ANALYSIS OF IEEE 802.11E FOR QOS SUPPORT IN WIRELESS LANS

STEFAN MANGOLD, PHILIPS RESEARCH
SUNGHYUN CHOI, SEOUL NATIONAL UNIVERSITY
GUIDO R. HIERTZ, OLE KLEIN, AND BERNHARD WALKE, AACHEN UNIVERSITY

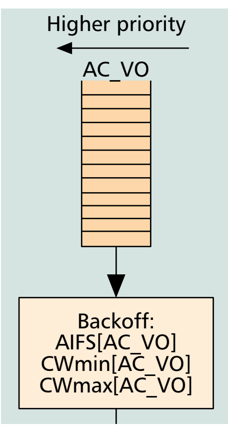

The IEEE 802.11 WLAN is being deployed widely and rapidly for many different environments including enterprise, home, and public access networking. The main characteristics of the 802.11 standard are simplicity and robustness against failures due to the distributed approach of its MAC protocol.

ABSTRACT

The IEEE 802.11e medium access control protocol is an emerging standard for wireless local area networks providing quality of service. An overview of this standard based on the current draft is presented in this article. We analyze the enhancements in 802.11e and compare its performance to the legacy 802.11 standard. The new hybrid coordination function of IEEE 802.11e with its contention-based and contention-free (controlled) medium access control schemes is evaluated. The capability to provide QoS support is discussed by means of simulations.

INTRODUCTION

The *IEEE 802.11 (802.11) wireless local area network (WLAN)* is being deployed widely and rapidly for many different environments, including enterprise, home, and public access networking [1]. The main characteristics of the 802.11 standard are simplicity and robustness against failures due to the distributed approach of its *medium access control (MAC)* protocol [2]. There is a diverse set of versions of WLANs in the market, which apply different transmission schemes and operate in different frequency bands. Operating in the *industrial, scientific, and medical (ISM)* band at 2.4 GHz, the 802.11b version provides data rates up to 11 Mb/s on the wireless medium, applying *complementary code keying (CCK)* and *direct sequence spread spectrum (DSSS)* as transmission schemes. The 802.11a version operates in the unlicensed 5 GHz band, and is able to achieve data rates up to 54 Mb/s on the wireless medium, applying the multicarrier technique *orthogonal frequency-division multiplexing (OFDM)* as the transmission scheme. The 802.11g version applies the same multicarrier transmission scheme as 802.11a, but operates in the 2.4 GHz ISM band like 802.11b. It is worth noting that due to channel conditions and protocol overhead, the maximum achievable throughput on the MAC layer is less than the data rate available on the wireless medium.

Today, 802.11 WLAN (referred to as *legacy 802.11* in this article) can be interpreted as a wireless version of Ethernet supporting best effort service. However, the interest in wireless networks supporting *quality of service (QoS)* has recently grown [1, 3–7]. There are already mechanisms available in the legacy 802.11 to support QoS that have not yet been implemented in real hardware because of their limitations. Accordingly, the 802.11 working group initiated an activity to enhance the current 802.11 MAC protocol to enable support of applications requiring QoS. Analysis of the main results of this activity is the goal of this article.

In this article we discuss the upcoming enhancements of the 802.11e standard as specified in the current draft [8]. Limitations of QoS support in the legacy 802.11 are discussed in the next section. We summarize the new mechanisms for QoS support being defined in 802.11e. A performance evaluation of the described mechanisms with the help of simulations is provided, followed by conclusions.

LEGACY 802.11

In this section the legacy 802.11 MAC protocol defined in [2] is summarized. Its limitations to support QoS are highlighted. We consider an infrastructure *basic service set (BSS)*, which is composed of an *access point (AP)* and a number of stations associated with the AP. The AP connects its stations with the infrastructure such as the Internet, and each associated station communicates only via the AP in legacy 802.11.

DISTRIBUTED COORDINATION FUNCTION

The basic 802.11 MAC protocol is referred to as the *distributed coordination function (DCF)* and operates as a listen-before-talk scheme, as explained in the following.

Carrier Sense Multiple Access — The listen-before-talk scheme of the DCF is based on *carrier sense multiple access* (CSMA). Applying the DCF, a station determines individually when to access the medium. Hence, the decision making process about medium access is distributed among all stations. The station service responsible for information exchange is referred to as *MAC ser-*

1536-1284/03/$17.00 © 2003 IEEE

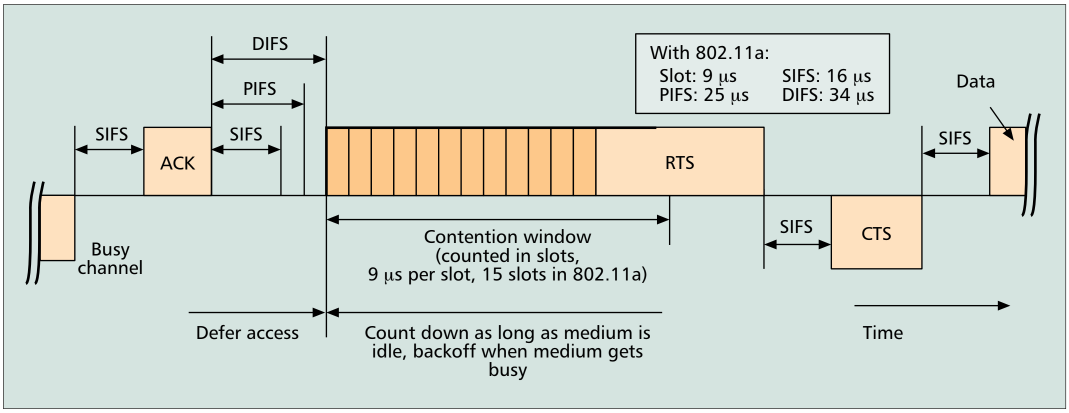

■ **Figure 1.** *[3] Interframe spaces and backoff procedure with random contention window size. Here the transmitting station uses CW = CWmin (15 slots) of 802.11a, and has selected a random backoff time of 12 slots.*

vice data unit delivery (MSDU delivery). MSDUs are transmitted with the help of a *MAC protocol data unit (MPDU)* or, if a station decides to fragment a long MSDU into a number of MPDUs, with the help of several MPDUs. Stations deliver *MSDUs*[1] of arbitrary lengths (up to 2304 bytes) ("talk"), after detecting that there are no other transmissions in progress on the wireless medium ("listen"). If two or more stations detect the medium as being idle at the same time, they may initiate their transmissions at the same time, and inevitably a collision occurs.

Collision Avoidance — To reduce the probability of collisions, the DCF applies a *collision avoidance (CA)* mechanism, where stations perform a so-called *backoff procedure* before initiating a transmission. After detecting the medium as idle for a minimum duration called *DCF interframe space (DIFS)* (34 μs in 802.11a), stations keep sensing the medium (listening) for an additional random time called *backoff time*. A station initiates its transmission only if the medium remains idle for this additional random time. The duration of this random time is determined by each station individually, as a multiple of a slot time (9 μs in 802.11a). A new independent random value is selected for each new transmission attempt. See Fig. 1 for an illustration of the backoff procedure. Since stations select the number of slots at random out of an interval between 0 and *contention window (CW)*, which is initially set to the minimum value CWmin (15 in 802.11a), it is less likely that collisions occur. All stations use the same value for CWmin, but select their random backoff time individually. Since all stations operate with the same CWmin, all stations have the same medium access priority in the DCF. This results in no mechanism to differentiate between stations and their traffic, and therefore no QoS support in the DCF.

If the medium gets busy due to interference or other transmissions while a station is downcounting its backoff counter (i.e., waiting until the random time has elapsed), the station stops downcounting and defers from medium access until the medium becomes idle for a DIFS again (Fig. 2). This occurs, for example, when the random backoff time of a station is longer than the random backoff time of at least one other station. Stations that deferred medium access because of detecting the medium as busy do not select a new random backoff time, but continue downcounting the time of the deferred backoff after sensing the medium as idle again.

Since collisions may occur due to the nature of the CSMA/CA protocol, a station that has transmitted an MPDU needs to be informed about the success of its transmission. Therefore, in legacy 802.11 using the DCF, each transmitted MPDU requires an acknowledgment (ACK). For each successful reception of an MPDU, a receiving station immediately acknowledges the frame reception by transmitting an ACK frame back to the transmitting station. If this ACK frame is not received by the transmitting station right after the MPDU transmission, this transmitting station concludes that the MPDU was not delivered successfully and may repeat the transmission.

The CW of a transmitting station increases when a transmission fails (i.e., the transmitted data frame is not acknowledged). After any unsuccessful transmission attempt, a new backoff procedure is performed with double-sized CW, up to a maximum value defined by CWmax (equals 1023 in 802.11a). The CW is now larger than before to reduce the probability of repeated collisions if there are multiple stations attempting to access the medium. The larger the CW, the smaller the collision probability.

After each successful transmission, the transmitting station initiates another random backoff, even if there is no other pending MSDU to be delivered. This is often referred to as *post-backoff*, as this backoff is done after, not before, a transmission. There is one exception to the rule that a backoff has to be performed before any MPDU transmission. If an MSDU from the higher layer arrives at the station when:

- The transmission queue is empty.
- The latest post-backoff has finished already.
- The medium has been idle for at least one DIFS; it may be delivered immediately without performing the backoff procedure.

The post-backoff guarantees that there is always one random backoff time between two consecutive frame exchanges, and therefore is an important mechanism to guarantee DCF functionality.

Since collisions may occur due to the nature of the CSMA/CA protocol, a station that has transmitted an MPDU needs to be informed about the success of its transmission. Therefore, in legacy 802.11 using the DCF, each transmitted MPDU requires an acknowledgment.

[1] The MSDU can be interpreted as the data frame arriving at the MAC from the higher layer, for example, from the IEEE 802.2 logical link control (LLC) sublayer. The MSDU is transmitted inside one or more MPDUs.

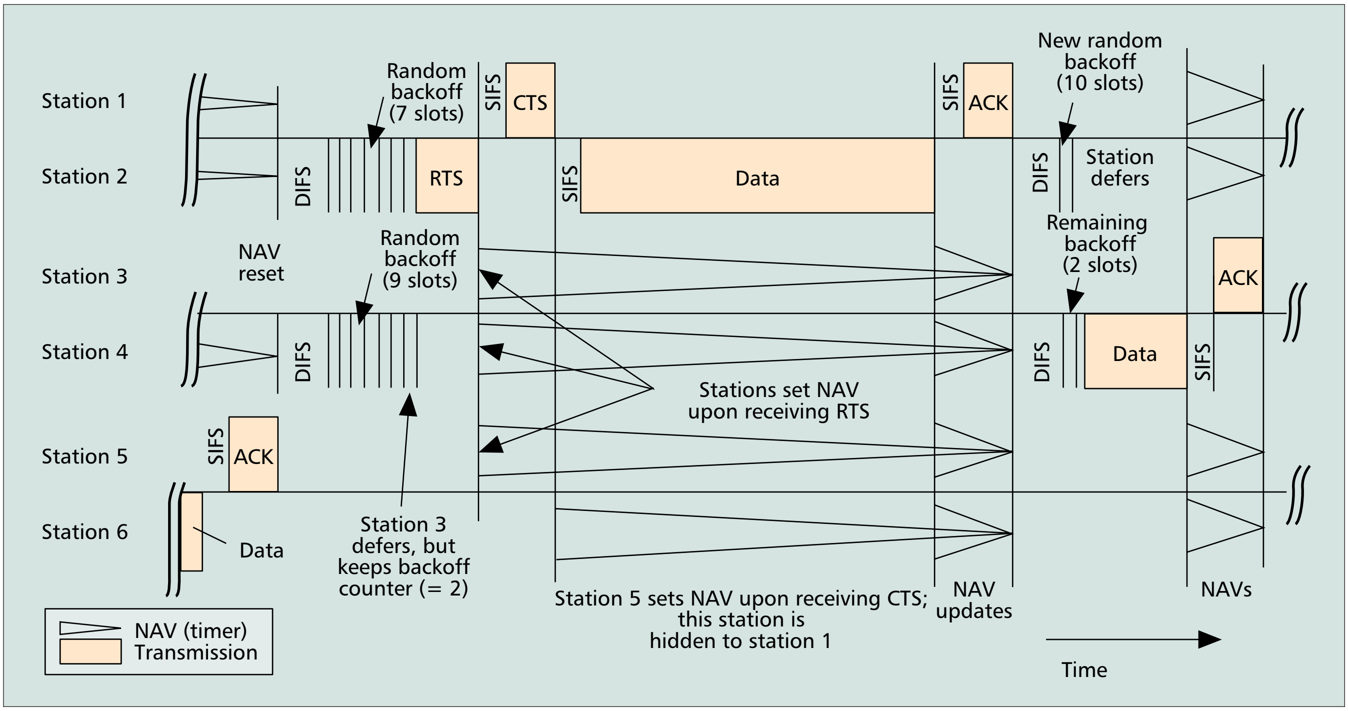

■ **Figure 2.** *Timing of frame exchanges and NAV settings of the 802.11 DCF [3]. Station 6 cannot detect the RTS frame of transmitting station 2, but receives the CTS frame of station 1. Although station 6 is hidden to station 1, it refrains from medium access because of the NAV.*

Optional Fragmentation and RTS/CTS — To reduce the loss of capacity due to collisions of long MPDUs, MSDUs may also be fragmented into multiple MPDUs. A large MSDU may be divided into several smaller MPDUs (i.e., fragments), which can then be transmitted one after the other as a fragment burst with individually acknowledged fragments. One backoff is performed before the transmission of the first fragment. All following fragments are transmitted subsequently without backoff. The benefit of fragmentation is that in case of a collision or other interference resulting in an error, less data has to be retransmitted. The obvious drawback is the increased overhead of fragmentation.

To reduce the hidden station problem, 802.11 uses a *request-to-send/clear-to-send* (RTS/CTS) mechanism, which can be used optionally before MPDU transmission. The hidden station problem refers to a scenario where a station (the "hidden" station) is not able to detect ongoing transmission of other stations because of radio channel conditions. The hidden station may initiate transmissions during the ongoing transmission because it determines the medium as idle, which may result in unacceptable levels of interference at the receiving station. The RTS/CTS works as follows. Before transmitting an MPDU, a station transmits a short RTS control frame, followed by the CTS control frame transmitted by the receiving station. The RTS and CTS frames include information on how long it will take to transmit the next data frame (in fragmentation the first fragment), and the corresponding ACK. Upon receiving either RTS or CTS, stations in transmission range to the transmitting station and hidden stations in transmission range to the receiving station set their timer, called *a network allocation vector* (NAV), with the duration announced within the RTS/CTS frames. The RTS/CTS frames protect the MPDU from interference. Stations that receive these frames will not start any transmission until this timer expires. NAVs are shown as triangles in Fig. 2.

Between two consecutive frames in the sequence of RTS, CTS, MPDU, and ACK frames, a *short interframe space* (SIFS) (16 μs in 802.11a) gives transceivers time to turn around. It is worth noting that SIFS is shorter than DIFS, which gives CTS and ACK always the highest priority access to the wireless medium. No other station can interrupt ongoing frame exchanges (e.g., before an ACK is transmitted) by initiating another transmission.

Limited QoS Support Applying the Point Coordination Function

The legacy 802.11 uses the *point coordination function (PCF)* to support QoS for time-bound services. The PCF provides mechanisms for prioritized access to the wireless medium, and is centrally coordinated by a station called the *point coordinator (PC)*. This station is typically the AP. PCF medium access has higher priority than medium access according to the DCF.

Contention-Free Period — With the PCF, a *contention-free period (CFP)* and a *contention period (CP)* alternate periodically over time, where a CFP and the following CP form one superframe.[2] The PCF is used for accessing the medium during the CFP, whereas the DCF is used during the CP. A superframe must include a CP of a minimum length that allows at least one *MSDU delivery* (one frame exchange) of

[2] *Originally, the phrase* superframe *was defined in the context of 802.11e, but can also be used for the legacy 802.11.*

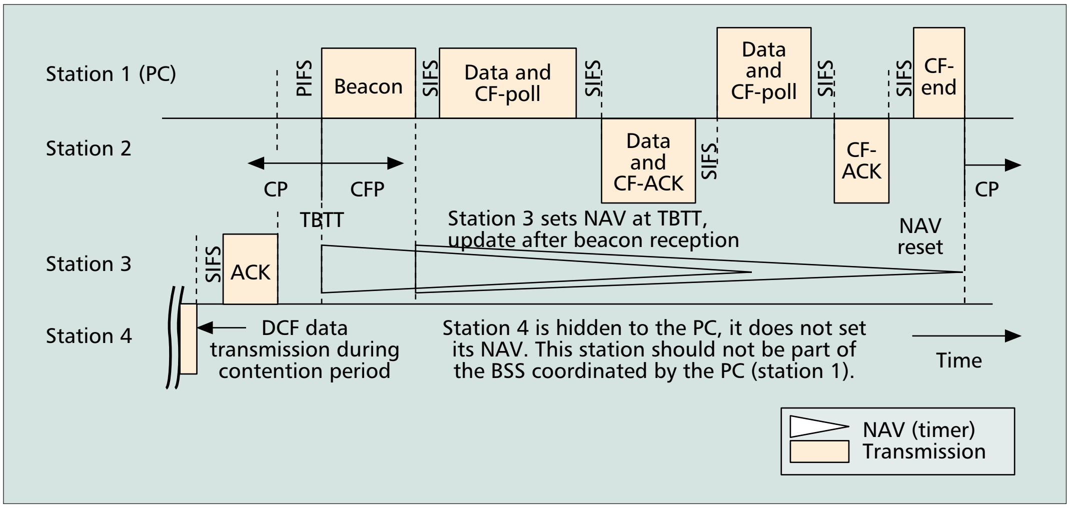

■ Figure 3. *An example of PCF operation [3]. Station 1 is the PC and polls station 2. Station 3 detects the beacon frame and updates the NAV to the whole CFP. Station 3 has learned from earlier beacons that a CFP starts after the TBTT shown here.*

A superframe starts with a beacon management frame transmitted by the AP. This beacon frame is transmitted irrespective of whether or not the PCF is used. Beacon frames are used to maintain the synchronization of the local timers in the stations and to deliver protocol-related parameters.

maximum size and at the slowest transmission rate under DCF.

A superframe starts with a beacon management frame transmitted by the AP. This beacon frame is transmitted irrespective of whether or not the PCF is used. Beacon frames are used to maintain synchronization of the local timers in the stations and to deliver protocol-related parameters. The AP generates beacon frames at regular beacon frame intervals. Every station knows when the next beacon frame will arrive. These points in time are referred to as *target beacon transmission time* (TBTTs) and are announced in the previous beacon frame.

During the CFP, there is no contention among stations; instead, stations are polled. See Fig. 3 for typical frame exchange sequences during a CFP. The PC polls a station asking for the transmission of a pending frame. Whenever the PC itself has a pending frame destined to this station, it uses a combined data and poll frame by piggybacking the CF-Poll frame onto the data frame.

Upon receiving the poll+data, the polled station acknowledges the successful data reception and piggybacks an MPDU. If the PC receives no response from a polled station after waiting for a *PCF interframe space* (PIFS) (25 μs in 802.11a) it polls the next station, or ends the CFP. Therefore, no idle period longer than a PIFS occurs during a CFP. Note that a PIFS is longer than a SIFS, but shorter than a DIFS. Because PIFS > SIFS, a poll is not issued, for example, between data and ACK, so a poll frame cannot interrupt an ongoing frame exchange. The PC continues with polling other stations until the CFP expires. A specific control frame, CF-End, is transmitted by the PC as the last frame within a CFP to indicate the end of the CFP.

Unsolved Problems of PCF — There are problems with the PCF that led to the current activities within the IEEE 802.11 working group to enhance the protocol. Among many others, those include unpredictable beacon delays and unknown transmission durations of the polled stations. At TBTT, a PC schedules the beacon as the next frame to be transmitted, but the beacon can only be transmitted when the medium has been determined to be idle for at least a PIFS. For the legacy 802.11 standard, stations can start their transmissions even if the MSDU delivery is not finished before the upcoming TBTT [7]. Depending on whether the wireless medium is idle or busy at TBTT, a delay of the beacon frame may occur. The time the beacon frame is delayed from TBTT determines the delay of the transmission of time-bounded MSDUs that have to be delivered in the CFP. This may severely affect the QoS as it introduces unpredictable time delays in each CFP. Beacon frame delays of around 4.9 ms are possible in 802.11a in the worst case (longest MSDU, fragmentation, RTS/CTS, most robust modulation and coding scheme).

Another problem with the PCF is the unknown transmission time of polled stations. A station that has been polled by the PC is allowed to deliver an MSDU that may be fragmented and of arbitrary length, up to the maximum of 2304 bytes (2312 bytes with encryption). Furthermore, different modulation and coding schemes are specified in 802.11a. As a result, the duration of MSDU delivery after polling is not under the control of the PC, which reduces the QoS provided to other stations polled during the rest of the CFP.

A detailed description of existing problems in QoS support in legacy 802.11 is provided in [4].

QoS Support Mechanisms of 802.11e

Enhancements to the above-described 802.11 MAC are currently under development. These enhancements will lead to the 802.11e extension of the 802.11 standard. 802.11e introduces the *hybrid coordination function* (HCF) for QoS support [8]. The HCF defines two medium access mechanisms:

- Contention-based channel access
- Controlled channel access (includes polling)

Note that 802.11e uses channel access as a synonym for medium access. Contention-based channel access is referred to as *enhanced dis-*

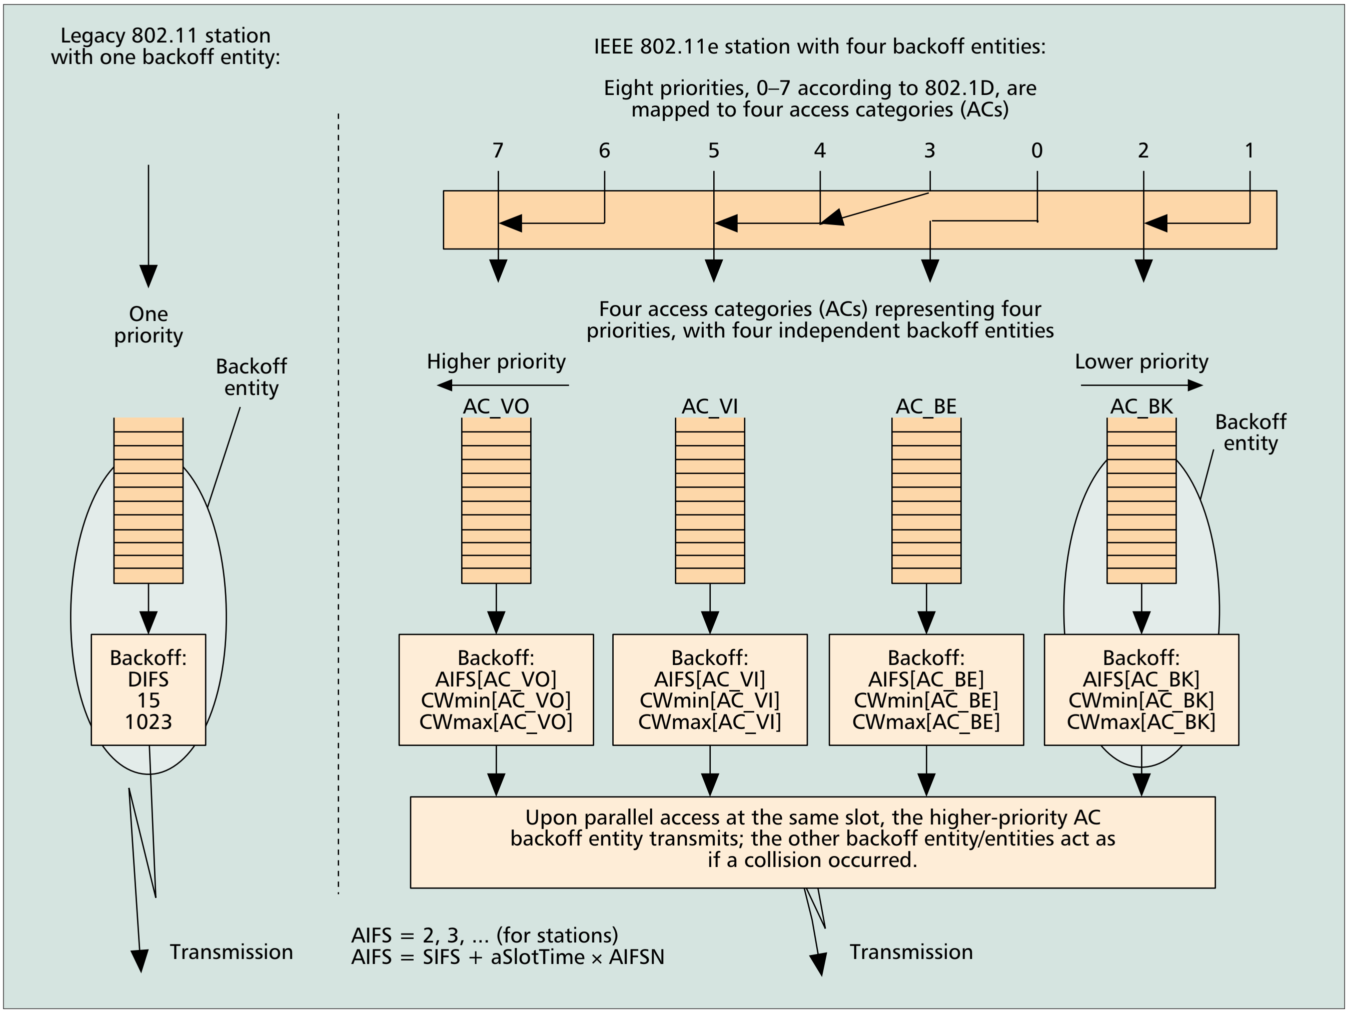

Figure 4. *[3] Legacy 802.11 station and 802.11e station with four ACs within one station.*

tributed channel access (EDCA), controlled channel access as *HCF controlled channel access* (HCCA). With 802.11e, there may still be the two phases of operation within a superframe (i.e., CP and CFP). The EDCA is used in the CP only, while the HCCA is used in both phases. The HCF combines methods of the PCF and DCF, which is the reason it is called hybrid.

Stations operating under the 802.11e protocol are referred to as 802.11e stations in this article. The station that operates as the central coordinator for all other stations within the same *QoS supporting BSS* (QBSS) is called the *hybrid coordinator* (HC). Similar to the PC, the HC resides within an 802.11e AP. A BSS that includes an 802.11e-compliant HC is referred to as a QBSS. There are multiple backoff processes operating in parallel within one 802.11e station, which will be explained later. Therefore, in the following we refer to backoff entities that attempt to deliver MSDUs instead of stations.

Basic Improvements of the Legacy 802.11 MAC

An 802.11e station (more precisely, a backoff entity) that obtains medium access must not utilize radio resources for a duration longer than a specified limit. This important new attribute of the 802.11e MAC is referred to as a *transmission opportunity* (TXOP). A TXOP is an interval of time during which a backoff entity has the right to deliver MSDUs. A TXOP is defined by its starting time and duration. TXOPs obtained via contention-based medium access are referred to as EDCA-TXOPs. Alternatively, a TXOP obtained by the HC via controlled medium access is referred to as HCCA-TXOP or polled TXOP.

The duration of an EDCA-TXOP is limited by a QBSS-wide parameter referred to as *TXOPlimit*. This TXOPlimit is distributed regularly by the HC within an information field of the beacon.[3] However, in the absence of legacy stations, the TXOPlimit allows control of the maximum time a backoff entity allocates the medium for MSDU delivery, and therefore is an important means to control MSDU delivery delay.

Another enhancement is that no backoff entity transmits across the TBTT. That is, a frame exchange is initiated only if it can be completed before the upcoming TBTT. This reduces the expected beacon delay, which gives the HC better control over the medium, especially if the optional CFP is used after beacon transmission.

Additionally, an 802.11e backoff entity is allowed to transmit frames directly to another backoff entity in a QBSS, without involving communication with the AP. In the legacy 802.11

[3] Legacy stations will only understand the fields known from the legacy standard, whereas 802.11e backoff entities additionally will understand all new information fields. The new information fields are ignored by legacy stations. Therefore, legacy stations may transmit for longer durations than allowed by the TXOPlimit.

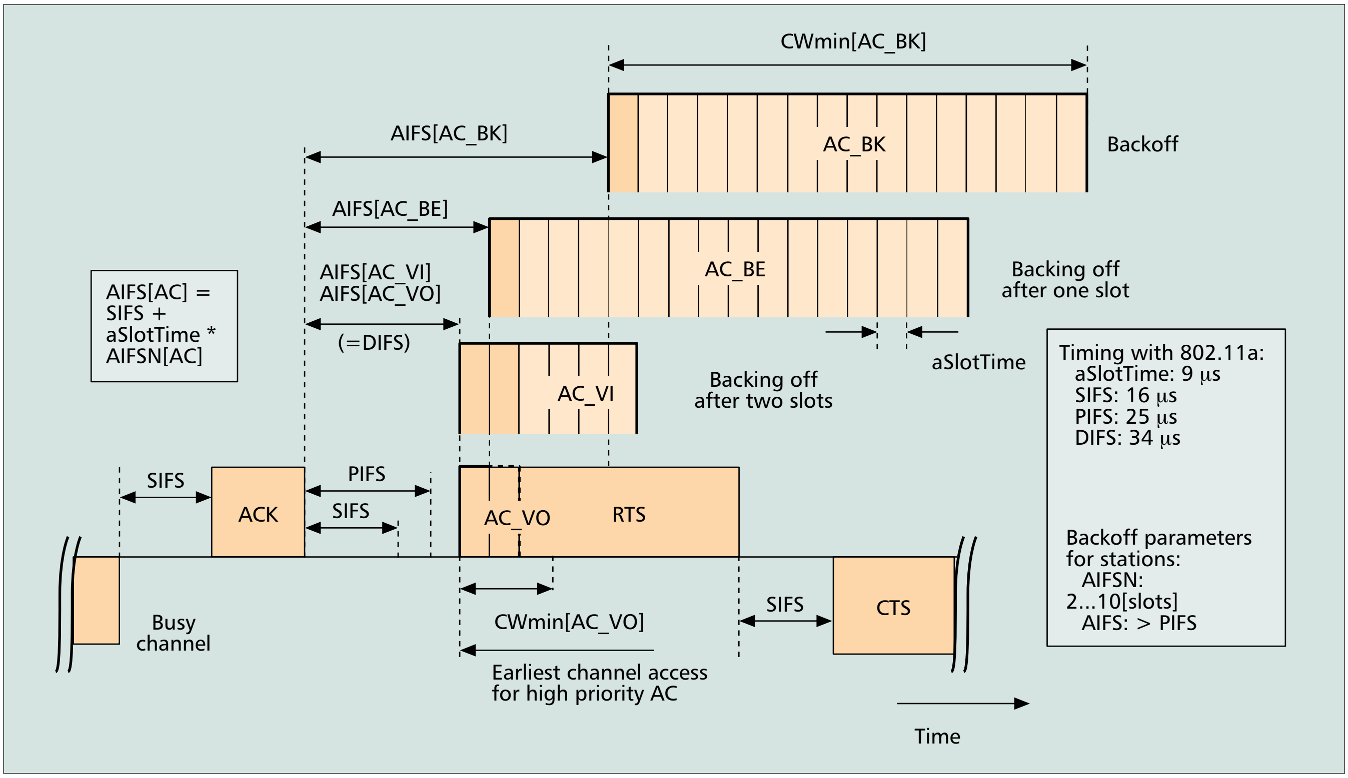

Figure 5. *In EDCA, multiple backoff entities contend for medium access with different priorities in parallel [3]. The earliest possible medium access time after a busy medium is DIFS.*

protocol, within an infrastructure-based BSS all data frames are either sent or received by the AP. For this purpose, an 802.11e station needs to establish a direct link with another 802.11e station using the *direct link protocol (DLP)* before initiating direct frame transmissions.

HCF Contention-Based Medium Access

The QoS support in EDCA is provided by the introduction of *access categories (ACs)* and multiple independent backoff entities. MSDUs are delivered by parallel backoff entities within one 802.11e station, where backoff entities are prioritized using AC-specific contention parameters, called EDCA parameter set. There are four ACs; thus, four backoff entities exist in every 802.11e station. The ACs are labeled according to their target application, i.e., AC_VO (voice), AC_VI (video), AC_BE (best effort), and AC_BK (background). See Fig. 4 for an illustration of the parallel backoff entities. The EDCA parameter set defines the priorities in medium access by setting individual interframe spaces, contention windows, and many other parameters per AC, as explained below.

EDCA Parameters per AC — Contention-based medium access is performed in every backoff entity by using different parameter values for the EDCA parameter set. Which values to be used by which backoff entity is defined by the HC. The EDCA parameter set can be modified over time by the HC, and is announced via information fields in beacon frames. The same EDCA parameter set is used by the backoff entities of the same AC in different stations. It is essential that the same values for the parameters are used by all backoff entities.

Each backoff entity within a station independently contends for a TXOP. It starts down-counting the backoff-counter after detecting the medium being idle for a duration defined by the *arbitration interframe space* (*AIFS[AC]*) instead of DIFS, which is used by legacy stations. The *AIFS[AC]* is at least DIFS, and can be enlarged per AC with the help of the *arbitration interframe space number (AIFSN[AC])*. The *AIFSN[AC]* defines the duration of *AIFS[AC]* according to

$$AIFS[AC] = SIFS + AIFSN[AC] \cdot aSlotTime, \quad AIFSN[AC] \geq 2.$$

AIFSN[AC] should be selected by the HC such that the earliest access time of EDCA stations is DIFS, equivalent to legacy 802.11. The parameter aSlotTime defines the duration of a slot. The smaller the *AIFSN[AC]*, the higher the medium access priority.

The minimum size of the contention window, *CWmin[AC]*, is another parameter dependent on the AC. The initial value for the backoff counter is a random number taken from an interval defined by the CW, similar to legacy DCF. See Fig. 5 for an illustration of the *AIFS[AC]* and *CWmin[AC]*. Three priorities are shown in the figure. The smaller the *CWmin[AC]*, the higher the priority in medium access. A big difference between legacy DCF and 802.11e EDCA in terms of the backoff countdown rule is as follows:

- The first backoff countdown occurs at the end of the *AIFSN[AC]* interval.
- A frame transmission is initiated after a slot from the moment when the backoff counter becomes zero.

However, the collision probability increases with

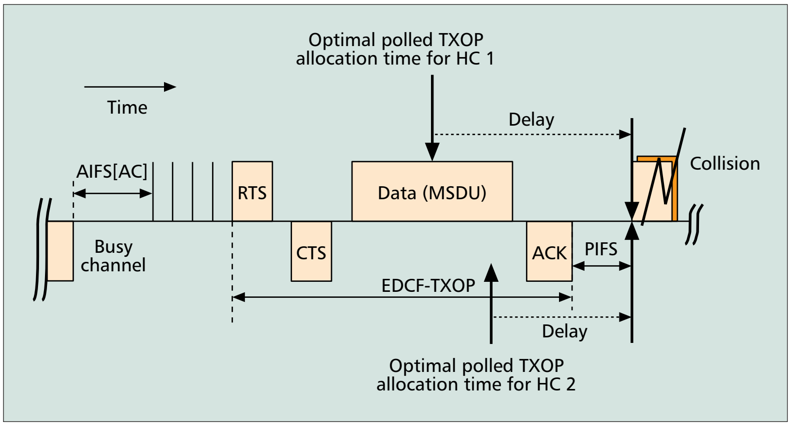

■ Figure 6. *A polled TXOP allocation [3]. Any 802.11e frame exchange will not take longer than the TXOPlimit, which is the limit for all EDCA-TXOPs and under control of the HC.*

	AC_VO	AC_VI	AC_BE	AC_BK	High (AC H)	Medium (AC M)	Low (AC L)
AIFSN:	2	2	3	7	2	4	7
CWmin:	3	7	15	15	7	10	15
CWmax:	7	15	1023	1023	7	31	255
	(Used for throughput evaluation, EDCA parameters from [8])				(Used for delay evaluation of OQBSS [1])		

■ **Table 1.** *Values for the EDCA parameter sets as used in simulations.*

smaller *CWmin[AC]* if there are more than one backoff entities of the respective AC operating in the QBSS. If *AIFSN[AC]* is selected such that the earliest medium access time is DIFS, priority over legacy stations can be supported by setting *CWmin[AC]* < 15 (for 802.11a).

The positions and sizes of the contention windows relative to each other, as defined per AC by the EDCA parameter set, are important factors to define relative priority in medium access per AC. The contention window increases upon unsuccessful frame exchanges, but never exceeds the value of *CWmax[AC]*. This parameter is defined per AC as part of the EDCA parameter set. The smaller the *CWmax[AC]*, the higher the medium access priority. However, a small *CWmax[AC]* may increase the collision probability. Furthermore, it should be highlighted that there are retry counters (similar to legacy 802.11) that limit the number of retransmissions. The 802.11e protocol also defines a maximum MSDU lifetime per AC, which specifies the maximum time a frame may remain in the MAC. Once the maximum lifetime has passed since a frame arrived at the MAC, the frame is dropped without being transmitted. This feature can be useful since transmitting a frame too late is not meaningful to many real-time applications.

In addition to the backoff parameters, the *TXOPlimit[AC]* is defined per AC as part of the EDCA parameter set. The larger *TXOPlimit[AC]* is, the larger the share of capacity for this AC. Once a TXOP is obtained using a backoff, a backoff entity may continue to deliver more than one MSDUs consecutively during the same TXOP, which may take up to the duration of *TXOPlimit[AC]*. This important concept in 802.11e is referred to as continuation of an EDCA-TXOP.

As described above, four backoff entities with different EDCA parameter sets reside inside an 802.11e station. During contention, when the counters of two or more backoff entities in the same station reach zero at the same time, a virtual collision occurs. Upon access to the same slot by more than one backoff entity of one station, the backoff entity with the higher priority will transmit, whereas all other backoff entities will act as if a collision occurred on the medium. It may still occur that the transmission of the backoff entity with the higher probability collides with another transmission initiated by other stations.

Table 1 is a short version of a table in [8], and shows some of the recommended default values for the EDCA parameter sets for the four ACs. A HC may use these values when setting up a QBSS, and may change them dynamically upon changes of medium and traffic conditions.

Hybrid Coordination Function, Controlled Medium Access

The controlled medium access of the HCF, referred to as *HCF controlled channel access* (HCCA) extends the EDCA access rules by allowing the highest priority medium access to the HC during the CFP and CP. The details about the controlled medium access are summarized in this section.

A TXOP can be obtained by the HC via the controlled medium access. The HC may allocate TXOPs to itself to initiate MSDU Deliveries whenever it requires, after detecting the medium as being idle for PIFS, and without backoff. To give the HC higher priority over legacy DCF and EDCA access, *AIFSN[AC]* must be selected such that the earliest medium access for EDCA stations is DIFS for any AC.

During CP, each TXOP of an 802.11e station begins either when the medium is determined to be available under the EDCA rules, that is, after *AIFS[AC]* plus the random backoff time, or when a backoff entity receives a polling frame, the QoS CF-Poll, from the HC. The QoS CF-Poll from the HC can be transmitted after a PIFS idle period, without any backoff, by the HC. During CFP, the starting time and maximum duration of each TXOP is also specified by the HC, again using the QoS CF-Poll frames. During CFP, 802.11e backoff entities will not attempt to access the medium without being explicitly polled, hence, only the HC can allocate TXOPs by transmitting QoS CF-Poll frames, or by immediately transmitting downlink data. During a polled TXOP, a polled station can transmit multiple frames that the station selects to transmit according to its scheduling algorithm, with a SIFS time gap between two consecutive frames as long as the entire frame exchange duration is not over the allocated maximum TXOPlimit.

Polled TXOP allocations may be delayed by the duration of an EDCA-TXOP, as illustrated in Fig. 6. The HC controls the maximum duration of EDCA-TXOPs within its QBSS by

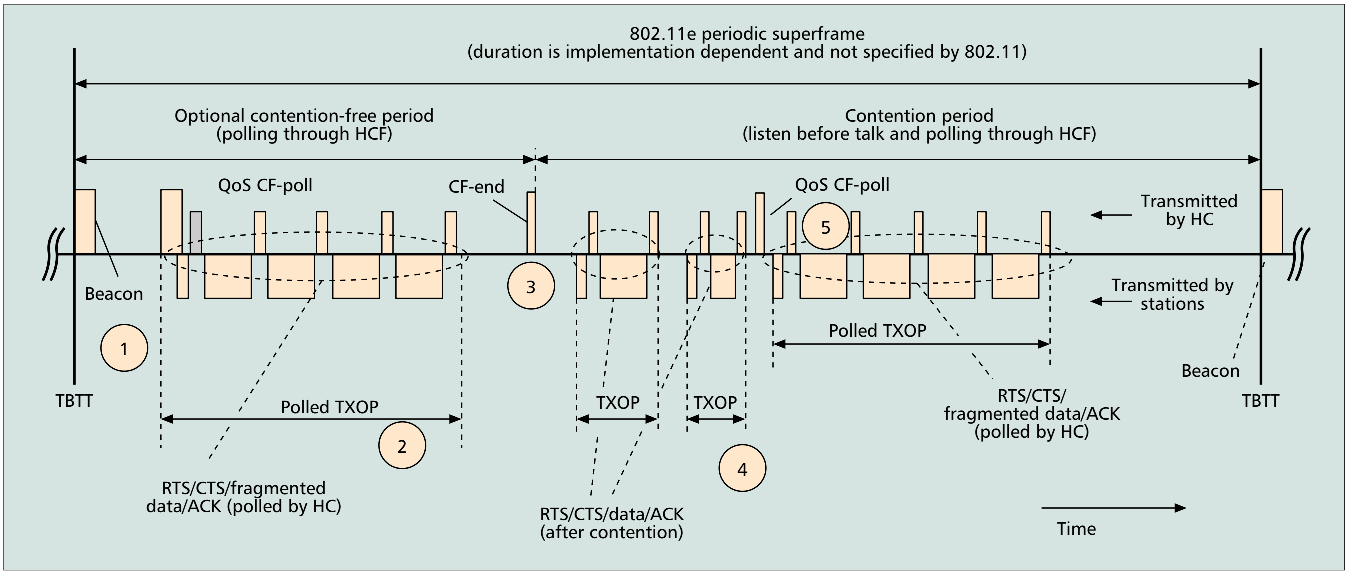

■ Figure 7. *An example of an 802.11e superframe where the HC grants TXOPs in contention-free period and contention period [3]. The duration of the superframe is not specified in the standard.*

announcing the *TXOPlimit[AC]* for every AC via the beacon. Therefore, it is able to allocate polled TXOPs at any time during the CP, and the optional CFP. When very small MSDU Delivery delays are required, CF-Polls may be transmitted a duration of *TXOPlimit[AC]* earlier than the optimal polled TXOP allocation time to avoid any MSDU delivery delay imposed by EDCA-TXOPs at all. However, the largest *TXOPlimit[AC]* of the four ACs must be considered.

Figure 7 illustrates an example of a superframe that includes a CFP and a CP. The superframe starts with a beacon transmitted by the HC (indicated with (1)). During the CFP (i.e., the first part of the superframe), the backoff entities only transmit upon being polled by the HC. Indicated with (2) is the transmission of a fragmented MSDU within the CFP. The CFP ends with the CF-End frame transmitted by the HC as shown at (3). During the following CP, all backoff entities attempt to transmit through the contention-based medium access of the HCF (i.e., the EDCA). EDCA-TXOPs are obtained through contention. Two such EDCA-TXOPs are indicated with (4). During the CP, the HC can also poll a station, which is different from the PCF of the legacy 802.11. This is shown as an example in the figure. Following the two EDCA-TXOPs, the HC polls a station to allocate a polled TXOP during which a fragmented MSDU is transmitted, as indicated with (5).

Improved Efficiency

More schemes that will improve the efficiency of the MAC protocol are becoming part of 802.11e, and are briefly summarized here.

Block Acknowledgment — With the optional block acknowledgment, the throughput efficiency of the protocol is improved. Block acknowledgments allow a backoff entity to deliver a number of MSDUs being delivered consecutively during one TXOP and transmitted without individual ACK frames. The MPDUs transmitted during the TXOP are referred to as a block of MPDUs. At the end of the block, or in a later TXOP, all MPDUs are acknowledged by a bit pattern transmitted in the block acknowledgment frame, thus reducing the overhead of control exchange sequences to a minimum of one acknowledgment frame per number of MPDUs delivered in a block.

Direct Link Protocol (DLP) — Any backoff entity can directly communicate with any other backoff entity in a QBSS, without communicating via the AP. In the legacy 802.11 protocol, within an infrastructure-based BSS (which is denoted as BSS), all data frames are sent to the AP, and received from the AP. This, however, consumes at least twice the channel capacity compared to the direct communication. Only in an independent BSS (which is denoted IBSS), station-to-station communication is allowed in the legacy protocol, due to the absence of the AP. The direct communication in 802.11e is referred to as *direct link* (DiL). A setup procedure, the *Direct Link Protocol* (DLP), is defined to establish a DiL between 802.11e backoff entities.

Evaluation

We use event-driven stochastic simulations to evaluate the performance of the 802.11e MAC for the 802.11a *physical layer (PHY)* at 5GHz that allows up to 54 Mb/s. For the delay results of the MSDU delivery, we give empirical *complementary cumulative distribution functions* (CDFs) of the resulting stochastic data, using the discrete *Limited-Relative-Error* (LRE) algorithm that also measures the local correlation of the stochastic data [9]. By measuring local correlations, the accuracy of empirical simulation results can be estimated. All results here are within a maximum limited relative error of 5 percent.

A radio channel error model as described in

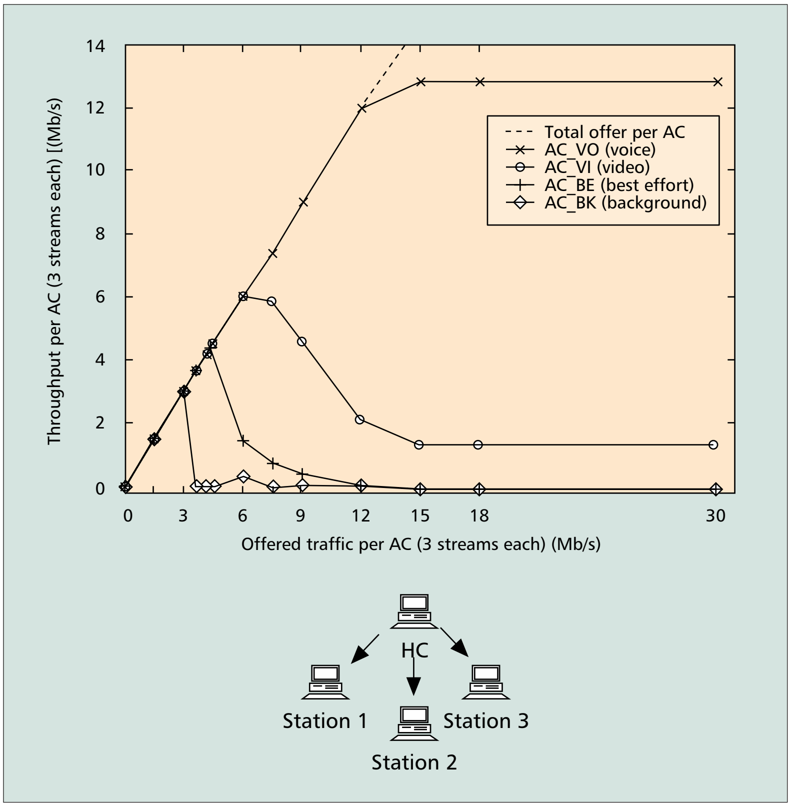

■ **Figure 8.** *Throughput per AC with increasing offered traffic per AC for the illustrated scenario.*

[5, 6] is used. Transmission powers and distances between stations are chosen in such a way that they are not hidden to each other with the selected PHY modes. If not stated otherwise, control frames are transmitted at 6 Mb/s, and data frames are transmitted at 24 Mb/s. Each station generates the same mixture of offered traffic of four data streams, which we label according to the known ACs. A simple traffic model is used to make sure that the characteristic of the traffic sources do not influence the results, which would lead us to misinterpretations (for example, a correlated packet arrival of realistic voice traffic would have impact on throughput results). If not stated otherwise, at all backoff entities, MSDUs of *512 byte* with negative-exponentially distributed inter-arrival times are offered. Each stream carries 250 kb/s.

We use the EDCA parameters as provided in Table 1. Neither RTS/CTS nor fragmentation is used. The duration of EDCA-TXOPs allows stations to transmit one data frame after winning the contention in EDCA. Beacon frames are not generated.

Achievable EDCA Throughput in an Isolated QBSS with Four Stations

The achievable (saturation) throughput depends on a large number of parameters. The results given here are valid only for the MAC and PHY settings we use in our simulation. In the simulation discussed in this and the following sections, all frames are transmitted at 24 Mb/s. Figure 8 is derived from a simulation with a single station (the AP) transmitting four streams to each of three stations, one stream per priority. Hence, the AP transmits 12 streams in total. The default EDCF parameters of the four ACs (AC_VI … AC_BK) are used (Table 1). One data frame per EDCF-TXOP is transmitted. The measured throughput per AC and the scenario are illustrated in Fig. 8.

It is known that the achievable throughput in 802.11 depends on the size of the delivered MSDUs, and the PHY modes used. The larger the MSDU size, the higher the achievable throughput, as long as the number of collisions and amount of interference are small. The higher the PHY mode, the higher the achievable throughput, as long as the channel conditions are sufficient. From Fig. 8 we observe that with increasing offered traffic, ACs with higher priority restrain the achievable throughput of the ACs with lower priority, thanks to smaller AIFSN, CWmin, and CWmax values. An analytical approximation of the achievable throughput per AC in such a scenario is described in [3]; there it is referred to as *share of capacity*.

Achievable EDCA Throughput with Increasing Number of Stations

In this section we evaluate QoS support with EDCA medium access in an isolated QBSS, as shown in Fig. 9. All stations rely again on prioritized medium access over the EDCA only (no polling-based access is utilized). Each station transmits using all four backoff entities ACs (250 kb/s each); hence each station attempts to carry 1 Mb/s in total.

To investigate the performance of 802.11e in hot spots, where it is likely that a larger number of stations are associated with an AP, we increase the number of contending stations up to 16, as indicated in Figure 9. Each station offers the same traffic. Figure 9 shows the measured throughput per AC, similar to the previous figure. It can be observed that with the default values for the EDCA parameter set, the throughput is reduced dramatically once the number of stations becomes larger. This is a result of the higher collision probabilities, particularly in AC_VO, where CWmin and CWmax are small numbers. This is clearly not a desirable result. However, in this case the HC should set new values for the EDCA parameters, which would then be used by all stations of the QBSS. It should be noted that the highest-priority AC of EDCA is expected to be used for voice-like calls in real-life scenarios, which means less offered traffic per station at AC_VO (assuming one call per station).

QoS Guarantees with Prioritized Access of HC

We now evaluate the effect of allowing the highest priority to the HC medium access, according to HCCA. This section is based on an earlier evaluation of 802.11e [1]. The values for the EDCA parameter set are here selected according to AC_H (high priority), AC_M (medium priority), and AC_L (low priority), as given in Table 1. Note that the HCCA will show the same tendencies if the values for the EDCA parameter set of the previous sections were selected for evaluation. At the high-priority AC, MSDUs of 80 bytes arrive at constant periods. The period depends on the offered traffic, and is 5 ms for offered traffic of 128 kb/s. To the medium- and low-priority ACs, MSDUs of 200 bytes with negative exponentially distributed interarrival times are offered. Each stream carries 160 kb/s. Data frames are transmitted at 24 Mb/s, control frames (ACKs) at 6 Mb/s. In contrast to the previous scenarios, where all stations including the AP contend for medium access via EDCA, the AP now carries an additional isochronous downstream (80 bytes/MSDU, 128 kb/s) delivered with the HCCA priority, which has higher priority than AC_H. That is, the data frames of this stream are immediately transmitted after a PIFS when the medium is detected as idle. Note that the HCCA achieves its strict delay requirements in the CP by setting a maximum TXOP duration for all other streams.

Figure 10 shows the resulting MSDU delivery delay distributions for an isolated QBSS and for overlapping QBSSs. The overlapping QBSS scenario is illustrated in the figure, as well as the resulting delays for the HCCA stream and AC_M. It can be seen that whereas the delays of the EDCA increase unpredictably with increased offered traffic, the HCCA delays remain below a certain threshold, which is defined by the TXOPlimit. Only the HCCA stream stays within its maximum delay limit.

There is a situation of significantly increased MSDU delivery delays even for the HCCA stream, which is not under the control of the HC and therefore is undesirable. That is, two (or more) QBSSs are co-located to each other, so that they interfere to each other. In this scenario, even polled data frames of highest priority suffer from an unpredictable delay and throughput degradation due to uncoordinated resource sharing between HCs. One result of such a scenario is given in Fig. 10. It can be seen that the delays of the high-priority stream exceed the TXOPlimit defined by the HC, which attempts to limit the MSDU delivery delays of the HCCA stream to 300 μs. Note that the given result is an example for a variety of delays and throughputs that can be observed in overlapping QBSSs. One observable example occurs if the two HCs always poll at the same time. Then all poll frames collide, and the HCCA throughput drops to zero.

For the overlapping QBSS problem, solution concepts are under discussion in the 802.11 working group. One solution would be to apply dynamic frequency selection to let a QBSS dynamically select a free medium. This approach is of interest for the 5 GHz band with its many available frequency channels. Other approaches are based on policies [3].

Summary and Conclusions

An overview of the QoS supporting protocol extensions of 802.11e is presented in this article. New mechanisms for QoS support (i.e., the EDCA and HCCA) are evaluated, and their performance is discussed with the help of simulation results. The upcoming 802.11e standard will be an efficient means of QoS support in WLANs for a wide variety of applications, although open

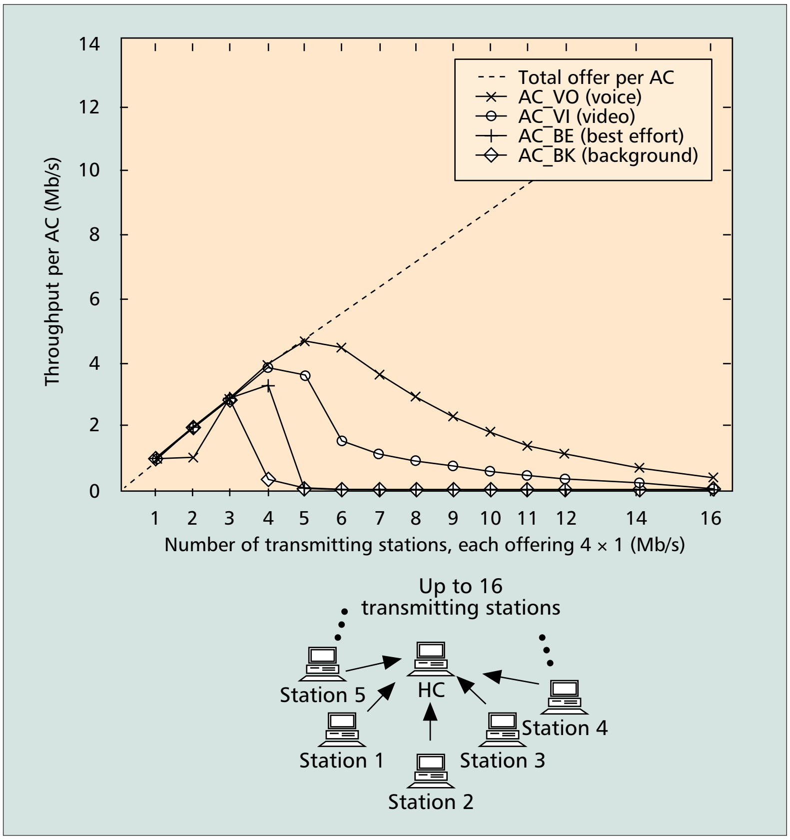

Figure 9. *Throughput per AC with increasing number of stations, and constant offered traffic per station for the illustrated scenario.*

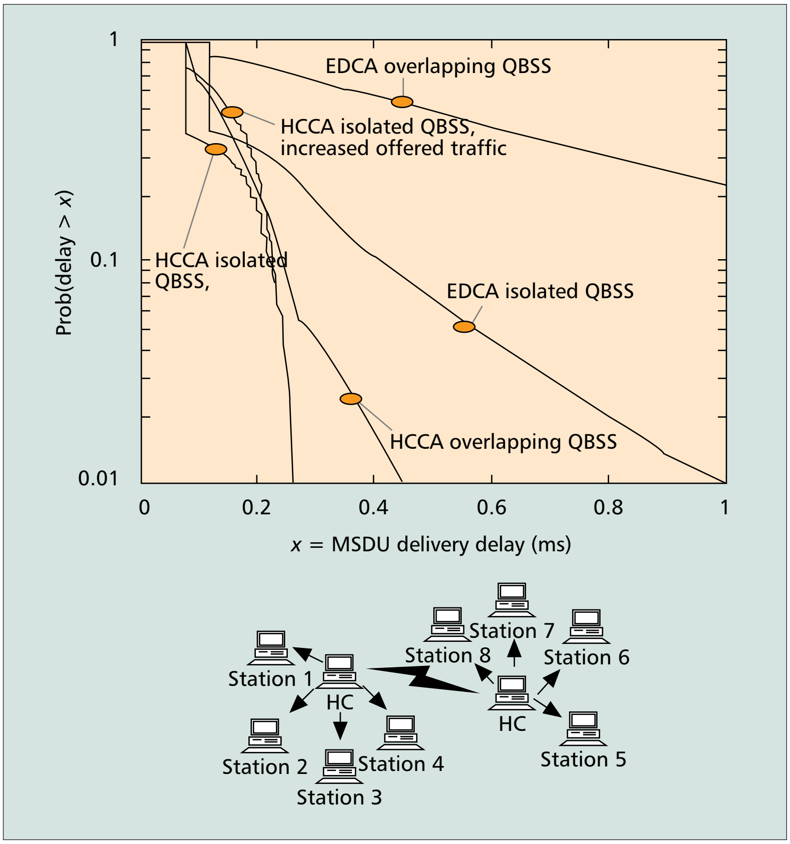

Figure 10. *MSDU delivery delays for AC_M (EDCA) and HCCA, in isolated and overlapping QBSSs. The increase of the HCCA delay in overlapping QBSSs is not under control of an HC, and therefore undesirable.*

The upcoming 802.11e standard will be an efficient means for QoS support in WLANs for a wide variety of applications, although open problems such as the overlapping QBSS still remain to be solved.

problems such as the overlapping QBSS still remain to be solved. The HCCA provides the means to deliver time-bounded traffic, but requires all stations within the range of the HC to follow its coordination.

ACKNOWLEDGMENTS

The authors wish to thank Javier del Prado and Kiran Challapali of Philips Research for their valuable comments and contributions to this work.

REFERENCES

[1] S. Mangold *et al.*, "IEEE 802.11e Wireless LAN for Quality of Service," *Proc. European Wireless '02*, Florence, Italy, Feb. 2002.
[2] "International Standard [for] Information Technology — Telecommunications and Information Exchange between Systems — Local and Metropolitan Area Networks — Specific Requirements — Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," IEEE 802.11 WG, ref. no. ISO/IEC 8802-11:1999(E) IEEE Std. 802.11, 1999.
[3] S. Mangold, "Analysis of IEEE 802.11e and Application of Game Models for Support of Quality-of-Service in Coexisting Wireless Networks," Ph.D. thesis, ComNets Aachen University, June 2003.
[4] S. Choi, "Emerging IEEE 802.11e WLAN for Quality-of-Service (QoS) Provisioning," *SK Telecom Telecommun. Rev.*, vol. 12, no. 6, Dec. 2002, pp. 894–906.
[5] S. Mangold, S. Choi, and N. Esseling, "An Error Model for Radio Transmissions of Wireless LANs at 5GHz," *Proc. Aachen Symp. 2001*, Aachen, Germany, Sept. 2001, pp. 209–14.
[6] D. Qiao, S. Choi, and K. G. Shin, "Goodput Analysis and Link Adaptation for IEEE 802.11a Wireless LANs," *IEEE Trans. Mobile Comp.*, vol. 1, no. 4, Oct.–Dec. 2002, pp. 278–92.
[7] B. Walke, *Mobile Radio Networks*, 2nd ed., Wiley, 2001.
[8] IEEE 802.11 WG, IEEE 802.11e/D5.0, "Draft Supplement to Standard for Telecommunications and Information Exchange Between Systems — LAN/MAN Specific Requirements — Part 11: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Medium Access Control (MAC) Enhancements for Quality of Service (QoS)," Aug. 2003.
[9] F. Schreiber, "Effective Control of Simulation Runs by a New Evaluation Algorithm for Correlated Random Sequences," *AEÜ Int'l. J. Elec. and Commun.*, vol. 42, no. 6, 1988, pp. 347–54.

BIOGRAPHIES

STEFAN MANGOLD (stefan.mangold@philips.com) is a senior member of research staff at Philips Research Laboratories, Briarcliff Manor, New York. Before joining Philips in March 2003, he worked at the Chair of Communication Networks (ComNets) at RWTH Aachen University, Germany. He received his Diploma degree (M.S.) and Dr.-Ing. degree (summa cum laude) in electrical engineering in 1997 and 2003, respectively, both from RWTH Aachen University, Germany. His current areas of research are spectrum agile radio, autonomous spectrum management, and radio spectrum etiquette. He is the author of about 35 technical papers and book chapters, and contributes to the standardization of IEEE 802.11.

SUNGHYUN CHOI [S'96, M'00] (schoi@snu.ac.kr) is an assistant professor at the School of Electrical Engineering of Seoul National University (SNU), Seoul, Korea. Before joining SNU in September 2002, he worked for Philips Research, Briarcliff Manor, New York, as a senior member of research staff for three years. He received his B.S. (summa cum laude) and M.S. degrees in electrical engineering from Korea Advanced Institute of Science and Technology in 1992 and 1994, respectively, and received a Ph.D. from the Department of Electrical Engineering and Computer Science of the University of Michigan, Ann Arbor in 1999. His current research interests are in the area of wireless/mobile networks with emphasis on QoS guarantee and adaptation, resource management, wireless LAN and PAN, next-generation mobile networks, data link layer protocols, and connection and mobility management. He has authored/co-authored over 40 technical papers and book chapters in the areas of wireless/mobile networks and communications. He is currently serving on program committees of a number of leading wireless and networking conferences. He is also an active participant and contributor of the IEEE 802.11 WLAN Working Group. He was a recipient of the Korea Foundation for Advanced Studies Scholarship and the Korean Government Overseas Scholarship during 1997–1999 and 1994–1997, respectively.

GUIDO R. HIERTZ [M] (grh@comnets.rwth-aachen.de) studied electrical engineering at RWTH Aachen University. During his studies he focused on communication networks, computer science, and engineering. In 2000 he installed one of the first German 802.11 networks connecting a student home to the university. In his diploma thesis he included 802.11e functionality in an advanced simulation tool and surveyed QoS supporting procedures. Since April 2002 he is with the Chair of Communication Networks (ComNets), RWTH Aachen University, working toward his Ph.D. His main research fields are protocols for gigabit WLANs with support for QoS in multihop networks. Therefore, he is working on decentralized access procedures and forwarding strategies. He has published several papers at IEEE conferences.

OLE KLEIN (ole@comnets.rwth-aachen.de) received his Diploma in electrical engineering in 2000 from RWTH Aachen University. In 2001 he joined ComNets, where he is working toward his Ph.D. He participated in the MultiHop project Cellular OFDM System with Extension Points for Increased Transmission Range (COVERAGE), where he was involved in the implementation and integration of the trail platform. His research interests include protocols for WLANs and new radio access systems.

BERNHARD H. WALKE [SM] (walke@comnets.rwth-aachen.de) has more than 35 years of activity in this area, and is well connected in industry and academia. For the last 13 years he has run the Chair for Communication Networks at Aachen University, where about 40 researchers work on topics like air interface design, formal specification of protocols, fixed network planning, development of tools for stochastic event-driven simulation, and analytical performance evaluation of services and protocols of 2G to 4G wireless systems. During that time he has supervised more than 650 Master's theses and 43 Ph.D. theses covering practically any aspects of fixed and mobile communication networks. Most of this work continuously has been funded from third party grants. He has published more than 110 scientific conference papers, 25 journal papers, and seven textbooks on the architecture, traffic performance evaluation, and design of future communication systems. His latest books are *Mobile Radio Networks* (2nd ed., Wiley, 2002) and *UMTS — A Comprehensive Introduction* (Wiley, 2003). He has been a board member of ITG/VDE, and has served for more than a decade as Program Committee and Steering Committee Chair of scientific European conference series like the European Personal Mobile Radio Communications Conference and European Wireless conferences that he co-founded. Since 1996 he has served as scientific coordinator for large German national research programs like Mobile ATM, Multi-Hop, and MOBIKOM. Since 1993 his group has participated in more than 15 large research projects funded by the European Commission and has substantially contributed to the specification of standards like ETSI/GPRS, ETSI/BRAN HiperLAN2, CEN TC 278, and IEEE 802.11e. Furthermore, the MAC protocols of IEEE 802.16a and 802.15.3 are based on his group's findings. In 2001 he was elected Chair of Working Group 4 (New Technologies) of the Wireless World Research Forum. WG4 is responsible to define interfaces for beyond 3G systems. Prior to joining academia, he worked in various positions for AEG Telefunken (now part of EADS AG). He holds a diploma degree and a Ph.D. in electrical engineering, both from the University of Stuttgart, Germany.