihara and D.Camps-Mur. "Analysis of the Integration of IEEE 802.11e Capabilities in Battery Limited Mobile Devices". Wireless Communications Magazine.

[10] X.Pérez-Costa and D.Camps-Mur. "APSM: Bounding the Downlink Delay for 802.11 Power Save Mode". In Proceedings of IEEE International Conference on Communications (ICC).

[11] OPNET Technologies. Http://www.opnet.com.

[12] F.H.P. Fitzek and M.Reisslein. "MPEG-4 and H.263 Video Traces for Network Performance Evaluation". IEEE Network,Vol. 15,No. 6,pages 40-54, November/December 2001.