RZ 3186 (# 93232) 11/22/99
Computer Science/Mathematics 18 pages

Research Report

IEEE 802.11 Wireless LAN Standard
A Technical Tutorial

Reto J. Hermann

IBM Research
Zurich Research Laboratory
8803 Rüschlikon
Switzerland

IBM Research
Almaden · Austin · Beijing · Haifa · T.J. Watson · Tokyo · Zurich

IEEE 802.11 Wireless LAN Standard
A Technical Tutorial

Reto J. Hermann

IBM Research, Zurich Research Laboratory, 8803 Rüschlikon, Switzerland

Abstract

The IEEE 802.11 standard for wireless local area networks (LANs) is the basis for a growing number of wireless networking products including network adapters and access points to wired networks. More and more of these products will be deployed as today's trend towards untethered computing accelerates and an increasing number of people will need to acquire a basic understanding of 802.11 wireless LANs. This tutorial is intended as a gentle introduction into the subject for the technically interested reader.

Contents

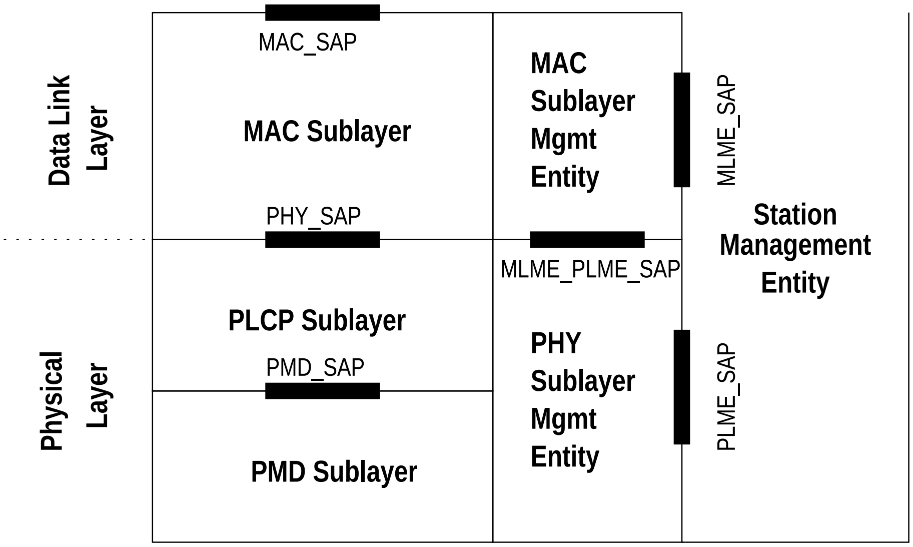

Figure 1: Protocol architecture and OSI reference model

.

1 Introduction

802.11 [1] is a standard developed by the IEEE 802 LAN/MAN Standards Committee for wireless local area networks (LAN) that operate in the gloabally available ISM (industrial, scientific & medical) band between 2.4 and 2.5 GHz. The ISM band exhibits regional variations because frequency allocations are under government regulation; exact frequency ranges and RF power levels may, therefore, differ from country to country. The standard specifies operation for both *ad hoc* and *infrastructure* networks with data rates ranging from 1 to 11 Mbits/s.

The 802.11 specification fits into the overall framework of 802 standards, which primarily address functionality in layers 1 and 2 of the OSI (Open System Interconnection) Reference Model. In particular, 802.11 specifies the media access control (MAC) of the data link layer and the physical layer (PHY). The carrier sense multiple access with collision avoidance (CSMA/CA) protocol can operate across multiple compatible physical layer transmission systems including frequency-hopping (FH) and direct-sequence (DS) spread-spectrum (SS) systems and infrared (IR)-based transmission systems. The architectural model is illustrated in Figure 1 and defines a number of protocol sublayers and protocol entities as well as the service interfaces between them.

The purpose of this tutorial is to provide an overview of 802.11 for the technically interested reader who wants to get a good understanding of the concepts and functional scope of this standard. We tried to be concise and still remain as accurate as possible. However, it was unavoidable to omit several subtle details; readers requiring full details are referred to the official standard document [1]. Regarding terminology, we adopt the standard's vocabulary (e.g., access point instead of base station) so that the reader who turns to the standard can read it with ease.

The tutorial is organized around concepts and function rather than around the architectural model as shown in Figure 1. Section 2 introduces the overall architecture and defines the basic terminology. In Section 3, we give detailed descriptions of the MAC mechanisms and MAC frame structure. Section 4 talks about networking aspects and Section 5 gives a short overview of the link level security concepts. Section 6 describes the power management mechanisms and, finally, section 7 very briefly addresses the PHY layer concentrating on the spread spectrum transmission systems.

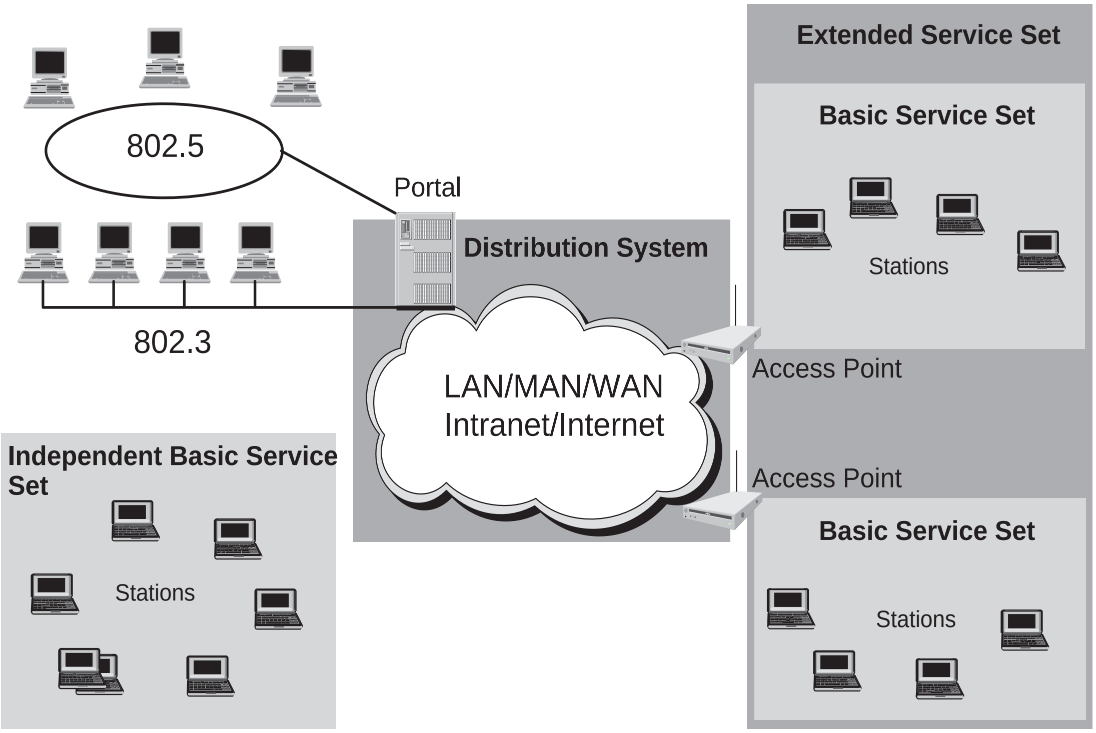

Figure 2: 802.11 architecture concepts

.

2 Architecture

This section presents fundamental concepts of the IEEE 802.11 architecture and introduces some of the terminology as it is used in the standard. Figure 2 illustrates these concepts and shows the complete picture of the architecture.

2.1 Stations

A *station* (STA) in 802.11 refers to the device that is addressed via an 802 address. In wired networks where devices are typically stationary, the address coincides with a fixed location. In wireless networks STAs are typically no longer stationary and the association of address and location breaks. 802.11 makes a distinction between *portable* and *mobile* stations. A portable station is one that can be moved from location to location but is in operation only when stationary, whereas mobile stations operate in transit. Although this distinction seems intuitive from the user's perspective, the difference does not really exist at the physical level due to the nature of radio propagation and the ways in which it may vary (for instance because of a person entering or leaving a room). In other words, a portable station that is operated for a certain period at a fixed location experiences a transient radio propagation environment in much the same way as a roaming mobile station. The communication entity in the station cannot tell the difference and therefore the protocols have to be designed to deal with intermittent radio connectivity.

2.2 Basic Service Set

The basic network building block of an 802.11 LAN is the *basic service set* (BSS). The basic service set consists of stations that can communicate with each other directly via the wireless medium. STAs that can communicate with each other must be within range of one another. A STA leaving the range of the other STAs in its BSS can no longer communicate with these STAs; we say it has left the BSS. Conversely, a STA coming within range of the STAs of a BSS may communicate with them; we say it has entered the BSS. As stations can be turned on and off or enter and leave the BSS, the association of a STA and a BSS is transient. BSSs may be spatially disjoint, partially overlapping or collocated. A BSS is identified via a 802 MAC address.

2.2.1 Independent BSS

The most basic configuration consists when two stations form a so-called *ad hoc network*, i.e. a network that does not require any configuration or auxiliary infrastructure. Ad hoc networks of two or more STAs are called an *independent BSS* (IBSS) in 802.11 parlance.

2.2.2 Extended Service Set

A network of interconnected BSSs forms an *extended service set* (ESS). The interconnection of the BSSs is accomplished by the *distribution system* (DS). The distribution system medium is logically separated from the wireless medium. The only requirement that must be met is that there be one STA in each BSS that serves as an *access point* (AP) to the DS. Networks requiring an accesss point are commonly called *infrastructure networks* (in contrast to ad hoc networks). The nature of the DS is not defined by 802.11. Non-802.11 LANs from the 802.x family are integrated with the DS via a logical component called a *portal*.

A common network configuration in practice is multiple BSSs interconnected via an 802.x wired LAN. In this case the DS medium is the 802.x wired LAN and the APs of the BSSs are all portals.

2.3 Architectural Services

In order to support the logical separation of the DS from the rest of the 802.11 architectural components, the interfaces between them are defined in terms of logical services.

2.3.1 Station Services

The service provided by a station is called *station service* (SS). It comprises

Authentication 802.11 provides link-level authentication with which stations establish their identity to stations with which they will communicate.

Deauthentication Deauthentication invalidates a previously established authentication between two stations.

Privacy 802.11 provides link level encryption by which stations can protect the information exchanged against eavesdropping.

MSDU service (MAC Service Data Unit) delivery.

2.3.2 Distribution System Services

The service provided by the DS is called *distribution system service* (DSS). It comprises

Association In a DS the distribution service must perform a routing function by forwarding a message to the proper AP for the given STA addressed. Association is the concept that provides this information under the assumptions that STAs remain within their original BSS. The association service provides the STA-to-AP mapping to the DS and is invoked by an STA prior to sending any messages.

Disassociation Disassociation revokes in the DS the STA-to-AP mapping established via the preceding association.

Distribution Distribution is the main service of the DS. It accomplishes the task of delivering a message from the AP of entrance to the appropriate exit AP for the destination specified.

Integration Integration is similar to distribution, but deals with messages exchanged between an STA in an 802.11 DS and non-802.11 LANs. For messages originating from an STA and distributed to a portal, the DS invokes the integration function. In the opposite direction, for messages originating in the non-802.11 LAN entering via a portal, the DS invokes the integration function prior to distributing the message.

Reassociation Reassociation complements the association service in situations where STAs can roam from one BSS of the DS to another. Essentially, reassociation "moves" the association from one AP to the other, thus keeping the STA-to-AP mapping of the DS current.

Note that the protocol between APs required to implement distribution, integration and reassociation services are beyond the scope of 802.11. As a result, interoperability between vendors products is usually limited even if they use the same PHY layer. Typically, STAs from one vendor and APs from another interoperate under the constraint that all STAs operate under a single AP. Furthermore, APs from different vendors usually do not interoperate.

3 Medium Access Control

The 802.11 MAC sublayer defines two related medium access coordination functions, the *distributed coordination function* (DCF) and the *point coordination function* (PCF). As the name suggests, the DCF is a decentralized access scheme and thus is available both in ad hoc and infrastructure networks. Architecturally, the PCF is layered on top of the DCF, that is, it uses the media access mechanisms of the DCF in providing its function.

The DCF uses a distributed algorithm called *carrier sense multiple access with collision avoidance* (CSMA/CA). CSMA is a well-known concept and we shall therefore emphasize here only those concepts that are special to 802.11. There are significant differences between wirebound and wireless media. With wirebound media all connected stations can hear each other. This is not necessarily the case in a wireless environment where finite propagation limits the extent of the medium relative to the sending station. In other words, for each station the extent of the medium and thus the set of "connected" stations varies. For exactly this reason, the sending station cannot reliably detect a collision that occurs at the receiving station due to another *hidden station* transmitting concurrently. As a consequence, the collision detection scheme used with wirebound media fails in the wireless environment. In CSMA/CA, this is compensated via the collision avoidance mechanisms and fast MAC-level acknowledgments.

The PCF uses a point coordinator, which controls access to the medium by periodically distributing access control information to STAs in the BSS. STAs operating under PCF thus gain contention-free access to the medium. PCF is available only with infrastructure networks.

DCF and PCF can operate concurrently within the same BSS. Essentially, total time is split into alternating *contention-free periods* (CFP) and *contention periods* (CP).

3.1 Beacon Mechanism

The *beacon* mechanism is central to the operation of 802.11 and provides for the periodic distribution of system parameters that are crucial to overall system operation. These parameters are contained in the so-called *beacon frames* and include:

Time Stamp The time stamp is used for synchronization between all STAs participating in a BSS.

Beacon interval The beacon interval defines the period with which beacon frames are scheduled to occur.

SSID The service set ID carries the ID of the ESS or IBSS.

Physical Parameter Set Physical-layer-dependent information that allows the synchronization of STAs.

Contention-Free Parameter Set Information to support the operation of the PCF.

IBSS Parameter Set Information to support the operation of an IBSS.

The time instance at which a beacon frame nominally occurs is called the *target beacon transmission time* (TBTT). Owing to the way medium access works, the actual transmission of beacon frames may be deferred. However, irrespective of deferred beacon transmissions, the TBTTs are on a schedule of equidistant time instances separated by the beacon interval. By definition, time zero is a TBTT.

In an infrastructure network (ESS), beacon generation is the obligation of the AP. In particular, the AP determines the beacon interval and time zero. All other STAs synchronize their timers and adopt the beacon interval when they join the network.

In an ad hoc network (IBSS), beacon generation is distributed. The beacon interval and time zero are defined by the STA that initiates the network. All STAs in the network share the responsibility of generating beacons. Essentially, each STA chooses a random length window (which is much shorter than the beacon period) starting at TBTTs. If no beacon is observed within the window, the STA generates one and transmits it itself.

3.2 Distributed Control Function - CSMA/CA

CSMA/CA relies on the existence of a time gap between contiguous frame sequences during which the medium is idle. This idle period must be asserted by an STA prior to starting transmission. If the STA asserts a busy medium it must defer the transmission until after the end of the current ongoing transmission. After deferral or after finishing transmission, STAs must select a random discrete-time backoff interval which they decrement while they assert an idle medium. Once its backoff interval is equal to zero a STA may start to contend for the medium. The random selection of the backoff interval makes collisions less likely. The probability of collisions can be further reduced by the RTS/CTS (Request-To-Send/Clear-To-Send) mechanism via which the sending/receiving STA use very short control frames to announce a busy period of the medium/acknowledge receipt of the announcement frame.

Interframe Spaces

The time gap between frames is called the *interframe space* (IFS). IFSs of varying length provide access priority levels. STAs have to wait for the IFS before they can attempt to capture the medium. As a consequence, an STA authorized to use an IFS of a given length I_S has priority over all STAs using an IFS $I_L > I_S$. The following four IFSs given in order of increasing length are defined:

SIFS The short IFS is the shortest IFS. STAs may use the SIFS once they have gained access to the medium and need to keep it to complete the frame exchange sequence started. Such a sequence is also called *fragment burst*. The SIFS cannot be shorter than the minimum amount of time it takes the transceiver to switch from transmit to receive mode.

PIFS The PCF IFS is used by STAs operating in PCF mode in order to gain prioritized access to the medium at the start of the contention-free period.

DIFS The DCF IFS is used by STAs operating in DCF mode for the transmission of data frames and management frames.

EIFS The extended IFS is used by STAs operating in DCF mode whenever their physical layer signals the receiption of an erroneous frame. As frame headers can contain duration information, which is the basis for correct operation of the virtual carrier sensing mechanism (see below), the STA that has received the erroneous frame may have an invalid NAV (*network allocation vector*). The duration of the EIFS is chosen such that another STA that has received the frame correctly has sufficient time to acknowledge the frame. The STA that has received the erronous frame can switch back to DIFS once it has received a frame correctly, ensuring that its virtual carrier sensing is based on valid duration information.

3.2.1 Carrier Sensing

CSMA/CA can run on top of different compatible physical layers, which all must provide a certain set of services comprising transmit, receive and carrier sense primitives. Crucial to the algorithm is the carrier sense mechanism. CSMA/CA distinguishes between two carrier sense schemes, physical carrier sensing and virtual carrier sensing. The medium is considered busy if either one of the schemes indicates a busy medium.

Physical carrier sensing With physical carrier sensing the state of the medium is asserted by the physical layer and signalled to the MAC sublayer via an indication. The actual sensing mechanism depends on the physical layer.

Virtual carrier sense With virtual carrier sensing, the state of the medium is asserted by the MAC sublayer based on the so-called *network allocation vector* (NAV) maintained in each STA, which is based on duration information distributed in RTS/CTS frames and headers of frames that are sent during the contention period.

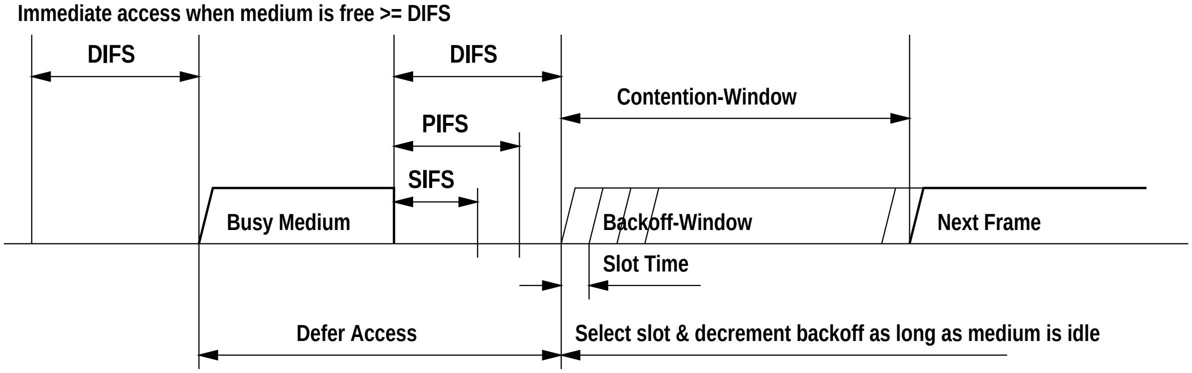
Immediate access when medium is free >= DIFS
DIFS
DIFS
Contention-Window
PIFS
SIFS
Busy Medium
Backoff-Window
Next Frame
Slot Time
Defer Access
Select slot & decrement backoff as long as medium is idle

Figure 3: Basic access method

A station that intends to transmit a packet first transmits a very short control frame called Request-To-Send (RTS). This frame contains the source, destination and duration of the subsequent exchange (packet and ACK). The destination station confirms receipt of the RTS frame with the so-called Clear-To-Send (CTS) control frame, which basically repeats the same duration information. All STAs receiving either one of the control frames update their NAV. The CTS is necessary because collisions at the receiving STA could be caused by an STA that cannot “hear” the STA transmitting the RTS. Announcing the duration information at both the transmitting and receiving STA lowers the probability of collisions at either location.

Virtual carrier sensing creates overhead due to the RTS/CTS frame exchanges. On the other hand, the small size of the frames reduces the likelihood that a collision will occur and, should a collision occur, less air time will have been wasted. Whether to use virtual carrier sensing depends on the nature of the data traffic. For long data packets it should be enabled, otherwise disabled. The threshold for the packet length is a parameter known as *RTS threshold* that can be configured for each STA. Note that STAs that themselves do not distribute duration information are still required to update their NAV when they receive duration information.

3.2.2 Basic Access and Random Backoff

Let us first look at the basic access method. An STA that wishes to transmit a frame senses the carrier and defers any attempt to access the medium until it has determined that the medium was idle for at least a DIFS or EIFS period (see above for an explanation of how EIFS differs from DIFS). Once it has asserted an idle medium for the required period, it starts transmitting. The problem with this basic approach is that if multiple STAs have been deferring their transmissions, all will start their transmissions at about the same time and collisions are therefore likely to occur. The random backoff procedure tries to remedy this situation.

With random backoff, in addition to asserting an idle medium for a DIFS or EIFS period, STAs have to defer their transmission by a random amount of time determined by the STA’s current backoff time. This is illustrated in Figure 3.

An STA with non-zero backoff time defers until it has sensed an idle state of the medium for at least the duration equal to the backoff time, where this duration need not be contiguous. An STA’s backoff time can be 0; this can only be the case if an STA has just been powered up and performs its first carrier sensing or if an STA has finished a previous backoff. The backoff time is initialized to a random value in the interval [0, CW] whenever an STA senses a busy medium and has zero backoff time. The value CW itself is adaptively chosen in an

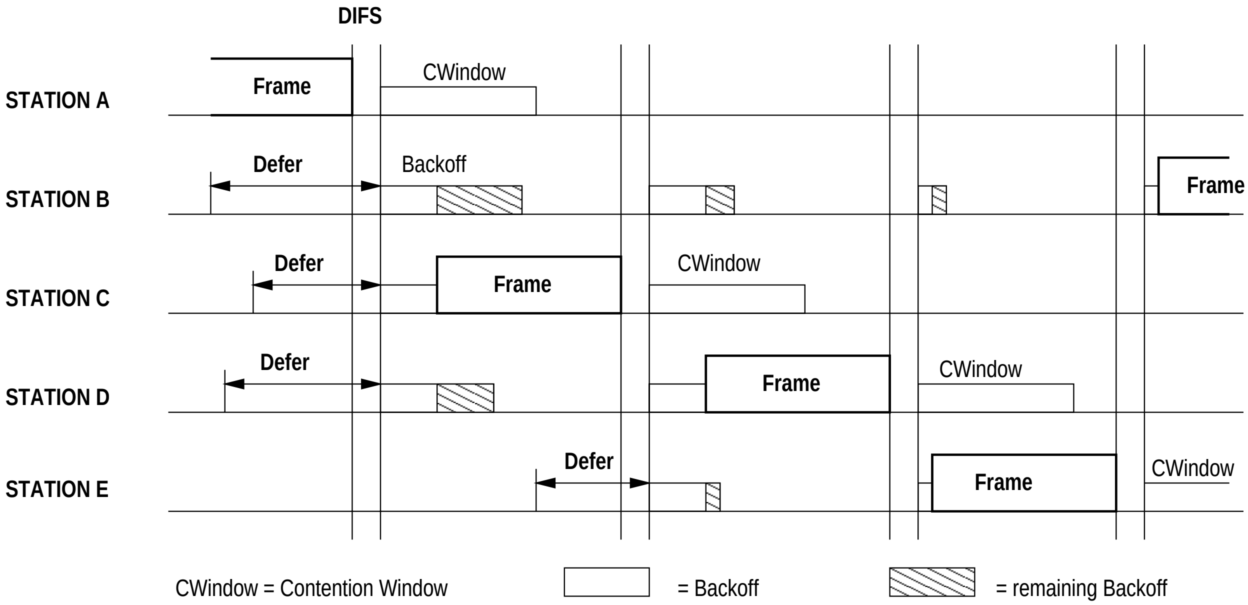

Figure 4: Backoff procedure

exponentially progressing sequence in the interval [CW$_{min}$,CW$_{max}$] depending on the success rate of previous packet transmissions. In essence, the net effect of this procedure in the situation where multiple STAs contend for the medium is that the STA selecting the smallest backoff time will win the contention.

Figure 4 illustrates this access scheme. STA A has just finished the transmission of a frame. It immediately starts contending again for the medium with the three STAs B, C, D. STA C has the shortest backoff time, wins and transmits the next frame. Meantime STA E has also entered the contest. At the point in time that the medium becomes idle again, STA D has the smallest backoff time and thus gets to transmit the next frame. Even though STA E has entered the contest last, its backoff time happened to be shorter than that of STA B. Therefore, STA E wins over STA B and transmits the next frame. Finally, STA B has the shortest backoff time and can transmit its frame.

3.3 Point Control Function – Polling

PCF relies on one STA performing a control function whereby it periodically gains prioritized access to the medium for the duration of the CFP during which it polls the other STAs in the BSS. This STA is called the *point coordinator* (PC) and must be an AP. Because the PCF is piggybacked onto the DCF, all STAs are in principle capable of operating correctly in the presence of a BSS in which a PC is active even if they do not implement polling related functionality (see below).

CFPs occur with a certain frequency determined by the *CFP repetition interval.* This interval is divided into the CFP and the CP. During the CFP, access is regulated by the PC whereas during the CP, pure DCF is used. The length of the CFP is variable but may not exceed a certain maximum duration. The PC adapts the length according to the amount of available traffic and the number of STAs to be polled. The actual start of the CFP may be delayed due to an ongoing fragment burst that is not interrupted (SIFS is less than PIFS).

At the beginning of the CFP the PC gains access to the medium via the basic DCF access mechanism using PIFS and then announces via a beacon frame the timing structure for this CFP, which is called the *CF parameter set.* This includes the CFP repetition interval, the

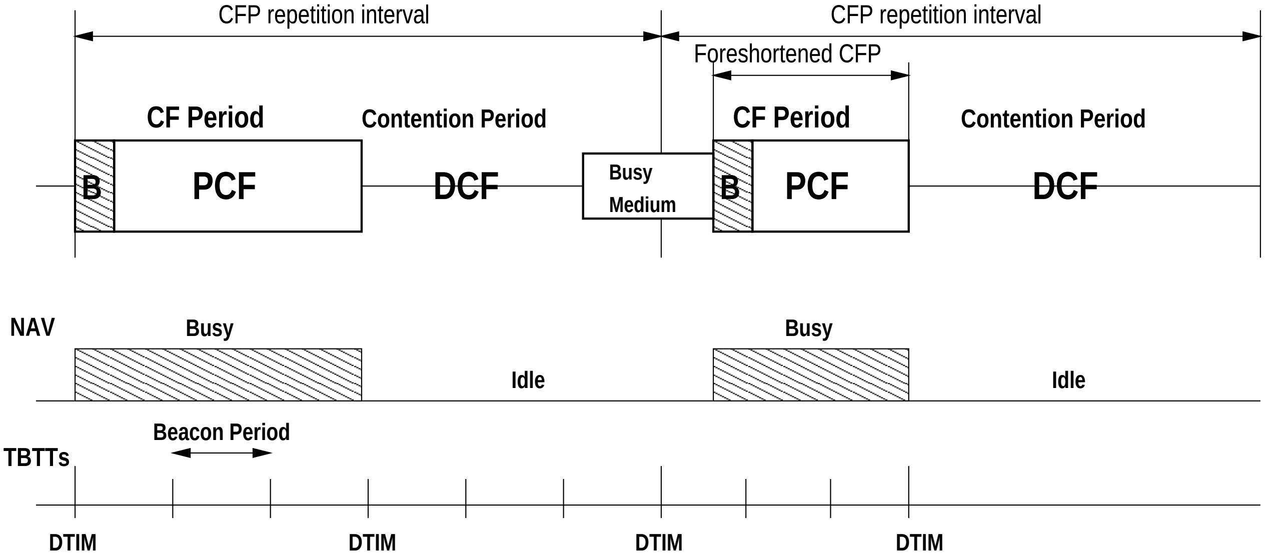

Figure 5: Alternation of CFP/CP

maximal duration of the CFP, the remaining duration of the CFP in the case of a delayed start (which must be chosen such that it ends before the next nominal CFP start time). In order to reduce the likelihood of collisions, non-PC STAs set their NAV at the nominal CFP start time corresponding to the maximal duration of the CFP irrespective of whether they have received the beacon frame or not. They update the NAV based on the information received in the CF parameter set of a correctly received beacon frame; in particular, they set the NAV corresponding to the remaining CFP duration time.

During the CFP, the PC regulates data traffic by sending polling frames to STAs. A pollable STA that receives a CF-Poll frame may transmit a single data frame one SIFS period after receiving the CF-Poll. A non-pollable STA that receives a CF-Poll frame simply responds with an acknowledgment.

The PC maintains a so-called *polling list* to select the STAs that can potentially be polled during the CFP. The duration of the CFP and the size of the polling list may not allow all STAs to be polled in each CFP. The PC therefore polls disjoint subsets of STAs on the polling list and thus cycles through the list in the course of multiple CFPs.

The polling list is established during the association procedure of a STA and the AP. A STA may indicate that it is non-pollable, pollable but not to be put on the list, pollable and to be put on the list, or never to be polled. Conversely, the AP may indicate that it is no PC, a PC for delivery only (no polling) or a PC for both delivery and polling. An STA may change the established polling behavior by performing a reassociation procedure.

Figure 5 illustrates the alternation between PCF and DCF operation. As access to the medium was deferred, the duration of the second CFP is shortened in order to meet the next scheduled start of the CP.

3.4 MAC Frame Format

3.4.1 Basic Frame Structure

The basic structure of the 802.11 MAC frame is shown in Figure 6. Accordingly, a MAC frame consists of a header, a frame body, and a frame check sum, which contains an IEEE 32-bit cyclic redundancy code. Let us look at some of the subfields relevant to the other material presented in this tutorial.

The frame header can have up to four address fields; not all frames carry all four address fields. The address fields can hold the BSSID, DA, SA, RA, or TA (see Section 4.3 below). The frame control field contains the protocol version to allow for proper protocol selection in the case of future deviating versions of the specification. The type field indicates the type of the frame: management, control, or data frame. Within each type there exists a number of subtypes, which are identified by the subtype field. Next, there is a set of flags: The ToDS/FromDS flag is set to 1 if the frame is destined to/coming from the DS and 0 otherwise. The power management flag indicates the power management mode (see Section 6) of the STA sending the frame. The MoreData field indicates to an STA in power-save mode that more data is buffered for it at the AP. The WEP flag is set to 1 if the content of the frame body has been processed by the WEP algorithms (see Section 5).

3.4.2 Selected Frame Formats

The complete set of frame formats is described in detail in the standard. Here, we discuss only a couple of frames that are crucial to the operation of the MAC.

Beacon and probe-response frames The beacon frame and the probe-response frame contain information relevant to the synchronization between STAs, which is the foundation of a network (see Section 4). This information includes:

time stamp to synchronize the timers of the STAs in a network (see Section 4.2).

beacon period the interval between beacons (see Section 3.1).

capability the capability information field provides details about offered or required capabilities such as whether an STA is an AP, whether an STA can be or wants to be polled, whether an AP supports PCF at all/for delivery/for delivery and polling, or whether WEP should be enabled in an IBSS.

SSID the service set identifier.

FH parameter set the frequency-hopping parameter set provides such details as the hop set, hopping pattern, dwell-time, etc. of the spread-spectrum system in the PHY.

DS parameter set the direct-sequence parameter set indicates the channel for the spread-spectrum system in the PHY (see Section 7.2).

CF parameter set the contention-free parameter set provides such details as the time to the start of the next CF period, the interval between CF periods, the maximal duration of CF periods and the duration of the next CF period (see Section 3.3).

IBSS parameter set contains the ATIM window parameter as used for power management in IBSSs (see Section 6.2).

Request-to-send/clear-to-send frames The RTS/CTS frames are used for virtual carrier sensing (see Section 3.2.1). They simply carry the duration required to complete a pending data or management frame exchange including the acknowledgment frame with the indicated addressee.

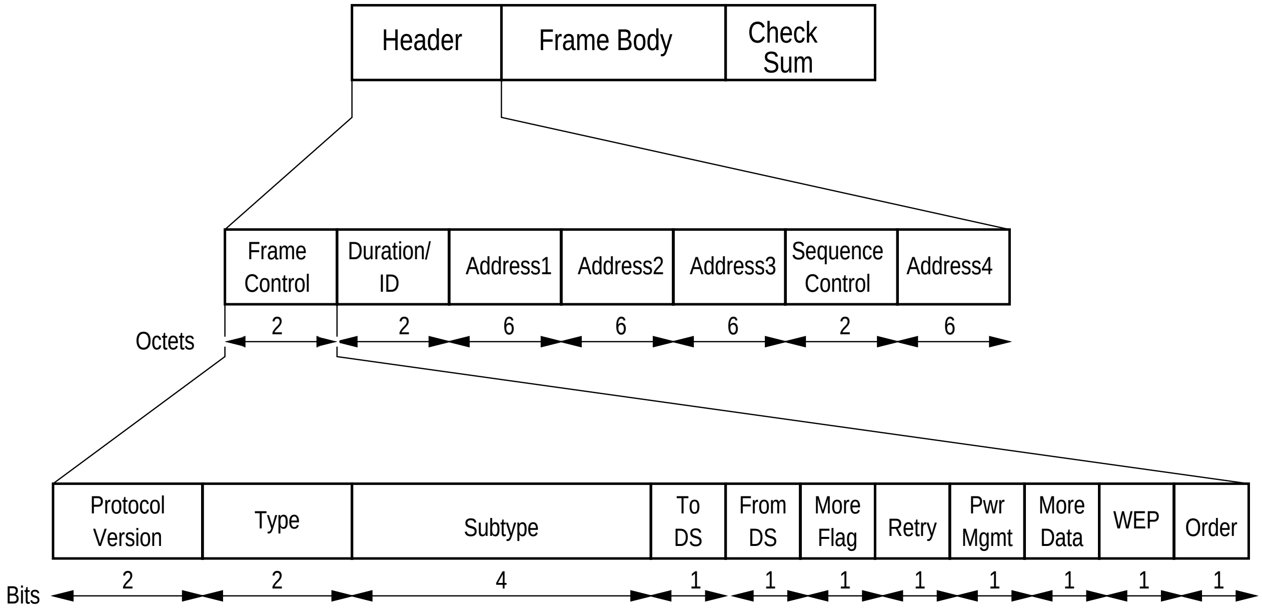

Figure 6: MAC frame structure

4 Networking Aspects

In 802.11 the notion of a *network* can roughly be equated with a BSS. In order to join a network, an STA must *synchronize* with the BSS. This implies operating on the same *channel*, running on synchronized *timers*, and using the same *BSS identifier* (BSSID). The channel is a logical concept provided to the MAC layer by the PHY layer. The physical manifestation of channels depends on the particular PHY layer.[1] The timers provide the time reference for the implementation of DCF and PCF. The BSSID logically identifies a BSS in a unique manner.

STAs obtain the information needed for synchronization by performing a *scanning procedure.*

4.1 Scanning

Scanning is the procedure via which an STA may probe the physical environment for the presence of existing BSSs. Scanning is used for finding and joining a network, for finding a new AP while roaming, or for initializing a new independent (ad hoc) BSS network. Scanning can be either *passive* or *active.*

4.1.1 Passive Scanning

With passive scanning, an STA tries to find a network simply by listening for beacons. Beacons include synchronization information such as time stamps, beacon interval, and SSID. The STA listens for beacons on one or several channels and at each channel scanned for at most the duration defined by the *channel time* parameter.

4.1.2 Active Scanning

With active scanning, an STA sends out a *probe* request frame and waits for the probe responses, essentially a solicited beacon. Probe responses contain the same synchronization

[1] spread-spectrum PHY it can be roughly equated with the hopping sequence, in a direct-sequence spread-spectrum PHY it is essentially the carrier frequency.

information as beacons, namely time stamps, beacon interval, and SSID. For each channel, the STA waits the requested *probe delay* time, then performs the basic medium access procedure, and sends the probe frame with the SSID and BSSID set to the broadcast address. It then starts the probe timer and waits for the responses. If no responses arrive before the *minimum channel time* has expired, it continues with the next channel, otherwise it waits until the *maximum channel time* has expired and then processes all the received probe responses.

In an infrastructure network, the probe response is generated by the AP. In an ad hoc network, the probe response is generated by the STA that transmitted the most recent beacon. This STA must remain in the *awake state* (see Section 6 on power management below) in order to answer probe requests until another STA in the same IBSS issued a beacon, and thus takes over.

4.2 Timing Synchronization

DCF and PCF in a BSS rely on the timing synchronization of the participating STAs. Therefore STAs support a *timing synchronization function* (TSF) and maintain a corresponding timer. Synchronization is achieved by using the information contained in periodically transmitted beacon frames. The beacon frame includes a *time stamp* and the beacon period. The time stamp recorded in the frame is adjusted by the TSF in the transmitter to account for propagation delays through the local physical layer. Similarly, the time stamp is modified at the receiving side to account for the propagation delays through the local physical layer up to the MAC prior to being used for adjusting the local TSF timer.

4.3 Addressing

STAs on a network are identified by a 48-bit address with a format as defined for IEEE 802 MAC addresses. Several types of addresses are identified:

Individual address designates one particular STA on a network.

Multicast-group address designates a collection of logically related stations (as defined by a higher-level entity).

Broadcast address is the address containing all 1's designating the set of all STAs actively connected to a medium.

Multicast-group and broadcast addresses are sometimes collectively referred to as *group addresses.*

Alternately, addresses are subdivided according to what they target:

BSSID A BSS is uniquely identified by the BSSID, a 48-bit field structured in the same way as an IEEE 802 MAC address. The value of the BSSID depends on the nature of the BSS; for an infrastructure network, the BSSID corresponds to the MAC address of the AP, for an ad hoc network it is formed of a 46-bit random number with the individual/group bit set to 1 and the universal/local bit set to 0. The broadcast BSSID can only be used with probe requests during active scanning.

SA The source address identifies the MAC entity that initiated the transfer of a frame.

DA The destination address identifies the MAC entity that is the final recipient of a frame.

TA The transmitter address identifies the STA that transmitted the frame onto the wireless medium.

RA The receiver address identifies the STA that is intended to receive the frame from the wireless medium.

Note that in the context of SA and DA, we speak of MAC entities, because these devices need not be physically connected to the wireless medium, whereas TA and RA always identify STAs participating in the same BSS.

5 Security

The security of wireless communication is adversely affected by the undefined boundaries of the transmission medium. In contrast to wired communication where signal propagation is physically constrained to a cable, wireless propagation is difficult to predict exactly; moreover, it varies. As a consequence, without further measures, any device in the actual range of the transmitting device could in principle eavesdrop an ongoing communication. 802.11 defines security services with the aim of raising the level of security at least to the physical security provided by wired links.

5.1 Authentication

802.11 provides link-level authentication by which stations establish their identity to stations with which they will communicate. A mutual authentication process establishes a trust relationship between two devices, be it in an infrastructure or an ad hoc network.

802.11 distinguishes between two subtypes of authentication, *open system* and *shared key.* Open system authentication in effect is a null authentication algorithm. By this, two STAs can agree or disagree to base their mutual trust relationship on zero evidence. The shared key authentication is based on a *challenge-response handshake protocol* requiring four frame exchanges between the requester and responder. Typically for these schemes, it is assumed that the shared secret key has been distributed to the authenticating STAs via a communication channel independent of 802.11. The shared secret is not transmitted across the air link. Encryption and decryption of the challenge and response, respectively, is based on the *wired equivalent privacy* algorithm.

5.2 Wired Equivalent Privacy

802.11 defines *wired equivalent privacy* (WEP) with the aim of protecting the communication between authenticated STAs from casual eavesdropping. Privacy is achieved by encrypting the contents of frames exchanged across the physical link. The WEP algorithm meets these design points: The cryptography is reasonably strong with key lengths meeting the export regulations of the U.S. Department of Commerce. It is self-sychronizing at the frame level, which is important in an environment where the frame loss rate can be high.

WEP combines an *integrity check* scheme and an *encryption* scheme. The integrity check is based on cyclic redundancy check using a generator polynomial of degree 32 (CRC-32), which produces the *integrity check vector* (ICV). The ICV is appended to the original MPDU for transmission and will also be encrypted. Encryption is based on the XOR combination of the key sequence and the plaintext sequence to produce the cipher text sequence. The key sequence is generated based on the *secret key* (which must be shared between the communicating STAs) and an *initialization vector* (IV). The initialization can be changed as frequently as every MPDU; it should be changed at a reasonable rate to prolong the useful lifetime of the secret key. The IV is transmitted with every MPDU and thus provides self-synchronous operation.

The data frame to be transmitted thus consists of the triplet {IV, frame body, ICV}, where both frame body and ICV have been encrypted. The IV field is a 4-octet field consisting of 24 bits of actual IV, 2 bits for key ID and 6 pad bits.

802.11 supports multiple key sets with multple (up to 4) keys in each set. The default key set is shared by all STAs and enables WEP for MPDUs with group addresses. Besides the default key set, key sets valid for a particular address pair TA/RA are also supported. Whenever such an address-specific key set exists, WEP uses keys from this set in the encryption process.

6 Power Management

Wireless LANs are typically used in combination with mobile computing devices such as laptops, palmtops or PDAs. As these devices are mostly battery powered battery lifetime is a concern. 802.11 addresses this issue by defining two different power states: in *awake state* the station is fully powered, in *doze state* the station is not able to transmit or receive and power consumption is at a minimum. Depending on whether an STA is in *active* or *power-save mode*, it is in the awake state or sometimes in doze state, thus preserving battery power.

In power-save mode, an STA needs to know the time intervals when it has to be in the awake state or the doze state. The way this is accomplished depends on the type of network, namely infrastructure or ad hoc.

6.1 Power Management in Infrastructure Networks

In infrastructure networks the AP is in charge of managing STAs in power save mode (PS STA). The AP knows about the mode of each STA in the BSS. When the AP sees traffic destined to PS STAs, it temporarily buffers these frames and records in the *traffic indication map* (TIM) the fact that there is pending traffic for this PS STA. The TIM is included in the beacons sent out by the AP.

PS STAs periodically switch every *listen interval* from doze state to awake state in order to receive beacons; the listen interval is chosen by the STA in multiples of the beacon interval and communicated to the AP during (re-)association. The PS STA receive and interpret the TIM and thus know about pending traffic waiting to be transmitted at the AP. If the TIM indicates pending traffic, the STA sends a short power-save poll frame to the AP, which responds by sending the pending MSDU to the STA. If the PS STA receives the TIM during a contention-free period, it does not send a poll frame but remains in the awake state until it has received the buffered MSDU.

Multicast and broadcast MSDUs are handled slightly differently. The AP transmits them immediately after the transmission of a *delivery TIM* (DTIM). Every TIM has a DTIM count field indicating the number of TIM periods until a DTIM occurs. The DTIM is simply a TIM whose count has a value of 0. The DTIM period is set by the AP. Thus, a PS STA listening to the TIM at its own listen interval frequency can always determine when the next DTIM occurs and can thus switch to the awake state for receipt of the broadcast/multicast MSDUs.

6.2 Power Management in Ad hoc Networks

With ad hoc networks there is no AP that can temporarily buffer traffic destined to a PS STA or knows about the STA's power state. As a consequence, STAs must themselves try to estimate the power-saving state of other STAs and announce messages destined for PS STAs by transmitting *ad hoc traffic information messages* (ATIM) during a defined period of time called the *ATIM window*.

802.11 does not specify the means by which STAs in an ad hoc network estimate the power state of other STAs. Possible passive means include "sniffing" into MAC frames from STAs to read their power management subfield in the frame control field or using statistics of failed transmissions. Possible active means include probing via the transmission of an RTS frame and assuming an STA is in power-save mode if a CTS from it is not received.

The ATIM window is a time period following a TBTT. The STA initiating the ad hoc network determines the ATIM window length, which must stay fixed for the duration of the ad hoc network. A length of zero implies that no power management is used. The ATIM length is transmitted with each beacon as part of the IBSS parameter set in the beacon frame. STAs joining the ad hoc network adopt this length.

During the ATIM window all STAs are in awake state. An STA that has an MSDU for another STA announces this MSDU during the ATIM window in an ATIM directed to that STA. The announcing STA may only send the respective MSDU if it has received a positive acknowledgment. An exception are ATIM frames to group addresses (broadcast or multicast), which must not be acknowledged. If a STA receives a directed ATIM frame during the window, it must acknowledge the ATIM frame and stay awake for the duration of the beacon interval in order to receive the announced MSDUs. If an STA does not receive an ATIM frame during the ATIM window, it may enter the doze state again.

7 Physical Transmission System

The 802.11 protocol architecture allows the MAC layer protocols to operate on top of different physical transmission systems. Currently, the following transmissions systems are specified: *frequency-hopping spread spectrum* (FHSS), *direct-sequence spread spectrum* (DSSS), and *infrared pulse-position* (IRPM).

7.1 Frequency Hopping Spread Spectrum

The FHSS provides data rates of 1 and 2 Mbits/s using *Gaussian frequency shift keying* (GFSK) at different modulation indices. Frequency spreading is achieved by hopping across 79 center frequencies (U.S. and Europe) and 23 center frequencies (Japan), respectively, in a periodic pseudo-random fashion.

7.2 Direct Sequence Spread Spectrum

The DSSS provides data rates of 1 and 2 Mbit/s using *differential binary phase shift keying* (DBPSK) and *differential quaternary phase shift keying* (DQPSK), respectively. Using the 11-chip barker code {+1, -1, +1, +1, -1, +1, +1, +1, -1, -1, -1} the system achieves a spreading factor of 11 with an effective bandwidth of 11 MHz. The DSSS operates in the 2.4 GHz ISM band across a total of 14 channels separated (with one exception) by 5 MHz between 2.412 and 2.484 GHz. The actual channels used depend on region- or country-specific regulatory rules: under the U.S. FCC and Canadian IC rules, 11 channels are used, under European ETSI rules (except in France and Spain), 13 channels are used, and under Japan's MPT rules, a single channel is used. In multiple cell network topologies, overlapping or adjacent cells using different channels can operate concurrently if their respective center frequencies are separated by at least 30 MHz.

References

[1] IEEE Standard 802.11-1997, *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications*, Institute of Electrical and Electronics Engineers, New York, November, 1997.

Acronyms

AP	Access point
ATIM	Announcement or Ad Hoc traffic indication map
BSS	Basic service set
BSSID	Basic service set identifier
CFP	Contention-free period
CP	Contention period
CRC	Cyclic redundancy check
CSMA	Carrier sensing multiple access
CTS	Clear to send
DA	Destination address
DBPSK	Differential binary phase-shift keying
DCF	Distributed control function
DQPSK	Differential quaternary phase shift keying
DSS	Distribution system service
DSSS	Direct sequence spread spectrum
DTIM	Delivery traffic indication map
ESS	Extended service set
ETSI	European Telecommunications Standards Institute
FCC	Federal Communications Commission
FHSS	Frequency hopping spread spectrum
GFSK	Gaussian minimum-shift keying

IBSS	Independent basic service set
IC	Industry Canada
ICV	Initial chaining vector
IEEE	Institute of Electrical and Electronics Engineers
IFS	Interframe space
ISM	Industrial, scientific, and medical
LAN	Local area network
MAC	Media access control
MAN	Metropolitan area network
MPT	Ministry of Post and Telecommunication
MSDU	MAC layer service data unit
NAV	Network allocation vector
OSI	Open systems interconnection
PC	Point control
PCF	Point control function
PDA	Personal digital assistant
PHY	Physical layer
RA	Receiver address
RTS	Request to send
SA	Source address
SS	Station service
SSID	Service set identifier
STA	Station
TA	Transmitter address
TBTT	Target beacon transmission time
TIM	Traffic indication map
TSF	Timing synchronization function
WEP	Wired equivalent privacy