An Overview of MAC Energy-saving Mechanisms in Wi-Fi

Esther Guérin[a,*], Thomas Begin[a] and Isabelle Guérin Lassous[a]

[a] *Univ Lyon, UCB Lyon 1, ENS Lyon, Inria, CNRS, LIP UMR 5668, Lyon, France*

ARTICLE INFO

Keywords:
IEEE 802.11
TIM
APSD
TWT
legacy PSM

ABSTRACT

The IEEE 802.11 standard has rapidly become the *de facto* standard for Wireless Local Area Networks (WLANs) under the commercial name Wi-Fi. Despite IEEE's regular efforts to improve the energy efficiency of Wi-Fi, many energy-saving mechanisms remain relatively unknown to the community, which contributed to hamper their large-scale adoption by chipset manufacturers.

In this paper, we describe the main MAC mechanisms that have been introduced by the IEEE 802.11 standards (and its amendments) for the sake of saving energy. This includes diverse mechanisms where either the stations (STAs) prompt the access point (AP), the AP notifies the STAs, or the STAs and the AP make an appointment. In every case, we illustrate the mechanism's principle through examples. We also present a set of metrics to compare the effectiveness of energy-saving mechanisms with each other and discuss how the energy-saving mechanisms of IEEE were adopted by the Wi-Fi Alliance, and hence by chip manufacturers.

Overall, we believe that this paper, which accurately details these mechanisms, can help researchers, manufacturers, and network operators to have a clear understanding of what Wi-Fi is readily capable of doing when it comes to energy efficiency, and thus foster their use in actual and upcoming WLANs.

1. Introduction

Following its release in 1999 by IEEE[1], and the commercialization efforts made by the Wi-Fi Alliance, the IEEE 802.11 standard has rapidly become the *de facto* standard for Wireless Local Area Networks (WLANs) under the commercial name Wi-Fi. Over the years, Wi-Fi has continuously been improved through multiple amendments to this standard, resulting in new usage in a wide variety of environments and growing demand from end-users. Lately, the 802.11ac and 802.11ax amendments have leveraged new technological progress to push the data rate of wireless links beyond the Gbps range. Overall, Wi-Fi has established itself as a versatile, mature, high throughput, economical, and commercial off-the-shelf technology which fulfills most, if not all, current needs of WLANs.

However, Wi-Fi needs to further evolve in order to contribute to the energy footprint reduction of Information Technology (IT) as well as better accommodate the development of Internet of Things (IoT) where a growing number of devices with limited energy capabilities must exchange data over networks. It is worth noting that, since its inception in 1999, efforts have been made to improve the energy efficiency of Wi-Fi. But these efforts were often of secondary importance and thereby overlooked by chipset manufacturers. A case in point is the 802.11ah amendment (marketed as Wi-Fi HaLow) which was designed in 2016 to handle IoT networks but was never really implemented by chipset manufacturers (with the exception of Newracom in 2020).

The mechanisms meant to improve energy efficiency in Wi-Fi are many and based on a variety of principles. They have appeared in various amendments spanning over more than 20 years, which hindered their understanding and interoperability, and hampered their large-scale adoption by chipset manufacturers. Furthermore, a lack of documentation summarizing these mechanisms has also prevented researchers, manufacturers, and network operators to have a clear understanding of what Wi-Fi is readily capable of doing when it comes to energy efficiency.

In this paper, we describe the main mechanisms that have been introduced by the IEEE 802.11 standards to improve energy efficiency. Given the immensity of the endeavor, we restrict our study to mechanisms that affect the way stations (STAs) and access points (APs) postpone their frames for transmission with the aim of reducing their energy consumption. Therefore, our study does not include mechanisms such as TXOP Power Save (also known as micro-sleep), which was introduced in IEEE WG 802.11ac (2013) (Section 10.2.1.19) and lets STAs doze as soon as

---

✉ `esther.guerin@ens-lyon.fr` (E. Guérin); `isabelle.guerin-lassous@univ-lyon1.fr` (I. Guérin Lassous)
ORCID(s):

[1] The first version of the IEEE 802.11 standard was released in 1997 and then revised in 1999.

they detect a transmission for another Wi-Fi devices, nor the wake-up radio, which will be introduced in the 802.11ba amendment.

The remainder of the paper is organized as follows. In Section 2, we briefly review the main principles of the MAC layer that are directly affected by the energy-saving mechanisms. Sections 3, 4 and 5 precisely describe the existing energy-saving mechanisms dedicated to the downstream traffic, upstream traffic, and both traffic, respectively. Section 6 introduces a set of metrics to ease the comparison of MAC energy-saving mechanisms, and discusses how the energy-saving mechanisms were adopted by the Wi-Fi Alliance. In Section 7, we discuss the related work. Section 8 concludes this paper.

2. Preliminaries on Wi-Fi MAC rules for energy saving

To ensure the sharing of the radio medium resources, Wi-Fi devices have to follow rules. For a clear and complete description of these rules, we refer the interested reader to Gast (2013). In this section, we rapidly recall the main rules that are directly affected by the energy-saving mechanisms.

Our study focuses on the *Medium Access Control* (MAC) layer because this is where the devices schedule their frame transmissions. The IEEE 802.11 standards propose several MAC options for WLANs, but we restrict our study to WLANs operated in *infrastructure mode* (described in Section 2.1) since it is the most common Wi-Fi infrastructure. For a WLAN in infrastructure mode, there are two functions to distribute access to the medium. Our study focuses on the *Distributed Coordination Function* (DCF), which means that the devices use the *Carrier Sense Multiple Access / Collision Avoidance* (CSMA/CA) technique to send a frame. We briefly describe CSMA/CA in Section 2.2. Finally, Section 2.3 presents the *Orthogonal Frequency-Division Multiple Access* (OFDMA), a transmission technique recently introduced in Wi-Fi that can be used alongside CSMA/CA. In appendix, Table 5 reports the main terms and abbreviations used in this article along with the section in which they are described.

2.1. Infrastructure mode

In infrastructure mode, WLANs are composed of a fleet of APs, each AP managing a set of associated STAs. Then, we distinguish two types of traffic: *downstream* (from the AP to a STA) and *upstream* (from a STA to the AP) traffic.

At the MAC layer level, a STA can save its power by temporarily disabling its radio: it is then said to be in the *doze state*, as opposed to the *awake state*. When a STA is dozing, it cannot emit nor receive any radio signal. In short, the rationale behind all the energy-saving mechanisms presented in this paper is to postpone frame transmissions to periods of time during which the concerned STA is awake. If done properly, the STA can safely doze in between these periods without missing any frame transmission. Note that in the IEEE standards, APs are assumed to be always awake and always ready to receive frames since they are plugged into the mains.

In infrastructure mode, the AP regularly broadcasts a *beacon* to advertise the WLAN to potential new STAs, but also to transmit pieces of information to the already associated STAs. Thus, STAs are required to listen regularly to beacons. In practice, each STA must define during its association with the AP a *Listen Interval* corresponding to the maximum number of beacons the STA may skip. Note that a STA may indicate a Listen Interval of 0 if it intends to receive all beacons.

2.2. CSMA/CA

The DCF uses the CSMA/CA technique. Its principle is that if a device wants to send a frame, it first listens to the medium. If the latter is idle, the device can send its frame immediately. Otherwise, the device will wait until the end of the current transmission, after which it will draw a random backoff to try to access the medium. Note that the medium can be busy either physically or virtually. Physically means that there is an ongoing radio signal on the medium. Virtually means that another Wi-Fi device previously reserved the medium for a transmission. Indeed, a STA may reserve the medium for a certain duration by exchanging special frames with the AP, called *Request To Send / Clear To send* (RTS/CTS). The duration of the reservation is indicated in these RTS/CTS frames. As soon as a STA receives either frame, it updates a timer called *Network Allocation Vector* (NAV) whose duration is set to the value contained in the received frame (provided that this value is greater than the current NAV value). As long as its NAV is not null, the STA must remain silent.

Clearly, a STA that just woke up (*i.e.*, that just switched from doze state to awake state) cannot know if there is an ongoing reservation made by an RTS/CTS. Therefore, a STA that just woke up has to wait either for the transmission of another device or for the expiration of the *Probe Delay* (IEEE WG 802.11 (1999), Section 11.2.1.1). On one hand,

Mechanism	Type of traffic	CSMA/CA or OFDMA	Amendment	Year
Ps-Poll	Downstream	CSMA/CA	Legacy	1999
TIM	Downstream	CSMA/CA	Legacy	1999
APSD	Downstream	CSMA/CA	802.11e	2005
TIM segmentation	Downstream	CSMA/CA	802.11ah	2016
RAW	Upstream	CSMA/CA	802.11ah	2016
Individual TWT	Upstream & downstream	CSMA/CA	802.11ah	2016
Trigger-enabled TWT	Upstream & downstream	OFDMA	802.11ax	2021
Broadcast TWT	Upstream & downstream	OFDMA	802.11ax	2021
rTWT	Upstream & downstream	OFDMA	802.11be	TBA

Table 1
Overview of MAC energy-saving mechanisms in Wi-Fi. The term "Legacy" refers to the first IEEE 802.11 standard.

the transmission of another device informs the STA that the medium is now reserved and the STA can then fall back to the usual procedure to access the medium. On the other hand, the expiration of the Probe Delay guarantees that any previous reservation has expired too. Note that the 802.11 standard does not specify any value for the Probe Delay. It is left to the manufacturers to find a compromise between delays and collision risks.

If two devices (STAs or AP) access the medium at the same time, their transmission may collide. In case of a collision, one or both frames become unintelligible, leading to retransmissions. Collisions represent a factor of energy inefficiency, in addition to being a performance issue as well. Note that when the probability of collision becomes non-negligible, the WLAN is said to be in contention.

2.3. OFDMA

In contrast with CSMA/CA whose aim is to have at most one transmission on a shared radio medium at any time, OFDMA is a transmission technique that bypasses this restriction under certain conditions described hereafter. The medium is divided into multiple blocks called *Resource Unit* (RU). One RU can be allocated to only one STA at a time, however different RUs inside an OFDMA transmission can be allocated to different STAs.

For downstream and upstream OFDMA transmissions, it is up to the AP to decide the scheduling of OFDMA transmissions as well as the number of RUs to allocate to each STA. In the case of downstream OFDMA transmissions, the AP knows exactly the amount of data needed to be sent to each STA and the AP can allocate the RUs accordingly. In the case of upstream OFDMA transmissions, the AP does not know the amount of data the STAs have to send. To address this issue, the STAs must notify the AP that they have frames to send, or the AP asks the STAs directly.

For the sake of backward compatibility – pivotal to Wi-Fi – the new OFDMA technique must be integrated into the old CSMA/CA procedure. In other words, the AP needs to make sure the medium is idle before initiating an OFDMA transmission. Thus, the AP applies the CSMA/CA technique to gain access to the medium, and then it sends an *OFDMA trigger frame* to schedule an OFDMA transmission. The OFDMA trigger frame notifies the STAs of the upcoming OFDMA transmission. It also carries information with respect to the allocations of the different RUs and the physical layer specifications to use them.

Besides allocating RUs, the AP can choose to leave some RUs "free": these RUs are for *Random Access* (RA). Any STA can try to use these so-called Random Access RUs following an ALOHA-like manner.

Note that the introduction of OFDMA in Wi-Fi is relatively recent, namely in 2021. Therefore, most of the existing energy-saving mechanisms presented hereafter rely on classical CSMA/CA. In fact, only Trigger-enabled TWT and Broadcast TWT (presented in Sections 5.3 and 5.4, respectively) are tailored to OFDMA.

2.4. Organization of the paper

The restrictions of the MAC energy-saving mechanisms may apply only to the downstream traffic, to the upstream traffic, or to both traffic. Based on this classification, we detail these mechanisms in Sections 3, 4 and 5, respectively. However, most of the MAC energy-saving mechanisms can operate simultaneously on the same WLAN with some STAs running a given mechanism and others running another mechanisms. Table 1 reports the different mechanisms presented hereafter, the type of traffic on which they operate, as well as the amendment in which they were first introduced.

3. Energy-saving mechanisms dedicated to downstream traffic

In this section, we describe energy-saving mechanisms that affect the schedule of the downstream traffic. Postponing the downstream traffic can help STAs save energy, but it often leads to lower throughput and longer delays. Note that STAs like laptops or smartphones generally favor performance and hence do not wish to delay the download of their frames. It is up to each STA to express its choice via the power management mode, which is a field in the MAC header of the frames. There are two power management modes: *active mode* and *Power Save (PS) mode*. If a STA declares itself in active mode, the AP may send the frames intended for the STA at any time (without additional delays). On the other hand, if a STA switches to the PS mode, the AP will not send the frames intended for the STA as soon as possible. Instead, the AP will buffer the frames and wait for a specific time to deliver them to the STA. There are two approaches for deciding when the frames must be delivered: either the STA prompts the AP, or the two of them agree on an appointment. The two approaches are described in Sections 3.1 and 3.2, respectively. With such approaches, a STA knows when the AP may send frames that are intended for it, and the STA can doze the rest of the time.

In addition to these two approaches, a third approach consists of the AP notifying a STA if it has pending frames in its buffer. This notification allows a STA that does not have pending frames to quickly return to the doze state. Section 3.3 describes this type of mechanisms.

3.1. Prompting the AP

The first approach for a STA to receive data frames pending in the buffer of the AP is to explicitly prompt the AP: the STA decides by itself when it is ready to receive incoming transmissions.

3.1.1. PS-Poll

PS-Poll is the first mechanism that was introduced in the legacy version of the 802.11 standards for the sake of energy saving (IEEE WG 802.11 (1999), Section 11.2.1.1). This mechanism must be implemented on any AP or STA that wishes to switch to PS mode.

When a STA wishes to receive its pending frames (if any) from the AP, the STA prompts the AP by sending a special frame called PS-Poll. When the AP receives a PS-Poll frame there are two possible scenarios. If there are pending frames for this STA in its buffer, then the AP sends one data frame to the STA (which will be acknowledged by the STA). Otherwise, the AP returns an ACK.

STAs can freely choose the time at which they send their PS-Poll frames with one restriction: each STA needs to prompt the AP at least once per Listen Interval. This obligation stems from the fact that the STA cannot know if there are any pending frames in the buffer of the AP. Thus, each STA must give the AP regular opportunities to send pending frames so that they do not stay for too long in the buffer of the AP. Of course, a STA may prompt the AP much more often than once per Listen Interval in order to shorten the delays. Overall, the STAs are in charge of setting the PS-Poll mechanism since the only parameter to be tuned is the Listen Interval whose value is set by the STAs themselves.

Figure 1 illustrates an example of exchanges between a STA and its AP. In the first Listen Interval, the STA wakes up, and since the medium is idle, the STA can send a PS-Poll frame after a Probe Delay. However the AP does not have any pending frames for this STA, so it sends a simple ACK. Then the STA returns to the doze state. The second Listen Interval displays the exchange when there are a pending frame for the STA. The STA sends a PS-Poll frame after a Probe Delay, to which the AP responds with one pending frame. In the third Listen Interval, the STA wakes up but the medium is busy (*e.g.*, due to transmission of another STA or of the AP). The STA must then wait to gain access to the medium before prompting the AP. The rest unfolds as in the second Listen Interval.

At the release of the legacy standard in 1999, there was no frame aggregation (every data frame was individually acknowledged), nor *Transmission Opportunity* (TXOP) (a device had to gain back the access to the medium after each ACK). Thus, with the PS-Poll mechanism, the AP answers a PS-Poll frame with a single pending frame, which in turn is acknowledged by the STA before the STA can try to send another PS-Poll frame to retrieve another pending frame. Even if it remains pending frames in the buffer of the AP, the STA must contend again for the medium for each PS-Poll frame. When the AP answers a PS-Poll frame with a data frame, it can indicate via the field “More Data” in the MAC header if there is at least another pending frame for the STA. Therefore, a STA receiving a data frame knows if it is up-to-date or if it needs to send yet another PS-Poll frame.

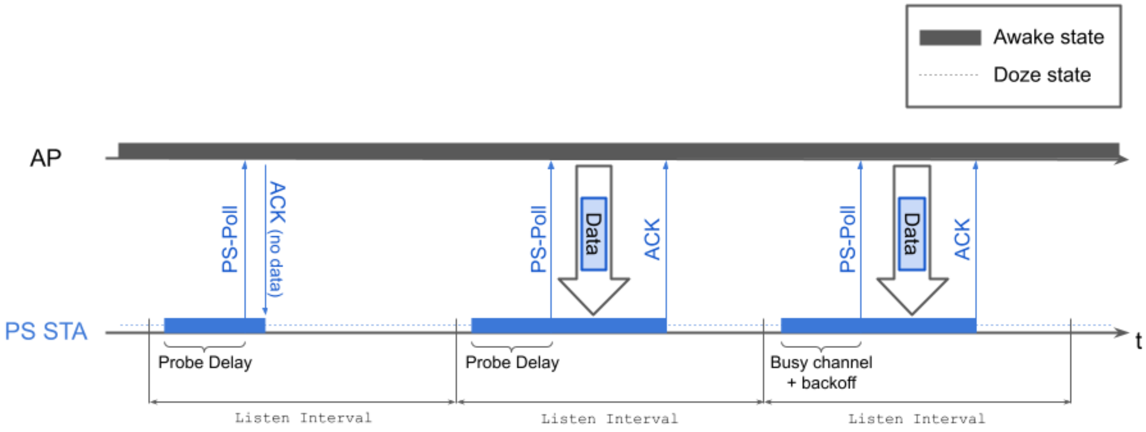

Figure 1: Example of frame exchanges with PS-Poll.

3.1.2. U-APSD

Introduced by the amendment 802.11e (IEEE WG 802.11e (2005), Section 11.2.1.4), *Automatic Power Save Delivery* (APSD) brought novel key concepts. APSD can come in two forms: scheduled (S-APSD) and unscheduled (U-APSD). Whereas U-APSD acts as an enhanced PS-Poll, S-APSD relies on a different approach described in Section 3.2.

The progress brought by U-APSD over PS-Poll is two-fold. First, U-APSD can use an alternative to the PS-Poll frame called *APSD trigger frame*. Recall that a PS-Poll frame is a MAC Management frame whose utility is solely to prompt the AP. With U-APSD, the AP does not need a special frame to send pending frames. Rather, it considers any upstream data frames as a trigger to send downstream pending frames (if any) to the associated STA. The term "APSD trigger frame" refers more to an abstraction than to a particular frame. The APSD trigger frame avoids the use of a separate PS-Poll frame. If the STA has no data to send to the AP, then it may use a *Null Data Packet* (NDP) to prompt the AP.

Second, U-APSD introduces the *Service Period* (SP): it is a period of time during which the STA stays awake to receive its frames. In Section 3.1.1, we explained that with the PS-Poll mechanism the AP must be prompted as many times as there are pending frames in the buffer of the AP. With U-APSD, a STA that sends an APSD trigger frame enters a SP, receives all of its pending frames, and exits the SP. The AP indicates the end of the SP via a flag called "*End of Service Period*" (EOSP): when it is raised, the STA knows that the AP will stop sending frames and the STA can return to the doze state. Note that the parameter "Max SP length", which is selected by the STA and indicated to the AP during the association, limits the maximum number of frames exchanged over a SP. If this parameter is set to zero, then the AP can send all the frames pending in its buffer during one SP.

Figure 2 illustrates the advantages of U-APSD over PS-Poll. With PS-Poll, the STA must prompt the AP three times, each time contending against the other STAs to gain access to the medium. With U-APSD, the STA has to gain access only once to prompt the AP and then it receives its three pending frames in a row.

Note that APSD was initially meant to be used alongside the Access Categories defined in the same amendment. In one STA, different Access Categories can use different energy-saving mechanisms, and each Access Category has its own set of parameters for the energy-saving mechanism it uses. For example, a STA can decide the following scheme: its first Access Category uses S-APSD, its second Access Category uses U-APSD whereas its third and fourth Access Categories use the PS-Poll mechanism.

In the rest of the paper, "prompting the AP" refers to either sending a PS-Poll frame or sending an APSD trigger frame.

3.2. Making an appointment between STAs and AP: the S-APSD mechanism

Another approach for a STA to retrieve its frames pending at the buffer of the AP is to make an appointment with the AP. The STA can agree with the AP beforehand on a time period during which the pending frames will be sent.

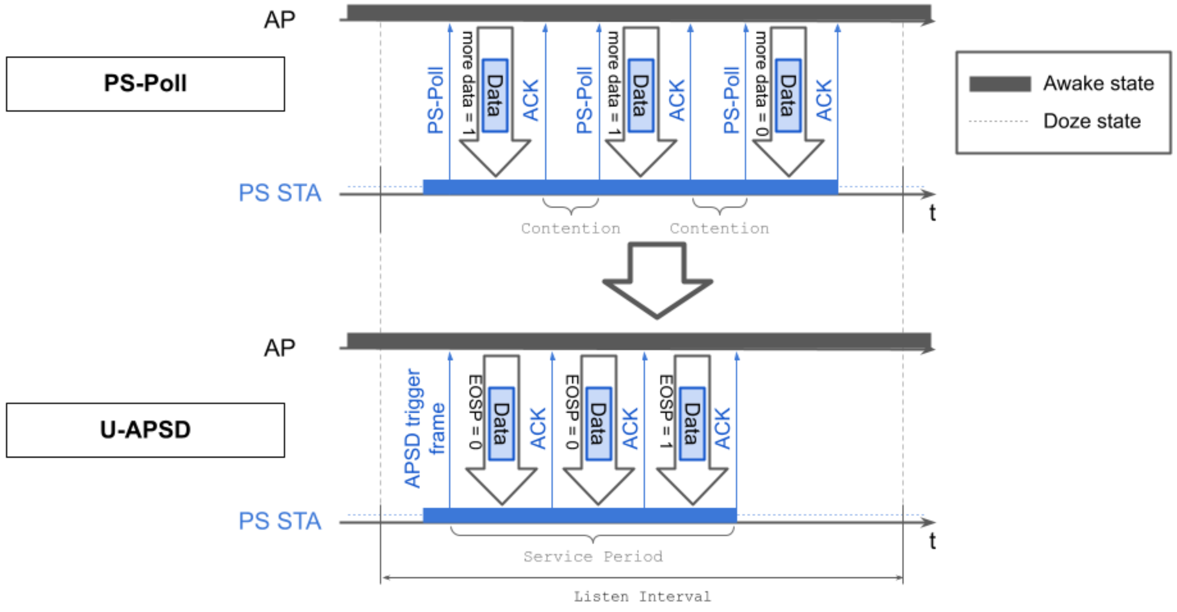

Figure 2: Comparison between PS-Poll and U-APSD.

The first and only attempt in this direction is the S-APSD mechanism (IEEE WG 802.11e (2005), Section 11.2.1.4). Similarly to U-APSD (see Section 3.1.2), S-APSD defines SPs which are periods of time during which all the pending frames are sent at once. However, contrary to U-APSD that triggers the SPs at the initiative of the STA, S-APSD schedules the SPs in advance, at the initiative of the AP. S-APSD can work in two modes: the *implicit mode* and the *explicit mode*.

In implicit mode, the appointments are separated by equal intervals of time. An appointment between a STA and the AP will result in a SP only if there are pending frames for this STA in the buffer of the AP. When the time of an appointment comes, the STA wakes up and waits for the AP, but only for a short time. If the AP does not start a transmission for the STA during this short time, the STA returns to doze state. Otherwise, the STA enters a SP, receives its pending frames, and returns to doze state when the SP ends. Since the appointments are regular in frequency, the STA can compute the time of the next appointment by itself. In a similar fashion, the AP can deduce when the STA is awake and able to receive frames, and when it is not. The appointment before the SP allows the STA not to waste energy by needlessly waiting for the AP.

In explicit mode, any appointment always turns into a SP. As soon as the STA wakes up for an appointment, the SP starts. If the AP have pending frames in its buffer, it can send them at any time during the SP. During the SP, the AP must also indicate the schedule of the next appointment to the STA. Therefore, the STA must wait for the AP to end the SP and schedule the next appointment before returning to doze state.

Whether in implicit or explicit mode, once the SP starts, it unfolds like the SPs in U-APSD. The AP sends all the pending frames – or the maximum number it is authorized to – to the STA and raises the EOSP flag on the last frame to indicate to the STA the end of the SP.

Regarding the parameters of S-APSD (in implicit or explicit mode alike), each STA may set bounds (lower and upper) to the time intervals between the SPs. However, their final value as well as the scheduling of the SPs are chosen by the AP. As a result, the actual needs of the STAs may not be precisely taken into account by the AP.

Figure 3 shows the process of S-APSD with implicit appointments. The STA wakes up at constant time intervals that were decided beforehand by the AP. The first time, the AP does not send any frames to the STA so it returns quickly to the doze state. The second time, the AP sends one pending frame with the EOSP flag raised. The STA returns an ACK before returning to the doze state. The third time, the AP sends two pending frames. The STA stays awake for the transmissions and then dozes right after receiving the second frame with the EOSP flag raised. Note that in implicit mode, the beginning of SPs occurs at a regular pace, but their ending may be irregular.

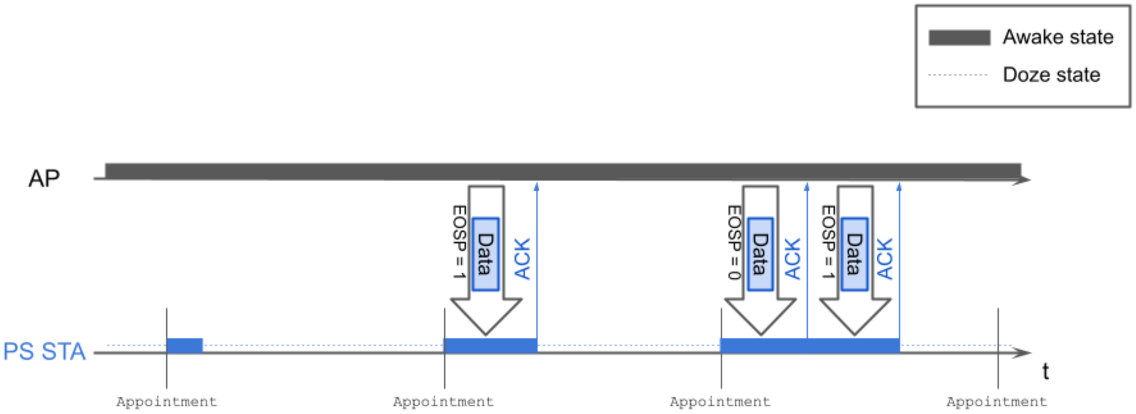

Figure 3: Example of frame exchanges with S-APSD in implicit mode.

3.3. Notifying the STAs of pending frames in the AP

Mechanisms relying on prompting the AP (like those discussed in Section 3.1) may result in STAs prompting often and unnecessarily the AP. This section describes the existing mechanisms in the 802.11 standards that let the AP notify the STAs of the presence or the absence of pending frames in the buffer of the AP.

3.3.1. TIM

The *Traffic Indication Map* (TIM) was introduced in the legacy standard as a complement to PS-Poll (IEEE WG 802.11 (1999), Section 11.2.1.3); the combination of TIM and PS-Poll is generally referred to as "*Legacy Power Save Mode* (PSM)". TIM is a bitmap broadcast by the AP inside its beacons. Each of its bits corresponds to a STA, and by setting it to 1 or 0 the AP indicates whether it has or not in its buffer pending frames for the corresponding STA. A STA can wake up just to receive a beacon and collect its bitmap. If the STA's corresponding bit is 0, the STA will return to the doze state since it does not have any pending frames. Otherwise, it tries to access the medium to prompt the AP.

Figure 4 shows how TIM can work with PS-Poll. STAa and STAb wake up to listen to the first beacon. The TIM bitmap of the first beacon indicates that only STAa has pending frames, so STAa sends a PS-Poll frame to retrieve its pending frames from the AP. At the same time, STAb returns to the doze state without sending a PS-Poll frame because STAb knows that it does not have any pending frames to retrieve. Both STAs wake up again for the second beacon. This time, both STAs have pending frames in the buffer of the AP. So they contend for the medium to send a PS-Poll frame. STAb gains access to the medium, sends a PS-Poll frame, and receives its pending frame. Then, STAa succeeds in sending its PS-Poll frame and receives its pending frame.

3.3.2. TIM segmentation

The *TIM segmentation* mechanism is defined and only used in the 802.11ah amendment which was designed for the Sub-1GHz radio band (IEEE WG 802.11ah (2016), Section 10.47). TIM segmentation extends the legacy TIM, by splitting the TIM bitmap in multiple chunks that are broadcast separately.

The 802.11ah amendment defines the notion of *short beacon* to segment TIM. A short beacon is a kind of beacon that carries only a fraction of the information of a regular beacon. In a 802.11ah WLAN, both regular and short beacons are broadcast by the AP. Short beacons are broadcast regularly with a time period shorter than that for the regular beacons. In other words, the regular beacon interval is divided into multiple short beacon intervals (of equal duration).

In parallel, STAs are divided into groups. For each group, the AP broadcasts one short beacon that carries a TIM bitmap dedicated to the STAs of this group. In between two regular beacons, there is one short beacon with a TIM bitmap for every group of STAs. The regular beacon itself carries another kind of TIM bitmap called *Delivery TIM* (DTIM). The bits of the DTIM bitmap do not correspond to a STA, but rather to a group of STAs. If the bit is set to 1,

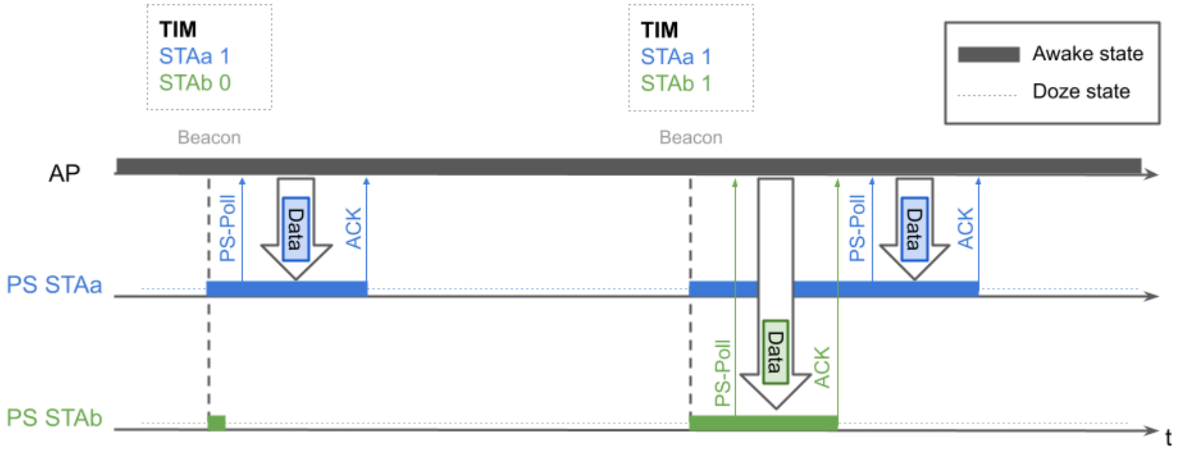

Figure 4: Example of frame exchanges with TIM bitmap. Note that STAs use PS-Poll frames in the figure, but they could also use an APSD trigger frame.

the AP indicates that at least one STA of the corresponding group has a pending frame. If the bit is set to 0, it means that no STA of the corresponding group has a pending frame.

The AP uses a mechanism named *Page Slicing* to divide the STAs into TIM groups. TIM groups contains STAs whose *Association ID* (AID) are consecutive. At each regular beacon, the AP informs the STAs of the TIM groups for the subsequent beacon interval, meaning that the TIM groups may change from one beacon interval to another. Alternatively, the Association ID of a STA can be changed through a reassociation process. Doing so can be particularly useful for the AP to efficiently gather STAs following departures or arrivals of STAs in the WLAN.

Figure 5 represents the TIM segmentation. In this example, there are two groups of STAs: STAa and STAb belong to Group1 (in blue), and STAc and STAd to Group2 (in green). Since there are two groups, there are two short beacons within a regular beacon interval. In the first regular beacon, the DTIM announces that the two groups have pending frames. So the STAs of Group1 wake up only for the first short beacon, and the STAs of Group2 wake up only for the second short beacon. The first short beacon has a TIM bitmap for Group1, and both STAs have pending frames. So STAa and STAb contend to prompt the AP. The second short beacon is for Group2, and its TIM bitmap indicates that only STAd has pending frames. Thus, STAd prompts the AP while STAc immediately returns to the doze state. The second regular beacon carries a DTIM that indicates that only Group1 has pending frames, and not Group2. Hence, STAs of Group1 wake up for the first short beacon, whereas the STAs of Group2 do not wake up for the second short beacon because they already know that none of them have pending frames.

The TIM segmentation was introduced in the 802.11ah amendment to increase the maximum number of associated STAs in WLANs, namely about four times more than in the other standards (see Tian, Santi, Seferagić, Lan and Famaey (2021)). However, the number of STAs may become too large for the corresponding TIM bitmap to be broadcast within a single beacon and this led to its segmentation.

4. Energy-saving mechanism dedicated to upstream traffic: RAW

The IEEE 802.11 Working Group has proposed a single energy-saving mechanism dedicated only to upstream traffic: the *Restricted Access Window* mechanism (RAW) (IEEE WG 802.11ah (2016), Section 10.22.5). RAW enforces speaking and silence time periods to STAs. In practice, the AP defines time periods, referred to as RAW slots, during which only a given group of STAs are allowed to emit. Said differently, a STA is not allowed to emit during another group's RAW slot. Note that the AP can decide of unrestricted time periods where every STA, regardless of their group, can access the medium. To a certain extent, the RAW mechanism can be viewed as TDMA overlaying CSMA/CA where STAs are allocated time periods to try to access the medium.

Like the TIM segmentation mechanism, RAW uses the Page Slicing to divide STAs into groups. More precisely, the RAW mechanism divides the set of associated STAs using a two-level separation: first into groups and then into

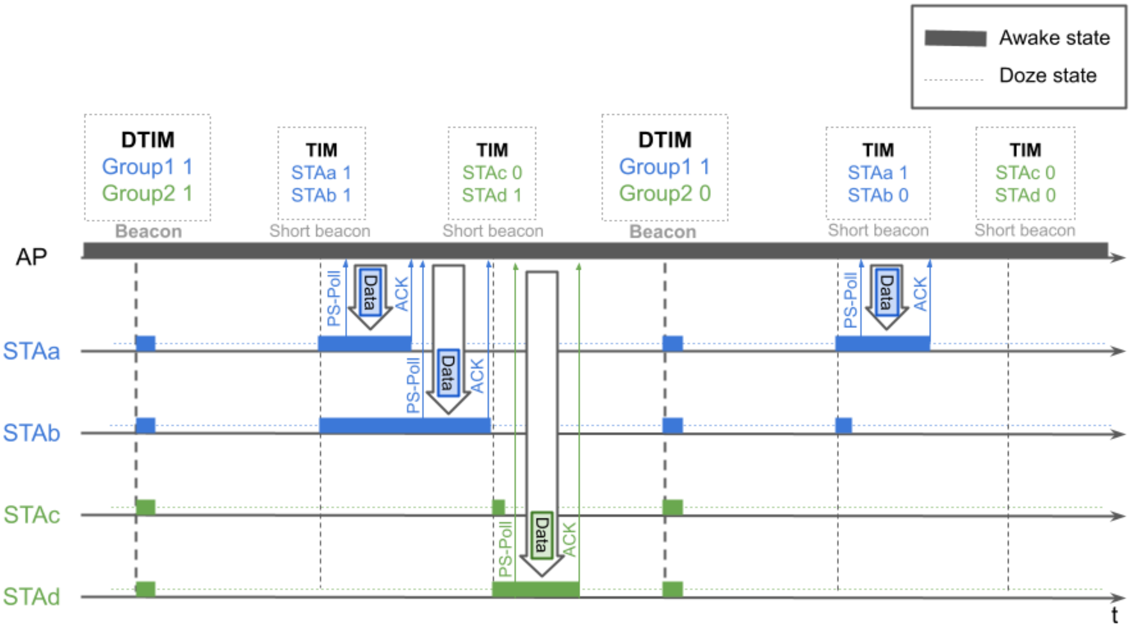

Figure 5: Example of segmented TIM bitmaps.

subgroups. Upon every beacon, the AP announces the composition of groups and subgroups for the upcoming beacon interval. Thus, the RAW groups may differ at each beacon interval. The AP defines a RAW group as a range of STAs having consecutive Association IDs, and then the number of RAW subgroups for each RAW group. Note that the STAs of a RAW group are distributed among its subgroups in a round-robin way. Finally, the duration of the beacon interval is divided between the groups to form the RAW slots, which are further themselves divided to form RAW subslots.

Figure 6 shows two ways of distributing 12 STAs in a beacon interval with RAW. During the first beacon interval, the AP divides the STAs into three groups, each group having four STAs. The first RAW slot is then further divided into four RAW subslots so that each STA of the first group has an exclusive subslot. The second RAW slot is divided into only two subslots. Thus, there are two STAs per RAW subslot. The third RAW slot is divided into three subslots. Following the round-robin way, the STAs with Association ID 9 and 12 have to share the first subslot. On the other hand, STAs 10 and 11 have exclusivity on their subslot. In the second beacon interval, the AP decides a different division with only two groups with respectively five and four subslots.

Note that a STA or the AP may change the Association ID by exchanging a (Re)association request / response (see Section 3.3.2).

A parameter named *Cross Slot Boundary* indicates if the AP authorizes the STAs to overflow into the next RAW slot. This parameter is contained in the beacons. When the Cross Slot Boundary bit is set to 0, a STA is not allowed to start a transmission unless this latter is guaranteed to finish before the end of its RAW slot. With this setting, STAs do not need to respect a Probe Delay at the start of their RAW slot. The STAs of a group can contend for the medium access right upon the beginning of their RAW slot since the STAs from the other RAW slots are guaranteed to be silent at that time. Otherwise, if crossing the slot boundary is allowed (*e.g.*, if the Cross Slot Boundary field is set to 1), the STAs of a group cannot contend for the medium directly at the beginning of their slots because of potential ongoing transmissions that started at the previous slot. In this case, STAs must wait for the Probe Delay to expire before emitting.

Besides the regular backoff, RAW introduces another kind of backoff, called *RAW backoff*. A STA uses the RAW backoff during its allocated RAW slots but the regular backoff during unrestricted periods (periods where any STAs can send frames upstream). Note that the RAW backoff is reset using the initial contention window at the beginning of each RAW slot. On the other hand, the regular backoff is put on hold during the RAW slots (including the ones for

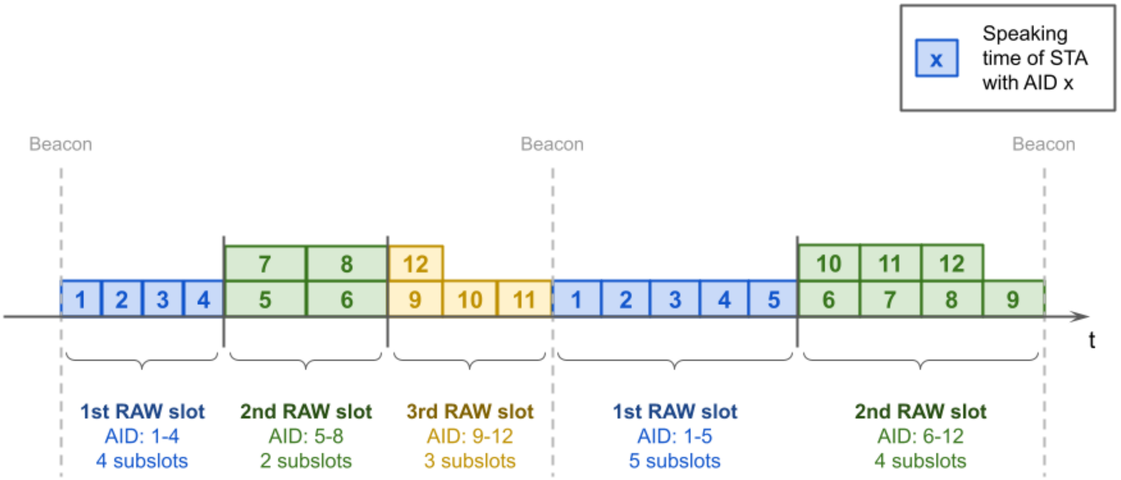

Figure 6: Examples of STAs divisions with RAW.

other STAs) and resumes at the beginning of the next unrestricted period. The RAW backoff ensures that the time of a RAW subslot is not wasted. The contention between RAW slots and unrestricted periods is different, and so should be their contention window.

5. Energy-saving mechanisms for both the downstream and the upstream traffic

The 802.11ah amendment was the first to propose a MAC energy-saving mechanism that handles both downstream and upstream traffic. It combines TIM segmentation and RAW to schedule the downstream and the upstream traffic, respectively. The details of this mechanism, which works only in the Sub-1GHz radio band, are presented in Section 5.1.

The 802.11ah amendment also introduced the Individual *Target Wake Time* (TWT) mechanism that manages both the downstream and the upstream traffics of STAs. Later, the 802.11ax amendment brought the Individual TWT mechanism into the 2.4 and 5GHz radio bands and added other types of TWT. Namely, there are three different types of TWT: Individual TWT is tailored to WLANs operating with CSMA/CA, whereas Trigger-enabled and Broadcast TWT leverage on OFDMA transmissions. These three types of TWT are described in Sections 5.2, 5.3 and 5.4, respectively.

5.1. RAW and TIM segmentation

Although RAW and TIM segmentation were described in separate sections (4 and 3.3.2), RAW was meant from the beginning to be used alongside the TIM segmentation.

Recall that both TIM segmentation and RAW rely on Page Slicing to divide STAs into different groups. TIM segmentation distributes STAs on a one-level hierarchy, whereas RAW uses two levels. The first level is identical between TIM and RAW: each TIM group has its dedicated RAW slot. Then, RAW further divides these slots into smaller subslots that are in turn allocated to the STAs of the corresponding group. Upon the reception of a short beacon, the STAs of the corresponding group discover whether they have pending frames in the buffer of the AP thanks to the segmented TIM. Then, subgroup per subgroup, the STAs can prompt the AP if necessary, and then send data frames upstream (if any).

Figure 7 shows how RAW and TIM segmentation are working together in a 802.11ah WLAN. We consider five STAs segmented in two groups, which correspond to RAW groups and TIM groups at the same time. STAa and STAb are in Group1 (in blue) whereas STAc, STAd and STAe are in Group2 (in green). For the sake of our example, we make the following assumptions: (i) Only STAa, STAc and STAd have data frames to send upstream; (ii) As for the downstream traffic, only STAc does not have pending frames in the buffer of the AP. The other STAs have exactly one pending frame in the buffer of the AP. At the beginning of our example, the AP broadcasts a regular beacon that notifies every STA of its respective (RAW and TIM) group and RAW subgroup. The regular beacon also carries the DTIM

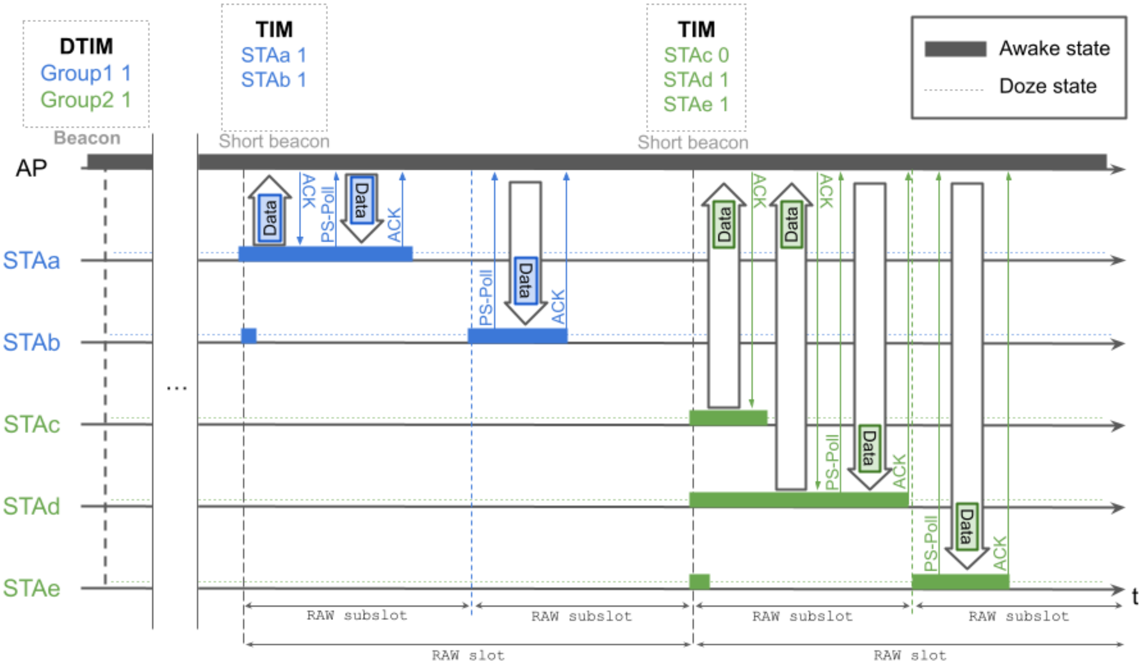

Figure 7: Example of RAW working with TIM segmentation.

bitmap that informs each group of the presence or not of pending frames in the buffer of the AP. In our example, both groups have pending frames, so every STA has to wake up again to listen to the short beacon of their group. The first short beacon is dedicated to Group1: it carries a segmented TIM bitmap for its STAs and also marks the beginning of a RAW slot. Group1 is divided into two RAW subgroups, which means that both STAa and STAb have their own RAW subslot. STAa accesses the medium without competition during the first RAW subslot. It sends a data frame upstream before sending a PS-Poll to retrieve its frames pending in the buffer of the AP. During the second RAW subslot, STAb also accesses the medium without competition to send a PS-Poll. The second short beacon ends the first RAW slot and marks the beginning of the second RAW slot, dedicated to Group2. Group2 is divided into two subgroups. Since there are three STAs in Group2, STAc and STAd share the same RAW subslot. Both STAc and STAd have data frames to send upstream, and STAd also have a pending frame in the buffer of the AP. As a consequence, both STAc and STAd try to access the medium during the RAW subslot they share. STAc accesses the medium first and sends its data frame upstream. Then STAd accesses the medium and sends a data frame upstream, before sending a PS-Poll to retrieve its pending frame in the buffer of the AP. During the second RAW subslot, STAe is the only STA allowed to access the medium so it sends a PS-Poll without competition.

To summarize, TIM notifies a group of STAs of the presence of pending frames in the buffer of the AP, and afterwards RAW regulates the prompting of the concerned STAs in addition to the data frames sent upstream.

5.2. Individual TWT

Individual TWT is a mechanism that relies on the notion of appointment between the STAs and the AP. Note that the idea of using appointments in Wi-Fi has been already introduced before but it was restricted to the case of downstream traffic (see S-APSD in Section 3.2). The Individual TWT mechanism was originally brought in the 802.11ah amendment (IEEE WG 802.11ah (2016), Section 10.43), but it has been later refined by the 802.11ax amendment (IEEE WG 802.11ax (2021), Sections 10.47, 26.8.2 and 26.8.3).

Individual TWT establishes a series of appointments between a STA and the AP for frame transmission, in an implicitly or explicitly way. The STA and the AP agree on a TWT SP during which the pending frames from each side are exchanged. The duration of TWT SPs is defined by the parameter “Minimum TWT Wake Duration”. Actually, TWT SPs may end earlier if the AP sends a frame with the EOSP flag raised and the STA has no frame to send upstream. If the Individual TWT is set up in implicit mode, the TWT SPs start at regular intervals of time. Otherwise,

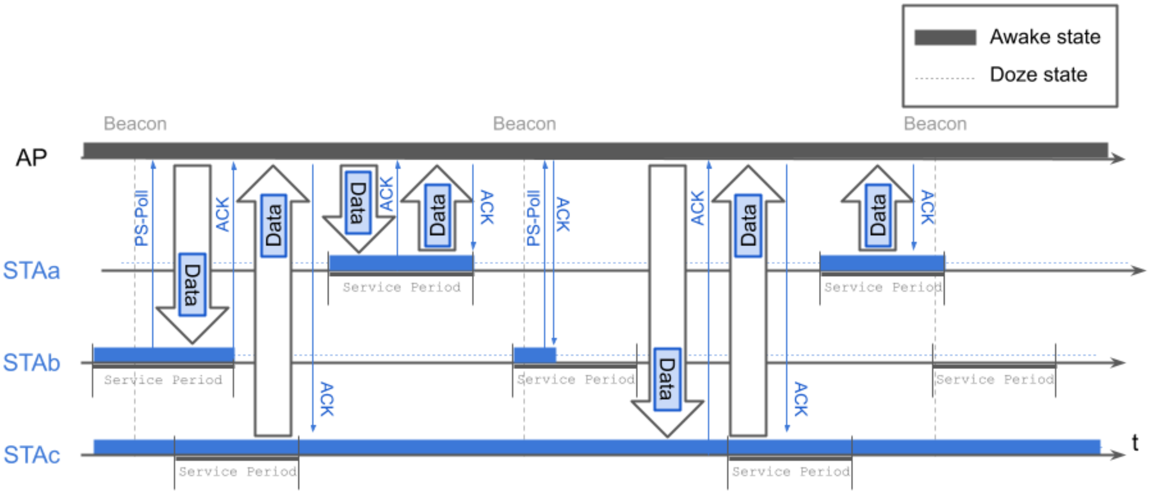

Figure 8: Example of frame exchanges when each STA agreed on an Individual TWT.

if the Individual TWT is running in explicit mode, the AP must notify the STA of the schedule of the next TWT SP during the current TWT SP. Note that these notions already exist with S-APSD. Therefore, for more details, we refer the reader to Section 3.2.

However, Individual TWT goes beyond S-APSD. First, Individual TWT enables the STAs and the AP to negotiate on their TWT SPs based on their respective objectives. Indeed, each STA has its own requirements (in terms of energy saving and delays) and the AP must balance the demands of every STA to ensure fairness and avoid contention. The negotiation uses special MAC management frames as detailed in Table 10-32a of IEEE WG 802.11ax (2021). These negotiation frames carry an entire set of parameters for Individual TWT. Each STA may have up to 8 different TWT agreements with the AP, each of them with their own set of parameters. With multiple TWT agreements, a STA can accommodate more efficiently different patterns of traffic (*e.g.*, flows with different periodicity). Both STA and AP can tear down a previously established TWT agreement at any time if it does not satisfy their objectives any longer.

Second, Individual TWT introduces a parameter called *flow type*. If this latter is set to 1, the AP can send frames downstream as soon as a SP starts, without checking whether the STA is actually in awake state. Otherwise, if the flow type parameter is set to 0, the AP waits for the prompt of the STA before sending pending frames downstream, even when an SP is currently ongoing. In this second case, the STA may skip some SP since the AP will not send frames downstream without being prompted, but this causes additional signaling traffic.

Unlike the mechanisms discussed so far with the exception of RAW, the use of Individual TWT is not restricted to energy saving. In particular, it also aims at reducing contention in the WLAN. Indeed, the AP may enforce an Individual TWT on a STA in order to limit the time a STA occupies the medium to send frames upstream. In this case, the AP entirely skips the negotiation process with the concerned STA. In addition, Individual TWT can be applied to any STAs, whether they are in active mode or in PS mode. For a STA in active mode, a TWT agreement postpones only the upstream traffic, meaning that it does not induce additional delays on the downstream traffic. By enforcing a TWT agreement on STAs in active mode, the AP can regulate the contention without downgrading their performance in downstream traffic.

Figure 8 presents different ways of using Individual TWT. We consider three STAs that each has previously agreed with the AP on a different setting for its Individual TWT. First, the frequency of the SPs differ from one STA to another. STAa is running in PS mode and has negotiated an Individual TWT with the flow type parameter set to 1. STAb is also in PS mode, but has negotiated an Individual TWT with the flow type parameter set to 0. Contrary to STAa and STAb, STAc is in active mode: it never switches to doze state, and as such it does not need the flow type parameter.

Because STAa is in PS mode, the AP sends frames downstream only during the SPs. Thus, STAa can switch to doze state outside its SPs. The flow type parameter set to 1 indicates that STAa will not prompt the AP. During the first SP, the AP has a pending frame for STAa and it sends this frame as soon as the SP starts. Then STAa can proceed to send upstream a data frame. During the second SP, the AP does not have pending frame for STAa in its buffer. So STAa waits for the Probe Delay to expire and then proceeds to send upstream its data frame. Note that STAa has to wake up at every SP in case the AP has pending frames in its buffer: the AP may send them right away at the beginning of the SP.

Contrary to STAa, STAb has a flow type parameter set to 0. During its first SP, the AP has a pending frame for STAb in its buffer, but it does not send it immediately. STAb prompts the AP by sending a PS-Poll and retrieves its pending frame. During the second SP, STAb prompts again the AP but this time the AP returns an ACK, meaning that there is no pending frame for STAb in its buffer. Since STAb has also no pending frame to send upstream, it returns to doze state before the end of the SP. During the third SP, STAb does not wake up at all. STAb can stay on doze state during this SP because with the flow type parameter set to 0 the AP does not send data frames downstream unless it was explicitly prompted.

In the case of STAc, which is in active mode, the TWT restricts only the upstream traffic. When STAc has frames to send upstream, it waits for its SP to do so because STAc is not allowed to emit outside its SP. However, the AP can send frames downstream to STAc at any time because STAc does not doze and is always ready to receive frames.

5.3. Trigger-enabled TWT

Trigger-enabled TWT is derived from Individual TWT in order to leverage OFDMA transmissions (IEEE WG 802.11ax (2021), Section 26.8.2) to make simultaneous transmissions. Recall that we introduced the bare minimum of OFDMA for the understanding of this paper in Section 2.3. During the Trigger-enabled TWT SPs, the AP schedules upstream OFDMA transmissions for the TWT-requesting STAs as well as downstream OFDMA transmissions for data frames pending at the AP. TWT-requesting STAs can use upstream OFDMA transmissions either to prompt the AP (with the goal to receive in return frames pending at the AP buffer) or to send data frames upstream.

With this mechanism, TWT-requesting STAs are not authorized to emit except inside upstream OFDMA transmissions. In return, the AP guarantees the allocation of RUs in the upstream OFDMA transmission of an TWT SP to TWT-requesting STAs.

Figure 9 depicts how STAs exchange their frames with a Trigger-enabled TWT mechanism. In this example, we assume that STAa, STAb and STAc have agreed beforehand with the AP of using the Trigger-enable TWT and that STAa, STAb, and STAc have 3, 2, and 0 data frames, respectively, to send upstream. On the other hand, the AP is assumed to have 8, 2 and 4 data frames to send to STAa, STAb, and STAc, respectively. Although the STAs have individually negotiated with the AP, the AP has synchronized their TWT SPs to let them start at the same time. During the TWT SP common to the three STAs, the AP first sends an OFDMA trigger frame to notify the STAs of their allocated RUs. Then STAs can use the upstream OFDMA transmission to send their pending data frames (if any). Note that the STAs have been each allocated four RUs. For instance, STAa uses its first RU to send an PS-Poll and the three remaining RUs to send data frames upstream. STAb proceeds in the same way, except that it has only two data frames to send upstream so that a RU is left unused. Lastly, in the case of STAc, it uses only one RU to send a PS-Poll, leaving the three other RUs void. After the upstream OFDMA transmission, the AP schedules a downstream OFDMA transmission which includes all the frames for STAa, STAb, and STAc pending in its buffer. Note that there are no RU left unused in the downstream OFDMA transmission. Indeed, the AP knows exactly the number of data frames that need to be sent and thus can allocate the right number of RUs for each STA.

Overall, the objective of Trigger-enabled TWT is to gather the emissions of all STAs to leverage OFDMA transmissions. In this sense, Trigger-enabled TWT is interesting only if there are enough STAs able to use OFDMA.

5.4. Broadcast TWT

Broadcast TWT (IEEE WG 802.11ax (2021), Section 26.8.3) is another mechanism originating from Individual TWT, which like Trigger-enabled TWT, leverages OFDMA transmissions. However, unlike Trigger-enabled TWT, Broadcast TWT uses Random Access to let STAs send their frames inside the upstream OFDMA transmissions. In other words, the AP does not allocate specific RUs for TWT-requesting STAs, but instead, STAs decide randomly on one or more RUs to send their frames upstream.

The AP announces the Broadcast TWT SPs using a TWT element transported by its beacons. Through TWT elements, the AP notifies the STAs of a set of parameters for Broadcast TWT. In particular, each TWT element indicates the starting time of a Broadcast TWT SP as well as its minimum duration. If the AP decides on multiple Broadcast TWT SPs in between its beacons, it can simply announce multiple TWT elements in its beacons. For each Broadcast TWT SP, the AP must schedule at least one upstream OFDMA transmission with Random Access slots. From the

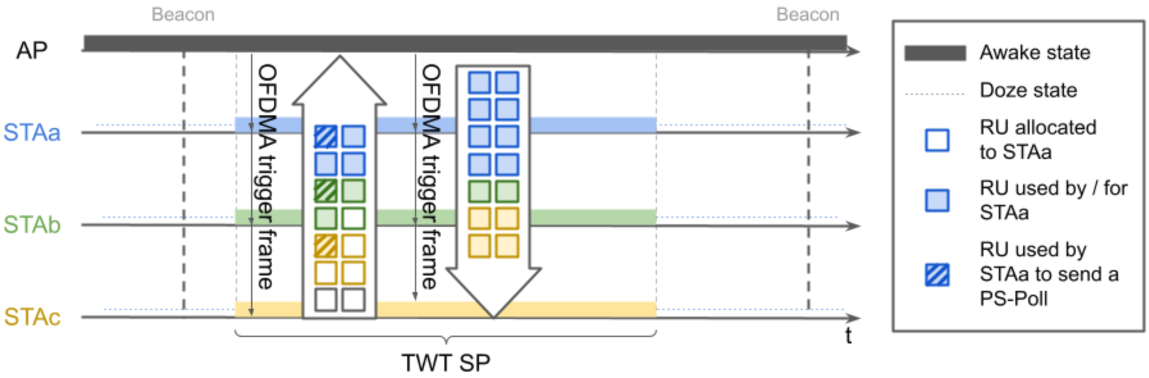

Figure 9: Example of frame exchange during a Trigger-enabled TWT SP. The colors are adapted for each STA.

standpoint of STAs, they just need to wake up in time for receiving the beacons in order to discover the schedule of the Broadcast TWT SPs. Then, for upstream traffic, each STA willing to send frames randomly chooses which Broadcast TWT SP to use, and inside this Broadcast TWT SP, randomly chooses a RU in Random Access. Note that in Random Access within an upstream OFDMA transmission, a STA can only use a single RU. For the downstream traffic, since STAs are not necessarily awake during every Broadcast TWT SP, they must prompt the AP, using a Random Access RU, to receive their frames pending at the AP.

Figure 10 illustrates an example of Broadcast TWT SPs. We assume that STAa, STAb, STAc, and STAd have all agreed to use Broadcast TWT. STAa, STAb, and STAd have all stringent requirements on delay and they do not skip any Broadcast TWT SP, whereas STAc has low requirements so it skips every other Broadcast TWT SP. We also assume that during the first Broadcast TWT SP, none of the STAs has frames to send upstream and that only STAd has a frame, which fits into one RU, to send upstream during the second Broadcast TWT SP. Moreover, we assume that the AP has many pending frames in its buffer for each of the four STAs. During the first Broadcast TWT SP, the AP schedules an upstream OFDMA transmission (depicted by the first upwards wide and white arrow) where every RU is accessible by Random Access. In other words, no RU is allocated to a given STA. All four STAs wake up and attempt to access one RU to send upstream a PS-Poll. Unfortunately, the PS-Polls of STAa and STAb overlap, making their reception unintelligible at the AP. Upon completion of the upstream OFDMA transmission, the AP schedules, within the same Broadcast TWT SP, a downstream OFDMA transmission (depicted by the first downwards white and wide arrow) to send pending frames to the STAs whose PS-Polls were successfully received. Therefore, only STAc and STAd receive their pending frames in the first Broadcast TWT SP. We now turn to the second Broadcast TWT SP. STAa, STAb, and STAd wake up and try to access the Random Access RUs in the upstream OFDMA transmission. STAa and STAb are trying to send a PS-Poll (since their previous ones were not received), whereas STAd has a data frame to send upstream. This time, no collision occurs. In the following downstream OFDMA transmission, the AP sends the pending frames for STAa and STAb.

Interestingly, Broadcast TWT can be coupled with a TIM bitmap and their combination is referred to as *Opportunistic Power Save* (OPS) (see IEEE WG 802.11ax (2021), Section 26.14.3). At the beginning of a Broadcast TWT SP, the AP sends a TIM bitmap which allows or forbids a given STA to switch to doze state during this SP. In the context of OPS, the TIM bitmap does not exactly indicate the presence of pending frames in the buffer of the AP, but rather reflects the intention of the AP. If the AP plans to send data frames downstream to a given STA, or to allocate RUs to this STA in an upstream OFDMA transmission, then the AP sets the bit corresponding to the STA to 1. This prevents the STA from switching to doze state before the end of the SP, which could lead the STA to miss data frames sent downstream. On the contrary, if the bit of the TIM bitmap is set to 0, the STA is allowed to skip the upcoming Broadcast TWT SP, but the STA is also allowed to stay awake during this Broadcast TWT SP and try to send its data frames upstream. The advantage brought by OPS over a regular Broadcast TWT is the absence of prompts. Since the AP can demand STAs to be awake during a specified SP, the STA does not need to prompt the AP since the AP knows

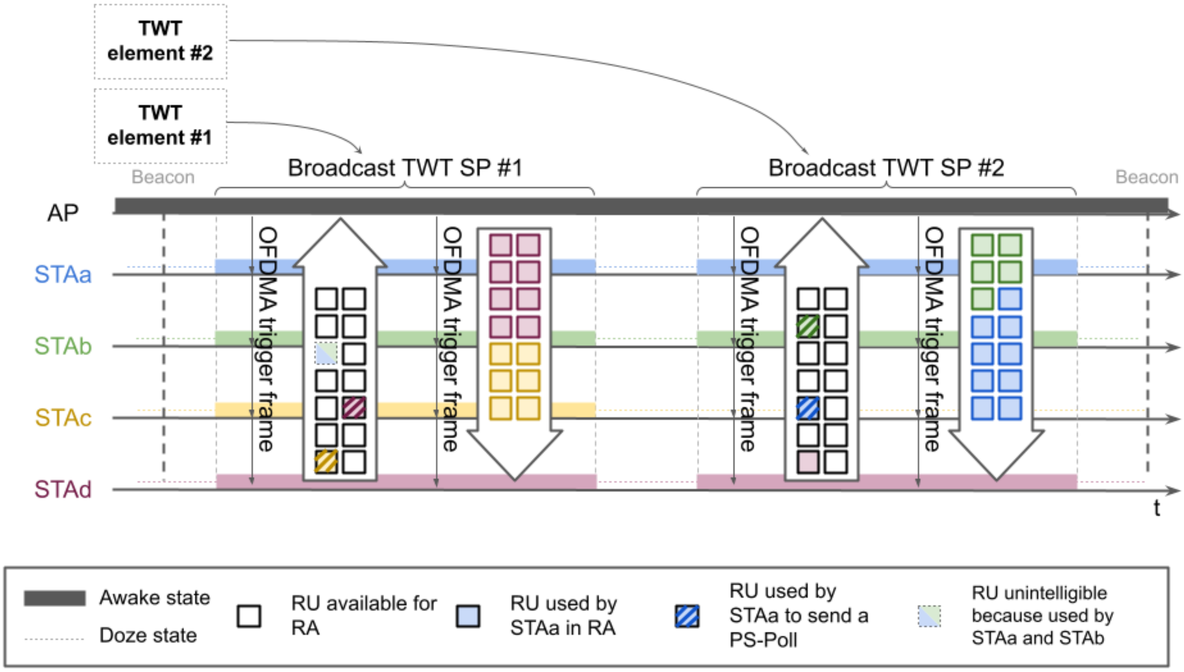

Figure 10: Example of OFDMA transmissions inside a Broadcast TWT SP. RA stands for Random Access.

that the STA is awake. On the downside, it forces the STAs to regularly wake up to listen to the TIM bitmap, namely at the beginning of each Broadcast TWT SP.

Overall, Broadcast TWT is more flexible than Trigger-enabled TWT. With Broadcast TWT, the AP does not need to forecast how many RUs each STA will need. Furthermore, the AP can schedule multiple Broadcast TWT SPs when needed. On the downside, overlapping frames may occur in Broadcast TWT since it relies on Random Access.

In the future, the 802.11be amendment will extend the Broadcast TWT with the *restricted TWT* or rTWT. The goal is to bound delays for a set of STAs that have strong real-time requirements. As said previously, the AP can schedule multiple Broadcast TWT SPs in one beacon with different TWT elements. With rTWT, some of these Broadcast TWT SPs will be reserved for certain STAs: these SPs will be referred to as rTWT SP. The STAs accessing rTWT SPs benefit from a lower contention and are more likely to send their frames upstream without collision. Even though there is no strict guarantee, the delays experienced by STAs in rTWT SPs are expected to be lower on average than with a regular Broadcast TWT.

6. Further considerations

Having described the existing energy-saving mechanisms in Wi-Fi, we now discuss the different aspects to be taken into account when facing the choice of an energy-saving mechanism. In Section 6.1, we establish a set of metrics to facilitate an objective comparison between the energy-saving mechanisms. Section 6.2 delivers a chronology of the IEEE 802.11 amendments and the associated energy-saving mechanisms. Finally, in Section 6.3, we discuss the existing certifications delivered by the Wi-Fi Alliance, which reveals the choices of chipset manufacturers for energy saving.

6.1. Metrics of interest

All energy-saving mechanisms operating at the MAC layer may reduce the network performance in order to save energy. However, the extent to which their performance as well as their energy consumption will be affected may widely differ across the energy-saving mechanisms and remains an open question. For example, an energy-saving mechanism may attempt to keep a high throughput while making the saving of energy a subsidiary goal. On the contrary, strongly reducing the energy consumption on a device may also significantly reduce its throughput. Overall, there is a complete range of trade-offs between these two latter extremes, and although none of them is likely to be better than the others, it is useful to be able to compare energy-saving mechanisms with each other.

To begin with, we distinguish two classes of metrics: quantitative versus qualitative metrics. Quantitative metrics can be accurately measured and compared. On the other hand, qualitative metrics refer to non-measurable performance. For instance, they may relate to ease-of-use of a mechanism or to its degree of security.

6.1.1. Quantitative metrics

Quantitative metrics are insightful to assess the quality of the trade-off between network performance and energy savings. Clearly, a single metric is not enough to reflect every aspect of the trade-off since network performance and energy savings are often (negatively) correlated. Thus, we suggest a set of relevant metrics that one may use to evaluate and compare energy-saving mechanisms.

- The relevant *network performance* are the throughput and the delay of the frames of the STAs, both upstream and downstream. We believe that other customary network metrics (*e.g.*, the jitter) are of secondary importance when energy saving is of utmost importance. Note that for the sake of the understanding of a mechanism, the *signaling load* and the *frame retransmission rate* may be helpful. The signaling load is the throughput used to send MAC management frames. The frame retransmission rate is the number of frames retransmissions as compared with the number of frames transmissions. Although not directly of interest to the end user, measuring these two latter metrics may help identify the flaws of a given energy-saving mechanism.
- In terms of energy saving, the most relevant metric is the time a STA spends in doze state, referred to as its *doze time*. Indeed, on the MAC layer, the ultimate means to save energy is to reduce the time during which a STA is needlessly being in awake state. This includes the times associated to collisions, the backoff, and the transmission of other devices. For example, a STA with a perfect energy-saving mechanism would be in awake state only to send or receive its frames, and would stay in doze state the rest of the time.

Energy efficiency is another energy metric worth of interest. It is generally defined as the average amount of energy consumed to successfully send one bit of information. This is obtained by dividing the average electric power by the average throughput. Therefore, the energy efficiency of a STA can be improved by reducing its energy consumption, but also by increasing its throughput. It is interesting to note that the throughput of a STA can be increased by orders of magnitude more than the energy consumption can be reduced. Indeed, according to Xiao, Cui, Savolainen, Siekkinen, Wang, Yang, Ylä-Jääski and Tarkoma (2014), the energy consumption of a STA in doze state is 0.042 W against 1.258 W in transmission state. The considered chip was complying to the IEEE 802.11n standard, which means its data rate could go as high as 600 Mbps. With these figures, it follows that a STA that constantly transmits frames at its maximum data rate will attain an energy efficiency approximately equal to 2.1 nJ/b. On the other hand, a STA that sends one frame per hour (corresponding to a throughput of about 3 bps) and dozes the rest of the time leads to an energy efficiency roughly equal to 14 mJ/b. This can lead STAs with very high energy consumption and high throughput to prevail, in terms of energy efficiency, over STAs with lower energy consumption but much smaller throughput. One needs to keep in mind this possible bias when studying MAC energy-saving mechanisms, and hence include the average throughput of STAs in their analysis.

6.1.2. Qualitative metrics

Besides the quantitative metrics, other aspects are key to assessing the efficiency of an energy-saving mechanism. First, *versatility* that reflects the ability of the mechanism to cover a wide range of trade-offs between network performance and energy saving. For example, a versatile energy-saving mechanism may accommodate simultaneously a laptop streaming in high video quality and a sensor sending one frame per hour. One advantage of a versatile mechanism is that a single energy-saving mechanism may suit the needs of a whole WLAN, thereby easing its maintenance.

Second, *self-adaptability* is the ability of a device to automatically adapt itself to the (varying) operating conditions of the WLAN. For example, a self-adaptable energy-saving mechanism lets a device return to doze state whenever it detects a high contention on the medium. Similarly, if the battery of the STA is running low, the self-adaptable energy-saving mechanism may space out the time periods during which the STA is in awake state. A self-adaptable mechanism can closely fit the needs of the devices without resorting to human intervention.

Third, *negotiability* may also be of some importance. In general, STAs and AP may have different objectives for energy saving. STAs may want to pursue a certain trade-off between their network performance and their energy

	Versatility	Self-adaptability	Negotiability	Safety	Backward compatibility
PS-Poll	high	high	low	no	high
U-APSD	high	high	low	yes	high
S-APSD	high	high/mid[a]	low	yes	high
TIM	low	low	low	no	high
TIM segmentation	low	low	low	no	low
RAW	low	low	low	yes	low
Individual TWT	high	high/mid[a]	high	yes	high
Trigger-enabled TWT	high	mid	high	yes	mid
Broadcast TWT	high	high	high	yes	mid

Table 2
Summary of the advantages and drawbacks of the different energy-saving mechanisms in Wi-Fi.
[a] Implicit mode on the left, explicit mode on the right.

consumption in favor of energy consumption. On the other hand, the AP's goal may be to ensure as much as possible a low contention among STAs and fair access to the medium. Energy-saving mechanisms that allow a negotiation before operating are naturally more capable to meet, at least partially, both parties' expectations.

Fourth, the *safety* relates both to the *security* and to the *robustness* of the energy-saving mechanisms. The security is the ability of an energy-saving mechanism to resist against potential malicious attacks. The robustness reflects the sensitivity of an energy-saving mechanism to unfavorable operating conditions (*e.g.*, poor radio channel).

Lastly, *backward compatibility* is a key component of Wi-Fi commercial success, which allows devices from various generations to coexist in WLANs. It is worth noting that the mechanisms defined in the 802.11ah amendment and its sub-1GHz band are incompatible with the 2.4 and 5GHz radio bands. Also, energy-saving mechanisms that rely on OFDMA are only partially compatible with those making use of the ordinary CSMA/CA.

6.1.3. Comparing the mechanisms

The values of quantitative metrics strongly depend on the scenario chosen to evaluate them. Unfortunately, to the best of our knowledge, no study compares the existing energy-saving mechanisms using the same metrics and the same conditions. This lack impedes any further solid discussion in terms of quantitative metrics.

On the other hand, for the qualitative metrics, Table 2 reports the pros and cons of each energy-saving mechanism. We detail below how we classify (low/mid/high or yes/no) each MAC energy-saving mechanism with regards to the five criteria listed above:

- Versatility: This is set to low for energy-saving mechanisms that are strongly tied to the pace at which beacons are broadcast (or any other fixed-frequency events).
- Self-adaptability: An energy-saving mechanism is classified as not self-adaptable if it requires a human intervention to change its parameters. On the other hand, if the mechanism can adapt the parameters by itself without the exchange of any frames, then we classify it as high in terms of self-adaptability. If the mechanism needs to explicitly exchange frames to change its parameters, then we classify it as mid.
- Negotiability: This is set to low if the parametrization of a mechanism is decided only by the STA, or alternately only by the AP.
- Safety: An energy-saving mechanism is classified as unsafe whenever a security or robustness failure has been reported.
- Backward-compatibility: An energy-saving mechanism is reported as low (resp. mid) in backward-compatibility if the sheer presence of a legacy STA in the WLAN precludes the use of (resp. significantly reduces the interest of using) this mechanism for every other associated STAs.

The PS-Poll and the U-APSD mechanisms restrict the downstream traffic by letting the STAs prompt the AP. Both mechanisms are entirely distributed: the STAs decide and the AP must follow, which makes these mechanisms low

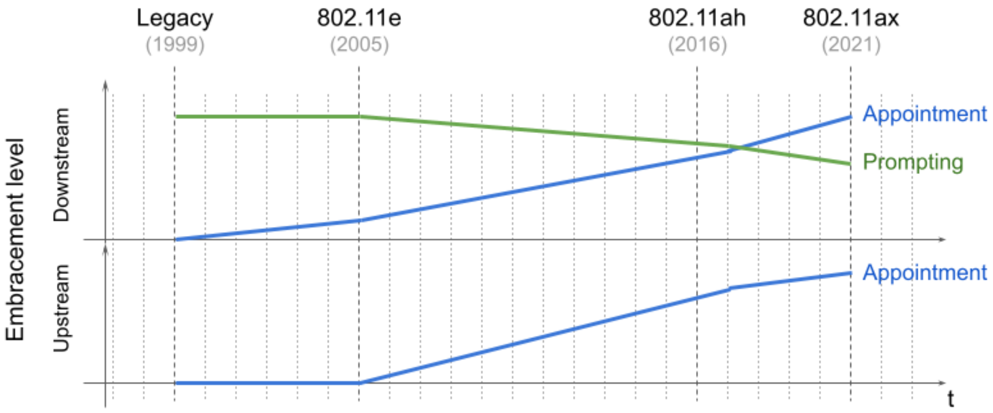

Figure 11: Embracement by IEEE Working Group 802.11 of the different types of MAC energy-saving mechanism.

in terms of negotiability. Therefore, the PS-Poll and U-APSD mechanisms perform well in versatility and in self-adaptability. Note that the PS-Poll frames are not authenticated and are thus prone to *Deny of Service* (DoS) attacks (and thus unsafe), which is not the case of APSD trigger frames (Souilah, Baadache and Bouallouche-Medjkoune (2014)).

The S-APSD and the Individual TWT mechanisms both rely on appointments and have a similar operating mode, although they do not restrict the same type of traffic. They allow a large range of frequency for appointments, making them highly versatile. As for the self-adaptability, they perform relatively well when configured in implicit mode: a STA may quickly return to doze state if the medium is busy and it will wait for the next appointment to exchange pending frames. However, none of these mechanisms automatically adapt the frequency of the appointments (*e.g.*, when the battery is running low). Note that Individual TWT brings a strong negotiation phase, which was not included in S-APSD.

The TIM mechanisms – segmented or not – rely on the beacons to transmit information to the STAs. Since the beacons are broadcast at fixed intervals of time, the TIM mechanisms are not versatile or self-adaptable. Note also that the TIM bitmaps are sensitive to poor radio conditions: the TIM bitmap is never retransmitted, even when some of its bits are received in error (Lee, Go and Lee (2020)), which makes these mechanisms unsafe.

The RAW mechanism behaves closely to the TIM mechanism as it also uses the beacons to transmit operating information to the STAs. Additionally, RAW is not backward compatible and cannot be used in the 2.4 and 5 GHz radio bands.

Trigger-enabled TWT and Broadcast TWT mechanisms have similar properties. They both rely on prompts to send frames making them versatile and self-adaptable. More precisely, we consider Trigger-enabled TWT a bit less versatile and self-adaptable than Broadcast TWT because of the induced cost at the WLAN scale. Indeed, unlike Broadcast TWT, Trigger-enabled TWT reserves RUs to STAs, which will be left unused if the STA decides to skip the SP. As for backward compatibility, Trigger-enabled TWT and Broadcast TWT can be used alongside older devices, but to be truly efficient they need multiple STAs in the WLAN to use the mechanism.

6.2. Chronology of the energy-saving mechanisms

In 1999, upon the very first steps of Wi-Fi, the legacy standard provided to the STAs the possibility to switch to PS mode, but only with the TIM and PS-Poll mechanisms. At this point, the sole fact that energy saving was considered was of a great importance. However, the mechanisms could not be optimized because nobody knew the exact needs of the STAs and the WLANs since Wi-Fi was not yet widely adopted. As a result, the TIM and the PS-Poll mechanisms were later revealed as flawed, although functional (see Souilah et al. (2014) and Lee et al. (2020)).

In 2005, the release of the 802.11e amendment brought multiple improvements in Wi-Fi, most of them pertaining to the user experience. Regarding energy saving, the 802.11e amendment introduced the U-APSD, which addresses a number of flaws of PS-Poll, as well as S-ASPD, which pioneered the notion of appointment in Wi-Fi (but did not draw as much attention as U-APSD).

In 2016, the release of the 802.11ah amendment which was dedicated to the IoT, marked a significant step in the development of energy-saving mechanisms. The 802.11ah amendment developed the notion of appointments with TWT, which stays the most "mature" energy-saving mechanism of IEEE 802.11 standards to this date. Indeed, it is the first and only energy-saving mechanism to acknowledge that STAs and AP may have different needs. In addition, TWT uses flexible appointments that allow a good versatility and self-adaptability while having lighter signaling overhead than a prompting mechanism. In addition, the schedule of the upstream traffic can be restricted with RAW, mostly with the aim of reducing the contention among STAs. However, the 802.11ah amendment was mostly a commercial failure probably because of several reasons including the use of a new radio band (sub 1 GHz) and the complexity of the amendment among others.

Lastly, in 2021, the release by IEEE of the 802.11ax amendment built on the previous attempts of the 802.11ah amendment and brought TWT to the 2.4 and 5 GHz radio bands. With TWT, the AP can restrict any upstream traffic to attempt to lessen the contention, and can also restrict the downstream traffic to help STAs that expressly wish to limit their energy consumption. Besides, with the introduction of OFDMA, the 802.11ax amendment opens the way towards new energy-saving mechanisms strongly leveraging OFDMA.

In the future, the next main IEEE 802.11 release will be the 802.11be amendment, which is already marketed as Wi-Fi 7. As it currently stands, Wi-Fi cannot give any deterministic guarantee in the delivery frame delays, and thus is mostly unsuited for time-sensitive applications. One objective of the 802.11be amendment is to enable such applications to run on Wi-Fi by reducing delivery delays in worst-case scenarios. For more information on what may cover the 802.11be amendment, we refer the interested reader to Khorov, Levitsky and Akyildiz (2020). With regards to the MAC energy-saving mechanisms, the 802.11be amendment does not plan on introducing novel mechanisms but rather on adapting the existing Broadcast TWT mechanism to time-sensitive applications: it will be then referred to as rTWT.

Overall, the different approaches of energy-saving mechanisms can be analyzed as follows. When operating on the downstream traffic, energy-saving mechanisms mainly help STAs reduce their consumption. On the downside, they may have harmful side effects at the WLAN scale (*e.g.,* a sudden peak of contention after a TIM beacon). On the other hand, energy-saving mechanisms dedicated to upstream traffic mostly contribute in reducing the contention and (wasteful) retransmissions, and, saving energy as a by-product. At last, combining restrictions both on the downstream and upstream traffic enables energy-saving mechanisms to find new trade-offs wherein STAs may reduce their energy consumption while preserving a high level of performance for the WLAN.

Figure 11 illustrates the relative adoption of each type of mechanisms by the IEEE standards throughout the years. Prompting the AP is probably the most straight-forward solution to restrict the downstream traffic. As such, energy-saving mechanisms relying on prompting the AP were the most widely adopted in the early years of the Wi-Fi. However, in the last years, IEEE has more and more promoted the appointments over the prompts to restrict the downstream traffic. This is due to the fact that a prompt frame may collide with a data frame or with another prompt frame. As a result, prompts put in contention the different downstream traffics. Figure 11 also shows that restricting the upstream traffic has gradually become an important issue, just as much as for the downstream traffic. This re-balancing may result from the increasing contention that WLANs may undergo.

6.3. Wi-Fi Alliance's certifications

The commercial market for Wi-Fi is hard to apprehend; relatively few technical details are disclosed about the available products. However, studying the Wi-Fi Alliance and its protocols can help appreciate which mechanisms have penetrated the market and those which have not.

The Wi-Fi Alliance is the industrial consortium that certifies a product as Wi-Fi compatible. To do so, the Wi-Fi Alliance defines a subset of features (*e.g.*, the beacons, the backoff procedure) from the IEEE standards. The features included in the tests are discussed among the stakeholders of the Wi-Fi Alliance, which encompass the major companies of chipset manufacturers (*e.g.*, NXP Semiconductors, Intel), device assemblers (*e.g.*, Huawei, Apple) and a few internet operators. When a feature is selected, a test protocol is also negotiated between the different parties. Any product that wishes to be labelled "Wi-Fi certified" must pass every test that were negotiated. This way, the Wi-Fi Alliance can certify of the interoperability between two "Wi-Fi certified" devices.

So far, the Wi-Fi Alliance issued three certification programs that include at least one MAC energy-saving mechanism as a feature: *WMM-Power Save* (Wi-Fi MultiMedia), *Wi-Fi 6* and *Wi-Fi HaLow*. The WMM-Power Save

Certification program	Year	AP features	STA features
WMM-Power Save	2005	PS-Poll, TIM, U-APSD	PS-Poll, TIM
Wi-Fi HaLow	2021	PS-Poll, TIM segmentation, RAW	PS-Poll, TIM segmentation, RAW
Wi-Fi 6	2021	Individual TWT, Broadcast TWT*	Individual TWT*

Table 3
Overview of the Wi-Fi Alliance certifications featuring at least one MAC energy-saving mechanism.
*Optional feature.

Categories	Number of certified devices	WMM-Power Save	Wi-Fi HaLow	Wi-Fi 6
AP (home or small offices)	4,101	16%	0%	10%
Laptop	1,103	37%	0%	26%
Smartphone	19,085	68%	0%	4%
Printer	4,015	0%	0%	0%
Smart home	601	0%	0%	0%
TV	10,628	7%	0%	1%

Table 4
Percentage of "Wi-Fi certified" devices that passed a certification program featuring a MAC energy-saving mechanism. Data extracted in November 2022.

certification program is optional: a device may not implement its features and still be labeled "Wi-Fi certified". The WMM-Power Save certification features the following energy-saving mechanisms: PS-Poll, TIM, and U-APSD for the APs, and PS-Poll and TIM for the STAs. Wi-Fi 6 (which follows the 802.11ax amendment) is the most recent mandatory certification program, and is still a work in progress as we were writing this article. In its first release, the Wi-Fi 6 certification program already includes the individual TWT mechanism, mandatory for the APs and optional for the STAs. In January 2022, the second release added the Broadcast TWT mechanism as an optional feature for the APs. Lastly, the Wi-Fi HaLow certification program follows the 802.11ah amendment, which was designed to fit the peculiar needs of IoT such as energy saving requirement. Wi-Fi HaLow includes the following MAC energy-saving mechanisms: PS-Poll, TIM segmentation and RAW, which are mandatory for both the AP and STAs. Table 3 summarizes the year of release and the associated MAC energy-saving mechanisms for each of these Wi-Fi Alliance certifications.

The website of the Wi-Fi Alliance enables one to search for "Wi-Fi certified" products and to see for any product which certification programs it passed (see Wi-Fi Alliance). In fact, the website lists every product that has ever been certified, including products from the very first generation. The website further divides products by categories and subcategories that correspond to the general purpose of the devices (*e.g.*, routers, computers, televisions). We collected these data for several subcategories and computed the percentage of devices that passed the three certification programs described in the former paragraph: WMM-Power Save, Wi-Fi 6, Wi-Fi HaLow. The corresponding results are presented in Table 4. With the exception of three APs, the table shows that no device passed the Wi-Fi HaLow certification program. For WMM-Power Save, 68% of smartphones passed this certification program. We believe that this owes to the strong requirement that smartphones have for their battery lifetime. However, keep in mind that a smartphone that passed the WMM-Power Save certification program may not be able to activate its MAC energy-saving mechanisms if the AP, to which the smartphone is associated, does not implement them. For Wi-Fi 6, laptops seem to have embraced this new certification program (26%) faster than smartphones (4%). However, we believe that this may owe to the very high throughput that Wi-Fi 6 may offer rather than the introduction of the Individual TWT mechanism.

7. Related work

Wi-Fi is defined by the IEEE 802.11 standard and its subsequent amendments, which describe in detail every existing energy-saving mechanism. The standards are the primary sources of information for anything regarding Wi-Fi, and Table 1 indicates the MAC energy-saving mechanisms discussed in this paper and reports their corresponding IEEE amendments. However, the reading of the standards is complex, often hindering the understanding of energy-saving mechanisms. Furthermore, the goal of each mechanism (beyond energy saving) is not clearly stated in the standards and their parametrization is viewed by IEEE as out of the scope of the standards. In other words, the standards offer possibilities but do not help in choosing among them.

In response to the length and complexity of the standards, a couple of papers attempted to ease the understanding of one amendment. In 2021, Tian et al. (2021) focused on the 802.11ah amendment, which adapts Wi-Fi to the Internet of Things, and a large part of the paper is dedicated to the description of energy-saving mechanisms introduced in the 802.11ah amendment. Similarly, in 2018, Khorov, Kiryanov, Lyakhov and Bianchi (2018) presented the 802.11ax amendment in detail. The paper is very comprehensive and describes energy-saving mechanisms such as TWT. The scope of each of these two papers is restricted to a couple of energy-saving mechanisms. Hence, they do not put into perspective all the Wi-Fi energy-saving mechanisms that have been designed and implemented over the successive amendments since 1999.

We end this section by discussing some of the main issues that have attracted attention in the past as well as those being active areas of research for each of the MAC energy-saving mechanisms in Wi-Fi.

Legacy Power Save Mode: PS-Poll and TIM

Back in 1999, when the legacy Power Save Mode was introduced, most studies were focusing on improving the network efficiency of the MAC layer rather than enhancing the energy saving brought by the PS mode (*e.g.*, Woesner, Ebert, Schläger and Wolisz (1998))[2]. In 2008, in Anastasi, Conti, Gregori and Passarella (2008), the authors present a mathematical model of the legacy Power Save Mode in scenarios with bursty downstream traffic. Their primary goal is to forecast the energy spent by the STAs to receive their pending frames. Then, they use their model to propose a cross-layer energy-saving mechanism that is able to detect and adapt itself to traffic bursts, thereby further improving the amount of energy saved by STAs. In their work, He and Yuan (2009) identify several other factors that limit the doze time of the STAs using the legacy Power Save Mode. For instance, they point out that the “one-poll-one-answer” policy prevents a STA from retrieving all its pending frames at once (when it has more than one). Another example is that although STAs know when they have pending data frames in the buffer of the AP thanks to TIM, the use of CSMA/CA makes the retrieval times uncertain. The authors design a new MAC energy-saving mechanism that aims at addressing these issues. It is worth noting that these issues have been separately addressed by various IEEE amendments.

APSD

APSD (be it scheduled or unscheduled) was introduced in the IEEE 802.11e amendment, whose primary goal was to provide Quality of Service. Hence, research studies on APSD were mostly concerned with its ability to guarantee the delay requirement for applications such as Voice over IP. For example, in the case of U-APSD, Pérez-Costa, Camps-Mur and Vidal (2007) compares the performance of U-APSD with those of the legacy Power Save Mode in a scenario with both Voice over IP and background traffic. For S-APSD, Lee and Hsieh (2013) proposes a solution to minimize the overlapping of the different SPs in a scenario with departures and arrivals of flows. Lastly, the authors of Pérez-Costa and Camps-Mur (2010) compare the legacy Power Save Mode with both U-APSD and S-APSD. They conclude that S-APSD generally outperforms the other mechanisms both in terms of network throughput and energy consumption. Overall, implementing the Quality of Service on STAs in PS mode is a challenging task since there are many ways to configure the MAC energy-saving mechanism for each one of the four Access Categories defined by IEEE 802.11e. In general, existing studies focus on a single configuration – overlooking all the other possible configurations – and then they evaluate its performance through simulation.

RAW

Configuring the three parameters of RAW (namely, the number of groups and subgroups, the distribution of STAs among the groups, and the duration of the slots associated with each group) is the main challenge when implementing RAW in a WLAN. Overall, the vast majority of studies agree that the RAW groups should be balanced, either in the number of STAs or in their slot duration. Otherwise, this could cause unfairness issues. For example, Dong, Wu, Gao and Zhao (2016) proposes a balancing algorithm that puts STAs from the same geographical region in the same group to avoid hidden node issues. They evaluated the efficiency of their solution through simulation. In Kureev, Bankov, Khorov and Lyakhov (2017), the authors propose an analytical model of RAW and TIM segmentation to ease the coexistence of STAs having low throughput and high energy-saving requirements with STAs having high throughput. In essence, the authors rely on RAW to separate STAs with low throughput from those with high throughput and then

---

[2] We remind that the legacy Power Save Mode corresponds to the MAC energy-saving mechanism combining PS-Poll and TIM, whereas PS mode indicates the possibility for a STA to switch to the doze state.

they carefully select the duration of the RAW slot. Note that these two latter studies focus on a specific scenario and do not formulate a global algorithm that could work in any scenario. In Santi, Tian, Khorov and Famaey (2019), the authors develop an analytical model to compare the efficiency of combining RAW and TIM segmentation with the one obtained with TWT. This leads them to identify the most adequate mechanism depending on the properties of a network scenario.

TWT

TWT has been investigated by researchers prior to the release of the IEEE 802.11ax amendment. These efforts positively contributed to the design of the IEEE amendment, but many of these studies had to make assumptions (sometimes in disagreement with the forthcoming amendment) about the way TWT would implement OFDMA. Overall, most researchers' efforts focused on the combination of TWT with multi-user transmissions, be it OFDMA or *Multi-User Multiple-Input Multiple-Output* (MU-MIMO). For example, Chen and Zhu (2021) and Nurchis and Bellalta (2019) design a mathematical model to optimize the number of parallel transmissions in a WLAN. Note that these studies mainly aim at maximizing the network throughput in dense scenarios with energy-saving being at most a secondary objective. In the case of individual TWT, Yang, Lee and Bahk (2021) propose to cast TWT as a *Time Division Multiple Access* problems, to which they apply classical schedulers, such as max-rate or proportional fairness schedulers. Through simulation, the authors demonstrate that TWT enhances throughput, fairness, and energy consumption compared to a WLAN with no MAC energy-saving mechanism, whichever scheduler is chosen. A strong edge of TWT is that the STAs may stay in doze state during much longer periods compared to the previous MAC energy-saving mechanisms (this can go up to several hours). However, this possibility may deepen the issue of clock drifting that affects any MAC energy-saving mechanism working with appointments. The wake-up radio, as proposed by Stepanova, Bankov, Khorov and Lyakhov (2020), can help address this issue.

8. Conclusions

In this paper, we have described the MAC mechanisms that the IEEE 802.11 Working Group developed over the years (since 1999) for the sake of saving energy. More specifically, our study has covered the MAC energy-saving mechanisms dedicated to downstream traffic, upstream traffic, or both direction. This has led us to describe diverse mechanisms where either STAs prompt the AP, the AP notifies the STAs, or the STAs and the AP make an appointment. In every case, we further illustrated the mechanism's principle through examples. We have also presented a set of metrics to compare the effectiveness of energy-saving mechanisms with each other and discussed how the energy-saving mechanisms of IEEE were adopted by the Wi-Fi Alliance, and hence by chip manufacturers. Our contribution is a complete overview of 20 years of energy saving in WLANs, covering the advantages and limitations of each mechanism, with their relevant parameters.

This work could be extended by describing the efforts made by IEEE on the physical layer of 802.11 (e.g., TXOP Power Save aka micro-sleep, wake-up radio) - but this would represent a separate work. Of course, the content of the paper will also need to be updated with the upcoming amendments such as 802.11be, whose release is expected in a couple of years.

Overall, we believe this paper can serve researchers, practitioners, and chip manufacturers alike in their efforts to better design and configure existing and forthcoming MAC energy-saving mechanisms, thereby enabling 802.11 to contribute to the energy-saving efforts that WLANs (or mankind) must do.

Acknowledgements

This work was funded by the French National Research Agency (ANR) under the project ANR FACTO (ANR-21-CE25-0020).

A. Notation

Table 5 reports the main notation used in this paper.

References

Anastasi, G., Conti, M., Gregori, E., Passarella, A., 2008. 802.11 power-saving mode for mobile computing in Wi-Fi hotspots: Limitations, enhancements and open issues. Wireless Networks 14, 745–768. doi:`10.1007/s11276-006-0010-9`.

Acronym	Definition	Related sections
-	Active mode / PS mode	Section 3
APSD	Automatic Power Save Delivery	Sections 3.1.2 and 3.2
CSMA/CA	Carrier Sense Multiple Access / Collision Avoidance	Section 2.2
DCF	Distributed Coordination Function	Section 2.2
DTIM	Delivery TIM	Section 3.3.2
EOSP	End Of Service Period	Section 3.1.2
-	Listen Interval	Section 2.1
MAC	Medium Access Control	Section 2
OFDMA	Orthogonal Frequency-Division Multiple Access	Section 2.3
NAV	Network Allocation Vector	Section 2.2
-	Page Slicing	Sections 3.3.2, 4 and 5.1
-	Probe Delay	Section 2.2
PS-Poll	Power-Saving Poll	Section 3.1.1
RAW	Restricted Access Window	Section 4
RU	Resource Unit	Section 2.3
RA	Random Access	Section 2.3
SP	Service Period	Sections 3.1.2 and 5.2
TWT	Target Wake Time	Sections 5.2, 5.3 and 5.4
TIM	Traffic Indication Map	Section 3.3.1

Table 5
Definition of the main keywords and their respective acronyms.

Chen, Q., Zhu, Y., 2021. Scheduling Channel Access Based on Target Wake Time Mechanism in 802.11ax WLANs. IEEE Trans. Wirel. Commun. 20, 1529–1543. URL: `https://doi.org/10.1109/TWC.2020.3034173`, doi:10.1109/TWC.2020.3034173.

Dong, M., Wu, Z., Gao, X., Zhao, H., 2016. An Efficient Spatial Group Restricted Access Window Scheme for IEEE 802.11ah Networks, in: 2016 Sixth International Conference on Information Science and Technology (ICIST), pp. 168–173. doi:10.1109/ICIST.2016.7483405.

Gast, M.S., 2013. 802.11 ac: a survival guide: Wi-Fi at gigabit and beyond. " O'Reilly Media, Inc.".

He, Y., Yuan, R., 2009. A Novel Scheduled Power Saving Mechanism for 802.11 Wireless LANs. IEEE Trans. Mob. Comput. 8, 1368–1383. URL: `https://doi.org/10.1109/TMC.2009.53`, doi:10.1109/TMC.2009.53.

IEEE WG 802.11, 1999. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification. Standard 802.11. IEEE WG 802.11.

IEEE WG 802.11ac, 2013. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification ; Enhancements for Very High Throughput for Operation in Bands Below 6 GHz. Standard 802.11ac. IEEE WG 802.11ac.

IEEE WG 802.11ah, 2016. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification ; Sub 1 GHz License Exempt Operation. Standard 802.11ah. IEEE WG 802.11ah.

IEEE WG 802.11ax, 2021. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification ; Enhancements for High-Efficiency WLAN. Standard 802.11ax. IEEE WG 802.11ax.

IEEE WG 802.11e, 2005. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification ; Medium Access Control (MAC) Quality of Service Enhancements. Standard 802.11e. IEEE WG 802.11e.

Khorov, E., Kiryanov, A., Lyakhov, A., Bianchi, G., 2018. A Tutorial on IEEE 802.11ax High Efficiency WLANs. IEEE Communications Surveys & Tutorials PP, 1–1. doi:10.1109/COMST.2018.2871099.

Khorov, E., Levitsky, I., Akyildiz, I.F., 2020. Current Status and Directions of IEEE 802.11be, the Future Wi-Fi 7. IEEE Access 8, 88664–88688. doi:10.1109/ACCESS.2020.2993448.

Kureev, A., Bankov, D., Khorov, E.M., Lyakhov, A.I., 2017. Improving Efficiency of Heterogeneous Wi-Fi Networks with Joint Usage of TIM Segmentation and Restricted Access Window, in: 28th IEEE Annual International Symposium on Personal, Indoor, and Mobile Radio Communications, PIMRC 2017, Montreal, QC, Canada, October 8-13, 2017, IEEE. pp. 1–5. URL: `https://doi.org/10.1109/PIMRC.2017.8292304`, doi:10.1109/PIMRC.2017.8292304.

Lee, I.G., Go, K., Lee, J., 2020. Battery Draining Attack and Defense against Power Saving Wireless LAN Devices. Sensors 20, 2043. doi:10.3390/s20072043.

Lee, T., Hsieh, J., 2013. Low-Complexity Class-Based Scheduling Algorithm for Scheduled Automatic Power-Save Delivery for Wireless LANs. IEEE Trans. Mob. Comput. 12, 571–580. URL: `https://doi.org/10.1109/TMC.2012.114`, doi:10.1109/TMC.2012.114.

Nurchis, M., Bellalta, B., 2019. Target Wake Time: Scheduled Access in IEEE 802.11ax WLANs. IEEE Wireless Communications 26, 142–150. doi:10.1109/MWC.2019.1800163.

Pérez-Costa, X., Camps-Mur, D., 2010. IEEE 802.11E QoS and power saving features overview and analysis of combined performance [Accepted from Open Call]. IEEE Wireless Communications 17, 88–96. doi:10.1109/MWC.2010.5547926.

Pérez-Costa, X., Camps-Mur, D., Vidal, A., 2007. On Distributed Power Saving Mechanisms of Wireless LANs 802.11e U-APSD vs 802.11 Power Save Mode. Comput. Networks 51, 2326–2344. URL: `https://doi.org/10.1016/j.comnet.2007.01.026`, doi:10.1016/j.comnet.

2007.01.026.

Santi, S., Tian, L., Khorov, E., Famaey, J., 2019. Accurate Energy Modeling and Characterization of IEEE 802.11ah RAW and TWT. Sensors 19. doi:10.3390/s19112614.

Souilah, H., Baadache, A., Bouallouche-Medjkoune, L., 2014. Coping with Spoofed PS-Poll Based DoS Attack in IEEE 802.11 Networks, in: VECoS.

Stepanova, E., Bankov, D., Khorov, E., Lyakhov, A., 2020. On the Joint Usage of Target Wake Time and 802.11ba Wake-up Radio. IEEE Access 8, 221061–221076.

Tian, L., Santi, S., Seferagić, A., Lan, J., Famaey, J., 2021. Wi-Fi HaLow for the Internet of Things: An up-to-date survey on IEEE 802.11 ah research. Journal of Network and Computer Applications 182, 103036.

Wi-Fi Alliance, . Wi-Fi Alliance Product Finder. URL: https://www.wi-fi.org/product-finder.

Woesner, H., Ebert, J., Schläger, M., Wolisz, A., 1998. Power-saving mechanisms in emerging standards for wireless LANs: the MAC level perspective. IEEE Wirel. Commun. 5, 40–48. URL: https://doi.org/10.1109/98.683738, doi:10.1109/98.683738.

Xiao, Y., Cui, Y., Savolainen, P., Siekkinen, M., Wang, A., Yang, L., Ylä-Jääski, A., Tarkoma, S., 2014. Modeling energy consumption of data transmission over wi-fi. IEEE Transactions on Mobile Computing 13, 1760–1773. doi:10.1109/TMC.2013.51.

Yang, C., Lee, J., Bahk, S., 2021. Target Wake Time Scheduling Strategies for Uplink Transmission in IEEE 802.11ax Networks, in: IEEE Wireless Communications and Networking Conference, WCNC 2021, Nanjing, China, March 29 - April 1, 2021, IEEE. pp. 1–6. URL: https://doi.org/10.1109/WCNC49053.2021.9417269, doi:10.1109/WCNC49053.2021.9417269.