CHAPTER 14

IEEE 802.11 Wireless LAN Standard

The most prominent specification for wireless LANs was developed by the IEEE 802.11 working group. We look first at the overall architecture of IEEE 802 standards and then at the specifics of IEEE 802.11.

14.1 IEEE 802 ARCHITECTURE

The architecture of a LAN is best described in terms of a layering of protocols that organize the basic functions of a LAN. This section opens with a description of the standardized protocol architecture for LANs, which encompasses physical, medium access control, and logical link control layers. We then look in more detail at medium access control and logical link control.

Protocol Architecture

Protocols defined specifically for LAN and MAN (metropolitan area network) transmission address issues relating to the transmission of blocks of data over the network. In OSI terms, higher-layer protocols (layer 3 or 4 and above) are independent of network architecture and are applicable to LANs, MANs, and WANs. Thus, a discussion of LAN protocols is concerned principally with lower layers of the OSI model.

Figure 14.1 relates the LAN protocols to the OSI architecture (Figure 4.3). This architecture was developed by the IEEE 802 committee and has been adopted by all organizations working on the specification of LAN standards. It is generally referred to as the IEEE 802 reference model.[1]

Working from the bottom up, the lowest layer of the IEEE 802 reference model corresponds to the **physical layer** of the OSI model and includes such functions as

- Encoding/decoding of signals
- Preamble generation/removal (for synchronization)
- Bit transmission/reception

In addition, the physical layer of the 802 model includes a specification of the transmission medium and the topology. Generally, this is considered "below" the lowest layer of the OSI model. However, the choice of transmission medium and topology is critical in LAN design, and so a specification of the medium is included.

Above the physical layer are the functions associated with providing service to LAN users. These include the following:

- On transmission, assemble data into a frame with address and error detection fields.
- On reception, disassemble frame, and perform address recognition and error detection.

[1] See Appendix A for a discussion of the IEEE 802 Standards Committee.

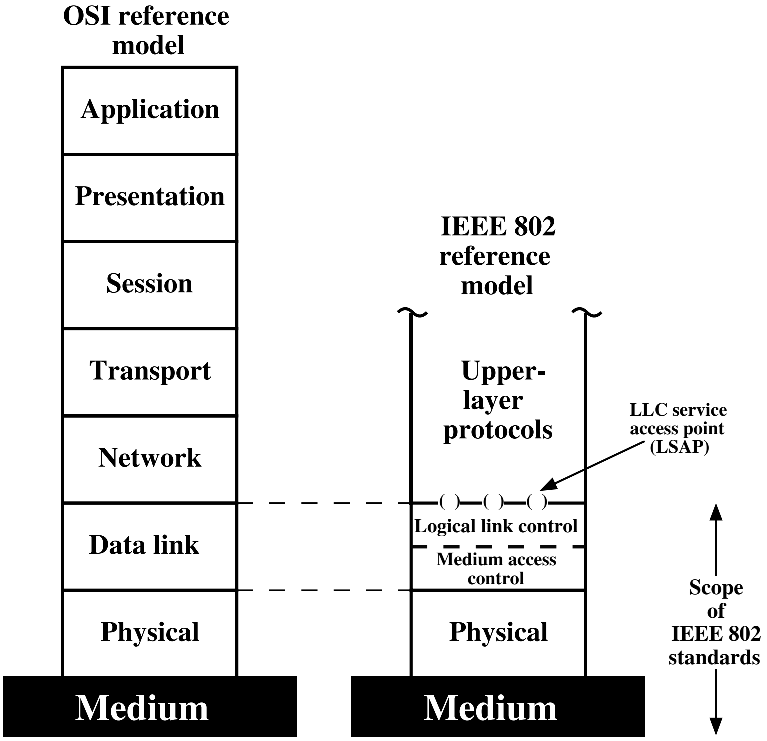

Figure 14.1 IEEE 802 Protocol Layers Compared to OSI Model

- Govern access to the LAN transmission medium.
- Provide an interface to higher layers and perform flow and error control.

These are functions typically associated with OSI layer 2. The set of functions in the last bullet item are grouped into a **logical link control (LLC)** layer. The functions in the first three bullet items are treated as a separate layer, called **medium access control (MAC)**. The separation is done for the following reasons:

- The logic required to manage access to a shared-access medium is not found in traditional layer 2 data link control.
- For the same LLC, several MAC options may be provided.

Figure 14.2, which reproduces Figure 11.16, illustrates the relationship between the levels of the architecture. Higher-level data are passed down to LLC, which appends control information as a header, creating an *LLC protocol data unit (PDU)*. This control information is used in the operation of the LLC protocol. The entire LLC PDU is then passed down to the MAC layer, which appends control information at the front and back of the packet, forming a *MAC frame*. Again, the control information in the frame is needed for the operation of the MAC protocol. For context, the figure also shows the use of TCP/IP and an application layer above the LAN protocols.

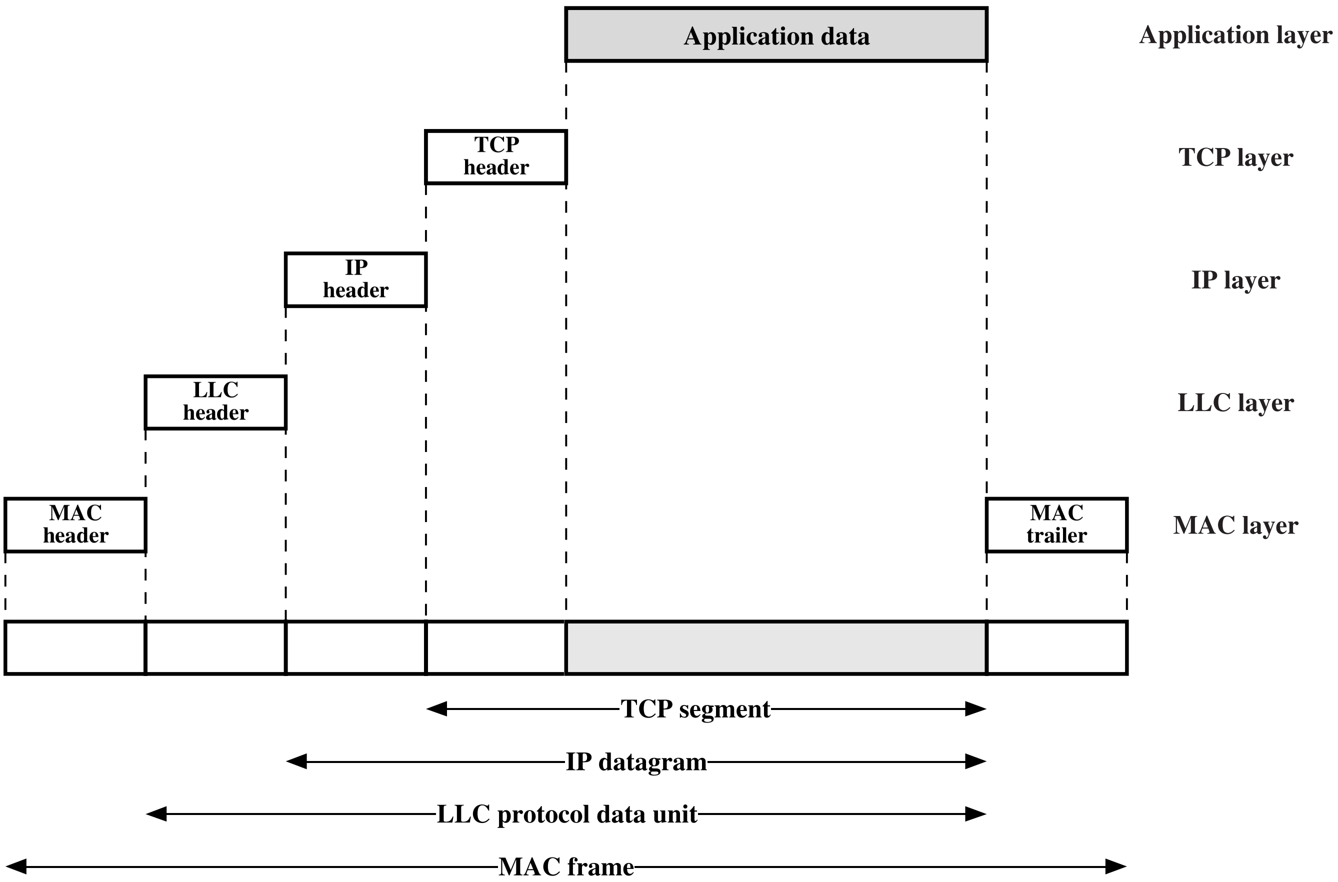

Figure 14.2 IEEE 802 Protocols in Context

MAC Frame Format

The MAC layer receives a block of data from the LLC layer and is responsible for performing functions related to medium access and for transmitting the data. As with other protocol layers, MAC implements these functions making use of a protocol data unit at its layer. In this case, the PDU is referred to as a MAC frame.

The exact format of the MAC frame differs somewhat for the various MAC protocols in use. In general, all of the MAC frames have a format similar to that of Figure 14.3. The fields of this frame are as follows:

- **MAC control:** This field contains any protocol control information needed for the functioning of the MAC protocol. For example, a priority level could be indicated here.
- **Destination MAC address:** The destination physical attachment point on the LAN for this frame.
- **Source MAC address:** The source physical attachment point on the LAN for this frame.
- **Data:** The body of the MAC frame. This may be LLC data from the next higher layer or control information relevant to the operatoin of the MAC protocol.
- **CRC:** The cyclic redundancy check field (also known as the frame check sequence, FCS, field). This is an error-detecting code, as we have seen in HDLC and other data link control protocols (Section 8.1).

In most data link control protocols, the data link protocol entity is responsible not only for detecting errors using the CRC but for recovering from those errors by retransmitting damaged frames. In the LAN protocol architecture, these two functions are split between the MAC and LLC layers. The MAC layer is responsible for detecting errors and discarding any frames that are in error. The LLC layer optionally keeps track of which frames have been successfully received and retransmits unsuccessful frames.

Logical Link Control

The LLC layer for LANs is similar in many respects to other link layers in common use. Like all link layers, LLC is concerned with the transmission of a link-level PDU between two stations, without the necessity of an intermediate switching node. LLC has two characteristics not shared by most other link control protocols:

1. It must support the multiaccess, shared-medium nature of the link (this differs from a multidrop line in that there is no primary node).
2. It is relieved of some details of link access by the MAC layer.

Addressing in LLC involves specifying the source and destination LLC users. Typically, a user is a higher-layer protocol or a network management function in the station. These LLC user addresses are referred to as service access points (SAPs), in keeping with OSI terminology for the user of a protocol layer.

We look first at the services that LLC provides to a higher-level user, and then at the LLC protocol.

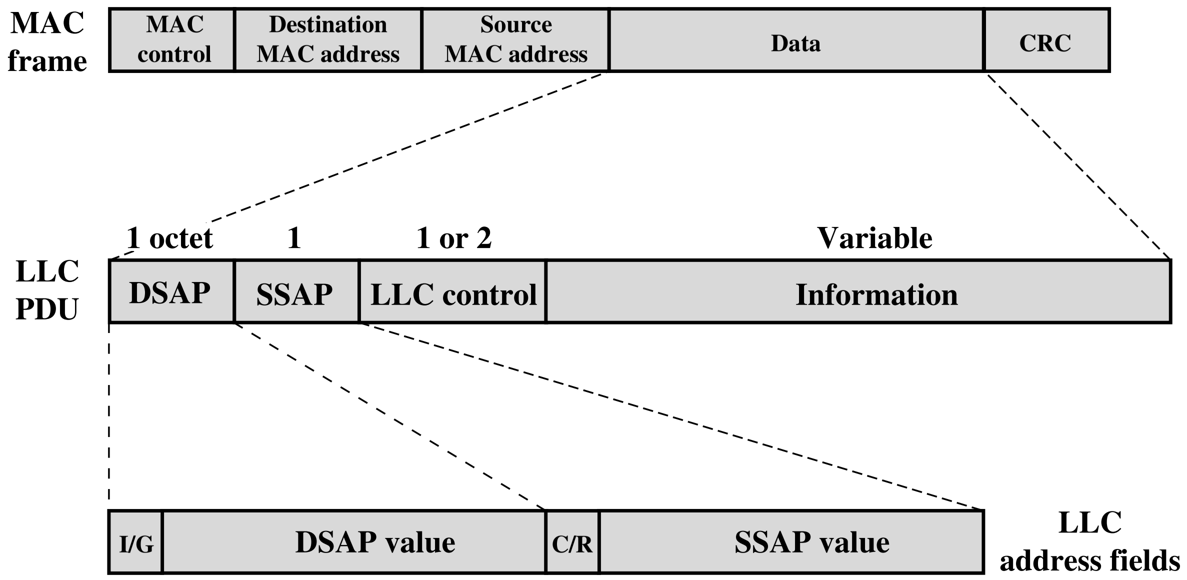

I/G = individual/group
C/R = command/response

DSAP = destination service access point
SSAP = source service access point

Figure 14.3 LLC PDU in a Generic MAC Frame Format

LLC Services

LLC specifies the mechanisms for addressing stations across the medium and for controlling the exchange of data between two users. The operation and format of this standard is based on HDLC. Three services are provided as alternatives for attached devices using LLC:

- **Unacknowledged connectionless service:** This service is a datagram-style service. It is a very simple service that does not involve any of the flow- and error-control mechanisms. Thus, the delivery of data is not guaranteed. However, in most devices, there will be some higher layer of software that deals with reliability issues.
- **Connection-mode service:** This service is similar to that offered by HDLC. A logical connection is set up between two users exchanging data, and flow control and error control are provided.
- **Acknowledged connectionless service:** This is a cross between the previous two services. It provides that datagrams are to be acknowledged, but no prior logical connection is set up.

Typically, a vendor will provide these services as options that the customer can select when purchasing the equipment. Alternatively, the customer can purchase equipment that provides two or all three services and select a specific service based on application.

The **unacknowledged connectionless service** requires minimum logic and is useful in two contexts. First, it will often be the case that higher layers of software will provide the necessary reliability and flow-control mechanism, and it is efficient to avoid duplicating them. For example, TCP could provide the mechanisms needed to ensure that data are delivered reliably. Second, there are instances in which the overhead of connection establishment and maintenance is unjustified or even counterproductive (for example, data collection activities that involve the periodic sampling of data sources, such as sensors and automatic self-test reports from security equipment or network components). In a monitoring application, the loss of an occasional data unit would not cause distress, as the next report should arrive shortly. Thus, in most cases, the unacknowledged connectionless service is the preferred option.

The **connection-mode service** could be used in very simple devices, such as terminal controllers, that have little software operating above this level. In these cases, it would provide the flow control and reliability mechanisms normally implemented at higher layers of the communications software.

The **acknowledged connectionless service** is useful in several contexts. With the connection-mode service, the logical link control software must maintain some sort of table for each active connection, to keep track of the status of that connection. If the user needs guaranteed delivery but there are a large number of destinations for data, then the connection-mode service may be impractical because of the large number of tables required. An example is a process control or automated factory environment where a central site may need to communicate with a large number of processors and programmable controllers. Another use of this is the handling of important and time-critical alarm or emergency control signals in a factory. Because of their importance, an acknowledgment is needed so that the sender can be assured that the signal got through. Because of the urgency of the signal, the user might not want to take the time first to establish a logical connection and then send the data.

LLC Protocol

The basic LLC protocol is modeled after HDLC and has similar functions and formats. The differences between the two protocols can be summarized as follows:

- LLC makes use of the asynchronous balanced mode of operation of HDLC, to support connection-mode LLC service; this is referred to as type 2 operation. The other HDLC modes are not employed.
- LLC supports an unacknowledged connectionless service using the unnumbered information PDU; this is known as type 1 operation.
- LLC supports an acknowledged connectionless service by using two new unnumbered PDUs; this is known as type 3 operation.
- LLC permits multiplexing by the use of LLC service access points (LSAPs).

All three LLC protocols employ the same PDU format (Figure 14.3), which consists of four fields. The DSAP and SSAP fields each contain a 7-bit address, which specify the destination and source users of LLC. One bit of the DSAP indicates whether the DSAP is an individual or group address. One bit of the SSAP indicates whether the PDU is a command or response PDU. The format of the LLC control field is identical to that of HDLC (Figure D.1, Appendix D), using extended (7-bit) sequence numbers.

For **type 1 operation**, which supports the unacknowledged connectionless service, the unnumbered information (UI) PDU is used to transfer user data. There is no acknowledgment, flow control, or error control. However, there is error detection and discard at the MAC level.

Two other PDU types, XID and TEST, are used to support management functions associated with all three types of operation. Both PDU types are used in the following fashion. An LLC entity may issue a command (C/R bit = 0) XID or TEST. The receiving LLC entity issues a corresponding XID or TEST in response. The XID PDU is used to exchange two types of information: types of operation supported and window size. The TEST PDU is used to conduct a loopback test of the transmission path between two LLC entities. Upon receipt of a TEST command PDU, the addressed LLC entity issues a TEST response PDU as soon as possible.

With **type 2 operation**, a data link connection is established between two LLC SAPs prior to data exchange. Connection establishment is attempted by the type 2 protocol in response to a request from a user. The LLC entity issues a SABME PDU[2] to request a logical connection with the other LLC entity. If the connection is accepted by the LLC user designated by the DSAP, then the destination LLC entity returns an unnumbered acknowledgment (UA) PDU. The connection is henceforth

[2] This stands for Set Asynchronous Balanced Mode Extended. It is used in HDLC to choose ABM and to select extended sequence numbers of 7 bits. Both ABM and 7-bit sequence numbers are mandatory in type 2 operation.

uniquely identified by the pair of user SAPs. If the destination LLC user rejects the connection request, its LLC entity returns a disconnected mode (DM) PDU.

Once the connection is established, data is exchanged using information PDUs, as in HDLC. The information PDUs include send and receive sequence numbers, for sequencing and flow control. The supervisory PDUs are used, as in HDLC, for flow control and error control. Either LLC entity can terminate a logical LLC connection by issuing a disconnect (DISC) PDU.

With **type 3 operation**, each transmitted PDU is acknowledged. A new (not found in HDLC) unnumbered PDU, the acknowledged connectionless (AC) information PDU, is defined. User data are sent in AC command PDUs and must be acknowledged using an AC response PDU. To guard against lost PDUs, a 1-bit sequence number is used. The sender alternates the use of 0 and 1 in its AC command PDU, and the receiver responds with an AC PDU with the opposite number of the corresponding command. Only one PDU in each direction may be outstanding at any time.

14.2 IEEE 802.11 ARCHITECTURE AND SERVICES

Work on wireless LANs within the IEEE 802 committee began in 1987 within the IEEE 802.4 group. The initial interest was in developing an ISM-based wireless LAN using the equivalent of a token-passing bus MAC protocol. After some work, it was decided that token bus was not suitable for controlling a radio medium without causing inefficient use of the radio frequency spectrum. IEEE 802 then decided in 1990 to form a new working group, IEEE 802.11, specifically devoted to wireless LANs, with a charter to develop a MAC protocol and physical medium specification. Table 14.1 briefly defines key terms used in the IEEE 802.11 standard.

Table 14.1 IEEE 802.11 Terminology

Access point (AP)	Any entity that has station functionality and provides access to the distribution system via the wireless medium for associated stations
Basic service set (BSS)	A set of stations controlled by a single coordination function.
Coordination function	The logical function that determines when a station operating within a BSS is permitted to transmit and may be able to receive PDUs.
Distribution System (DS)	A system used to interconnect a set of BSSs and integrated LANs to create an ESS.
Extended service set (ESS)	A set of one or more interconnected BSSs and integrated LANs that appear as a single BSS to the LLC layer at any station associated with one of these BSSs.
MAC protocol data unit (MPDU)	The unit of data exchanged between two peer MAC entites using the services of the physical layer.
MAC service data unit (MSDU)	Information that is delivered as a unit between MAC users.
Station	Any device that contains an IEEE 802.11 conformant MAC and physical layer.

IEEE 802.11 Architecture

Figure 14.4 illustrates the model developed by the 802.11 working group. The smallest building block of a wireless LAN is a basic service set (BSS), which consists of some number of stations executing the same MAC protocol and competing for access to the same shared wireless medium. A BSS may be isolated or it may connect to a backbone distribution system (DS) through an access point (AP). The access point functions as a bridge. The MAC protocol may be fully distributed or controlled by a central coordination function housed in the access point. The BSS generally corresponds to what is referred to as a cell in the literature. The DS can be a switch, a wired network, or a wireless network.

The simplest configuration is shown in Figure 14.4, in which each station belongs to a single BSS; that is, each station is within wireless range only of other stations within the same BSS. It is also possible for two BSSs to overlap geographically, so that a single station could pariticipate in more than one BSS. Further, the association between a station and a BSS is dynamic. Stations may turn off, come within range, and go out of range.

An extended service set (ESS) consists of two or more basic service sets interconnected by a distribution system. Typically, the distribution system is a wired backbone LAN but can be any communications network. The extended service set appears as a single logical LAN to the logical link control (LLC) level.

Figure 14.4 indicates that an access point (AP) is implemented as part of a station; the AP is the logic within a station that provides access to the DS by providing DS services in addition to acting as a station. To integrate the IEEE 802.11 architecture with a traditional wired LAN, a portal is used. The portal logic is imple-

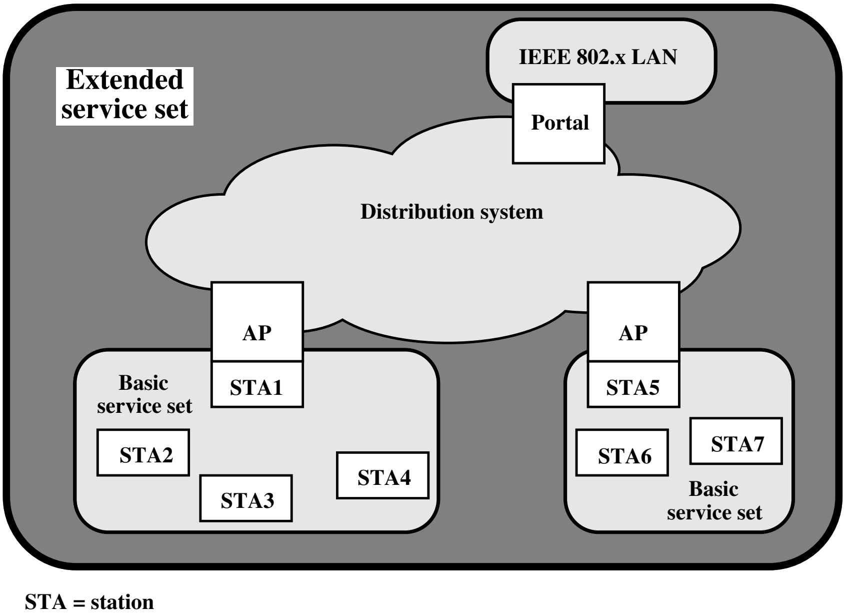

STA = station

Figure 14.4 IEEE 802.11 Architecture

mented in a device, such as a bridge or router, that is part of the wired LAN and that is attached to the DS.

IEEE 802.11 Services

IEEE 802.11 defines nine services that need to be provided by the wireless LAN to provide functionality equivalent to that which is inherent to wired LANs. Table 14.2 lists the services and indicates two ways of categorizing them.

1. The service provider can be either the station or the distribution system (DS). Station services are implemented in every 802.11 station, including access point (AP) stations. Distribition services are provided between basic service sets (BSSs); these services may be implemented in an AP or in another special-purpose device attaced to the distribution system.
2. Three of the services are used to control IEEE 802.11 LAN access and confidentiality. Six of the services are used to support delivery of MAC service data units (MSDUs) between stations. The MSDU is a the block of data passed down from the MAC user to the MAC layer; typically this is a LLC PDU. If the MSDU is too large to be transmitted in a single MAC frame, it may be fragmented and transmitted in a series of MAC frames. Fragmentation is discussed in Section 14.3.

Following the IEEE 802.11 document, we next discuss the services in an order designed to clarify the operation of an IEEE 802.11 ESS network. **MSDU delivery**, which is the basic service, has already been mentioned.

Distribution of Messages Within a DS

The two services involved with the distribution of messages within a DS are distribution and integration. **Distribution** is the primary service used by stations to exchange MAC frames when the frame must traverse the DS to get from a station in one BSS to a station in another BSS. For example, suppose a frame is to be sent

Table 14.2 IEEE 802.11 Services

Service	Provider	Used to support
Association	Distribution system	MSDU delivery
Authentication	Station	LAN access and security
Deauthentication	Station	LAN access and security
Dissassociation	Distribution system	MSDU delivery
Distribution	Distribution system	MSDU delivery
Integration	Distribution system	MSDU delivery
MSDU delivery	Station	MSDU delivery
Privacy	Station	LAN access and security
Reassociation	Distribution system	MSDU delivery

from station 2 (STA 2) to STA 7 in Figure 14.4. The frame is sent from STA 2 to STA 1, which is the AP for this BSS. The AP gives the frame to the DS, which has the job of directing the frame to the AP associated with STA 5 in the target BSS. STA 5 receives the frame and forwards it to STA 7. How the message is transported through the DS is beyond the scope of the IEEE 802.11 standard.

If the two stations that are communicating are within the same BSS, then the distribution service logically goes through the single AP of that BSS.

The **integration** service enables transfer of data between a station on an IEEE 802.11 LAN and a station on an integrated IEEE 802.x LAN. The term *integrated* refers to a wired LAN that is physically connected to the DS and whose stations may be logically connected to an IEEE 802.11 LAN via the integration service. The integration service takes care of any address translation and media conversion logic required for the exchange of data.

Association-Related Services

The primary purpose of the MAC layer is to transfer MSDUs between MAC entities; this purpose is fulfilled by the distribution service. For that service to function, it requires information about stations within the ESS that is provided by the association-related services. Before the distribution service can deliver data to or accept data from a station, that station must be *associated*. Before looking at the concept of association, we need to describe the concept of mobility. The standard defines three transition types of based on mobility:

- **No transition:** A station of this type is either stationary or moves only within the direct communication range of the communicating stations of a single BSS.
- **BSS transition:** This is defined as a station movement from one BSS to another BSS within the same ESS. In this case, delivery of data to the station requires that the addressing capability be able to recognize the new location of the station.
- **ESS transition:** This is defined as a station movement from a BSS in one ESS to a BSS within another ESS. This case is supported only in the sense that the station can move. Maintenance of upper-layer connections supported by 802.11 cannot be guaranteed. In fact, disruption of service is likely to occur.

To deliver a message within a DS, the distribution service needs to know where the destination station is located. Specifically, the DS needs to know the identity of the AP to which the message should be delivered in order for that message to reach the destination station. To meet this requirement, a station must maintain an association with the AP within its current BSS. Three services relate to this requirement:

- **Association:** Establishes an initial association between a station and an AP. Before a station can transmit or receive frames on a wireless LAN, its identity and address must be known. For this purpose, a station must establish an association with an AP within a particular BSS. The AP can then communicate this information to other APs within the ESS to facilitate routing and delivery of addressed frames.

- **Reassociation:** Enables an established association to be transferred from one AP to another, allowing a mobile station to move from one BSS to another.
- **Disassociation:** A notification from either a station or an AP that an existing association is terminated. A station should give this notification before leaving an ESS or shutting down. However, the MAC management facility protects itself against stations that disappear without notification.

Access and Privacy Services

There are two characteristics of a wired LAN that are not inherent in a wireless LAN.

1. In order to transmit over a wired LAN, a station must be physically connected to the LAN. On the other hand, with a wireless LAN, any station within radio range of the other devices on the LAN can transmit. In a sense, there is a form of authentication with a wired LAN, in that it requires some positive and presumably observable action to connect a station to a wired LAN.
2. Similarly, in order to receive a transmission from a station that is part of a wired LAN, the receiving station must also be attached to the wired LAN. On the other hand, with a wireless LAN, any station within radio range can receive. Thus, a wired LAN provides a degree of privacy, limiting reception of data to stations connected to the LAN.

IEEE 802.11 defines three services that provide a wireless LAN with these two features:

- **Authentication:** Used to establish the identity of stations to each other. In a wired LAN, it is generally assumed that access to a physical connection conveys authority to connect to the LAN. This is not a valid assumption for a wireless LAN, in which connectivity is achieved simply by having an attached antenna that is properly tuned. The authentication service is used by stations to establish their identity with stations they wish to communicate with. IEEE 802.11 supports several authentication schemes and allows for expansion of the functionality of these schemes. The standard does not mandate any particular authentication scheme, which could range from relatively unsecure handshaking to public-key encryption schemes. However, IEEE 802.11 requires mutually acceptable, successful authentication before a station can establish an association with an AP.
- **Deathentication:** This service is invoked whenever an existing authentication is to be terminated.
- **Privacy:** Used to prevent the contents of messages from being read by other than the intended recipient. The standard provides for the optional use of encryption to assure privacy. The algorithm specified in the standard is WEP, which is described in Section 14.3.

14.3 IEEE 802.11 MEDIUM ACCESS CONTROL

The IEEE 802.11 MAC layer covers three functional areas: reliable data delivery, access control, and security. We look at each of these in turn.

Reliable Data Delivery

As with any wireless network, a wireless LAN using the IEEE 802.11 physical and MAC layers is subject to considerable unreliability. Noise, interference, and other propagation effects result in the loss of a significant number of frames. Even with error-correction codes, a number of MAC frames may not successfully be received. This situation can be dealt with by reliability mechanisms at a higher layer, such as TCP. However, timers used for retransmission at higher layers are typically on the order of seconds. It is therefore more efficient to deal with errors at the MAC level. For this purpose, IEEE 802.11 includes a frame exchange protocol. When a station receives a data frame from another station it returns an acknowledgment (ACK) frame to the source station. This exchange is treated as an atomic unit, not to be interrupted by a transmission from any other station. If the source does not receive an ACK within a short period of time, either because its data frame was damaged or because the returning ACK was damaged, the source retransmits the frame.

Thus, the basic data transfer mechanism in IEEE 802.11 involves an exchange of two frames. To further enhance reliability, a four-frame exchange may be used. In this scheme, a source first issues a request to send (RTS) frame to the destination. The destination then responds with a clear to send (CTS). After receiving the CTS, the source transmits the data frame, and the destination responds with an ACK. The RTS alerts all stations that are within reception range of the source that an exchange is under way; these stations refrain from transmission in order to avoid a collision between two frames transmitted at the same time. Similarly, the CTS alerts all stations that are within reception range of the destination that an exchange is under way. The RTS/CTS portion of the exchange is a required function of the MAC but may be disabled.

Access Control

The 802.11 working group considered two types of proposals for a MAC algorithm: distributed access protocols, which, like Ethernet, distribute the decision to transmit over all the nodes using a carrier-sense mechanism; and centralized access protocols, which involve regulation of transmission by a centralized decision maker. A distributed access protocol makes sense for an ad hoc network of peer workstations and may also be attractive in other wireless LAN configurations that consist primarily of bursty traffic. A centralized access protocol is natural for configurations in which a number of wireless stations are interconnected with each other and some sort of base station that attaches to a backbone wired LAN; it is especially useful if some of the data is time sensitive or high priority.

The end result for 802.11 is a MAC algorithm called DFWMAC (distributed foundation wireless MAC) that provides a distributed access control mechanism with an optional centralized control built on top of that. Figure 14.5 illustrates the

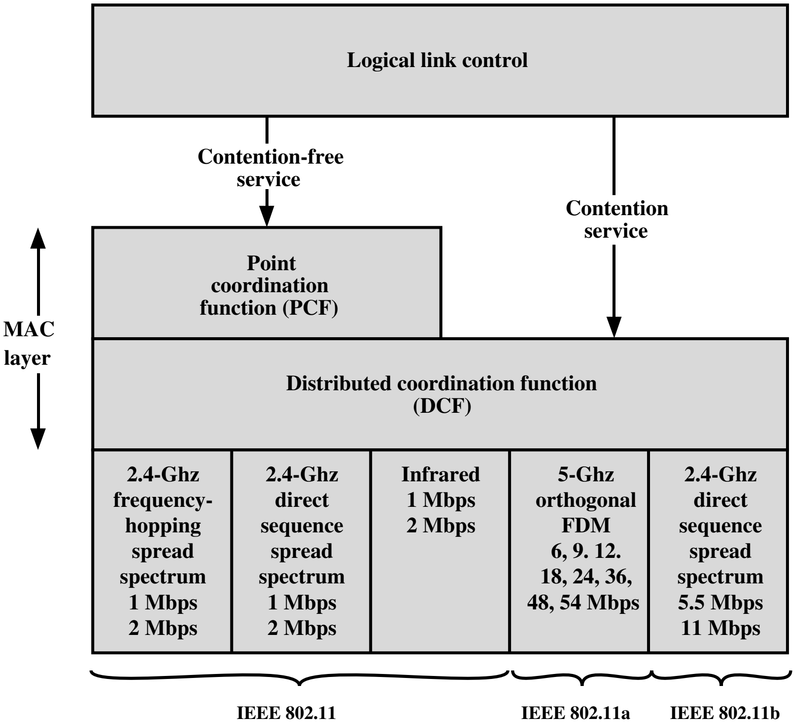
Logical link control
Contention-free service
Contention service
MAC layer
Point coordination function (PCF)
Distributed coordination function (DCF)
2.4-Ghz frequency-hopping spread spectrum 1 Mbps 2 Mbps
2.4-Ghz direct sequence spread spectrum 1 Mbps 2 Mbps
Infrared 1 Mbps 2 Mbps
5-Ghz orthogonal FDM 6, 9. 12. 18, 24, 36, 48, 54 Mbps
2.4-Ghz direct sequence spread spectrum 5.5 Mbps 11 Mbps
IEEE 802.11
IEEE 802.11a
IEEE 802.11b

Figure 14.5 IEEE 802.11 Protocol Architecture

architecture. The lower sublayer of the MAC layer is the distributed coordination function (DCF). DCF uses a contention algorithm to provide access to all traffic. Ordinary asynchronous traffic directly uses DCF. The point coordination function (PCF) is a centralized MAC algorithm used to provide contention-free service. PCF is built on top of DCF and exploits features of DCF to assure access for its users. Let us consider these two sublayers in turn.

Distributed Coordination Function

The DCF sublayer makes use of a simple CSMA (carrier sense multiple access) algorithm. If a station has a MAC frame to transmit, it listens to the medium. If the medium is idle, the station may transmit; otherwise the station must wait until the current transmission is complete before transmitting. The DCF does not include a collision detection function (i.e., CSMA/CD) because collision detection is not practical on a wireless network. The dynamic range of the signals on the medium is very large, so that a transmitting station cannot effectively distinguish incoming weak signals from noise and the effects of its own transmission.

To ensure the smooth and fair functioning of this algorithm, DCF includes a set of delays that amounts to a priority scheme. Let us start by considering a single delay known as an interframe space (IFS). In fact, there are three different IFS values, but the algorithm is best explained by initially ignoring this detail. Using an IFS, the rules for CSMA access are as follows (Figure 14.6):

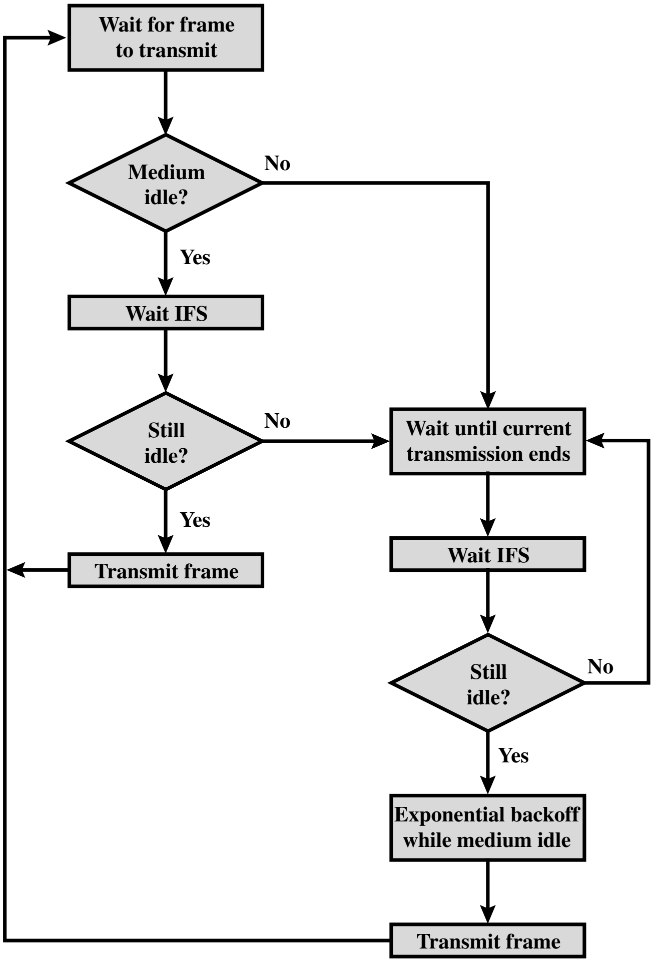
Wait for frame to transmit
Medium idle?
No
Yes
Wait IFS
Still idle?
No
Yes
Transmit frame
Wait until current transmission ends
Wait IFS
Still idle?
No
Yes
Exponential backoff while medium idle
Transmit frame

Figure 14.6 IEEE 802.11 Medium Access Control Logic

1. A station with a frame to transmit senses the medium. If the medium is idle, it waits to see if the medium remains idle for a time equal to IFS. If so, the station may transmit immediately.
2. If the medium is busy (either because the station initially finds the medium busy or because the medium becomes busy during the IFS idle time), the station defers transmission and continues to monitor the medium until the current transmission is over.
3. Once the current transmission is over, the station delays another IFS. If the medium remains idle for this period, then the station backs off a random amount of time and again senses the medium. If the medium is still idle, the station may transmit. During the backoff time, if the medium becomes busy, the backoff timer is halted and resumes when the medium becomes idle.

To ensure that backoff maintains stability, a technique known as binary exponential backoff is used. A station will attempt to transmit repeatedly in the face of repeated collisions, but after each collision, the mean value of the random delay is doubled. The binary exponential backoff provides a means of handling a heavy load. Repeated failed attempts to transmit result in longer and longer backoff times, which helps to smooth out the load. Without such a backoff, the following situation could occur. Two or more stations attempt to transmit at the same time, causing a collision. These stations then immediately attempt to retransmit, causing a new collision.

The preceding scheme is refined for DCF to provide priority-based access by the simple expedient of using three values for IFS:

- **SIFS (short IFS):** The shortest IFS, used for all immediate response actions, as explained in the following discussion
- **PIFS (point coordination function IFS):** A midlength IFS, used by the centralized controller in the PCF scheme when issuing polls
- **DIFS (distributed coordination function IFS):** The longest IFS, used as a minimum delay for asynchronous frames contending for access

Figure 14.7a illustrates the use of these time values. Consider first the SIFS. Any station using SIFS to determine transmission opportunity has, in effect, the highest

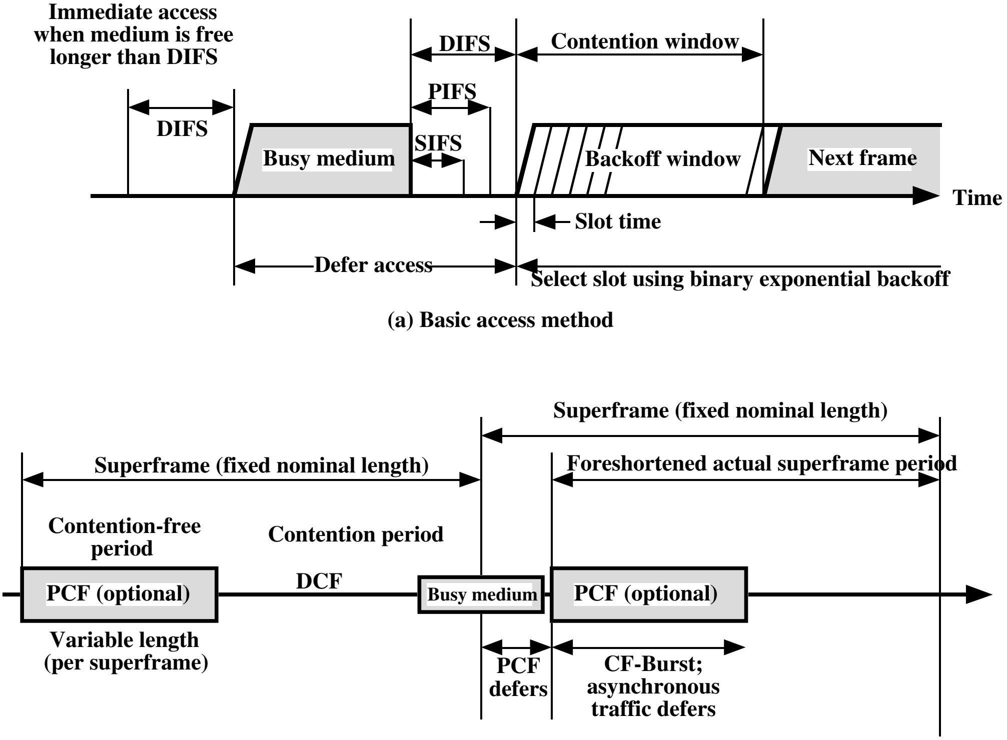
Immediate access when medium is free longer than DIFS
DIFS
DIFS
Contention window
PIFS
SIFS
Busy medium
Backoff window
Next frame
Time
Slot time
Defer access
Select slot using binary exponential backoff
Superframe (fixed nominal length)
Superframe (fixed nominal length)
Foreshortened actual superframe period
Contention-free period
Contention period
DCF
PCF (optional)
Busy medium
PCF (optional)
Variable length (per superframe)
PCF defers
CF-Burst; asynchronous traffic defers

(a) Basic access method

(b) PCF superframe construction

Figure 14.7 IEEE 802.11 MAC Timing

priority, because it will always gain access in preference to a station waiting an amount of time equal to PIFS or DIFS. The SIFS is used in the following circumstances:

- **Acknowledgment (ACK):** When a station receives a frame addressed only to itself (not multicast or broadcast) it responds with an ACK frame after waiting only for an SIFS gap. This has two desirable effects. First, because collision detection is not used, the likelihood of collisions is greater than with CSMA/CD, and the MAC-level ACK provides for efficient collision recovery. Second, the SIFS can be used to provide efficient delivery of an LLC protocol data unit (PDU) that requires multiple MAC frames. In this case, the following scenario occurs. A station with a multiframe LLC PDU to transmit sends out the MAC frames one at a time. Each frame is acknowledged after SIFS by the recipient. When the source receives an ACK, it immediately (after SIFS) sends the next frame in the sequence. The result is that once a station has contended for the channel, it will maintain control of the channel until it has sent all of the fragments of an LLC PDU.
- **Clear to Send (CTS):** A station can ensure that its data frame will get through by first issuing a small Request to Send (RTS) frame. The station to which this frame is addressed should immediately respond with a CTS frame if it is ready to receive. All other stations receive the RTS and defer using the medium.
- **Poll response:** This is explained in the following discussion of PCF.

The next longest IFS interval is the PIFS. This is used by the centralized controller in issuing polls and takes precedence over normal contention traffic. However, those frames transmitted using SIFS have precedence over a PCF poll.

Finally, the DIFS interval is used for all ordinary asynchronous traffic.

Point Coordination Function

PCF is an alternative access method implemented on top of the DCF. The operation consists of polling by the centralized polling master (point coordinator). The point coordinator makes use of PIFS when issuing polls. Because PIFS is smaller than DIFS, the point coordinator can seize the medium and lock out all asynchronous traffic while it issues polls and receives responses.

As an extreme, consider the following possible scenario. A wireless network is configured so that a number of stations with time-sensitive traffic are controlled by the point coordinator while remaining traffic contends for access using CSMA. The point coordinator could issue polls in a round-robin fashion to all stations configured for polling. When a poll is issued, the polled station may respond using SIFS. If the point coordinator receives a response, it issues another poll using PIFS. If no response is received during the expected turnaround time, the coordinator issues a poll.

If the discipline of the preceding paragraph were implemented, the point coordinator would lock out all asynchronous traffic by repeatedly issuing polls. To prevent this, an interval known as the superframe is defined. During the first part of this interval, the point coordinator issues polls in a round-robin fashion to all stations configured for polling. The point coordinator then idles for the remainder of the superframe, allowing a contention period for asynchronous access.

Figure 14.7b illustrates the use of the superframe. At the beginning of a superframe, the point coordinator may optionally seize control and issues polls for a give period of time. This interval varies because of the variable frame size issued by responding stations. The remainder of the superframe is available for contention-based access. At the end of the superframe interval, the point coordinator contends for access to the medium using PIFS. If the medium is idle, the point coordinator gains immediate access and a full superframe period follows. However, the medium may be busy at the end of a superframe. In this case, the point coordinator must wait until the medium is idle to gain access; this results in a foreshortened superframe period for the next cycle.

MAC Frame

Figure 14.8a shows the 802.11 frame format. This general format is used for all data and control frames, but not all fields are used in all contexts. The fields are as follows:

- **Frame control:** Indicates the type of frame and provides control information, as explained presently.
- **Duration/connection ID:** If used as a duration field, indicates the time (in microseconds) the channel will be allocated for successful transmission of a MAC frame. In some control frames, this field contains an association, or connection, identifier.

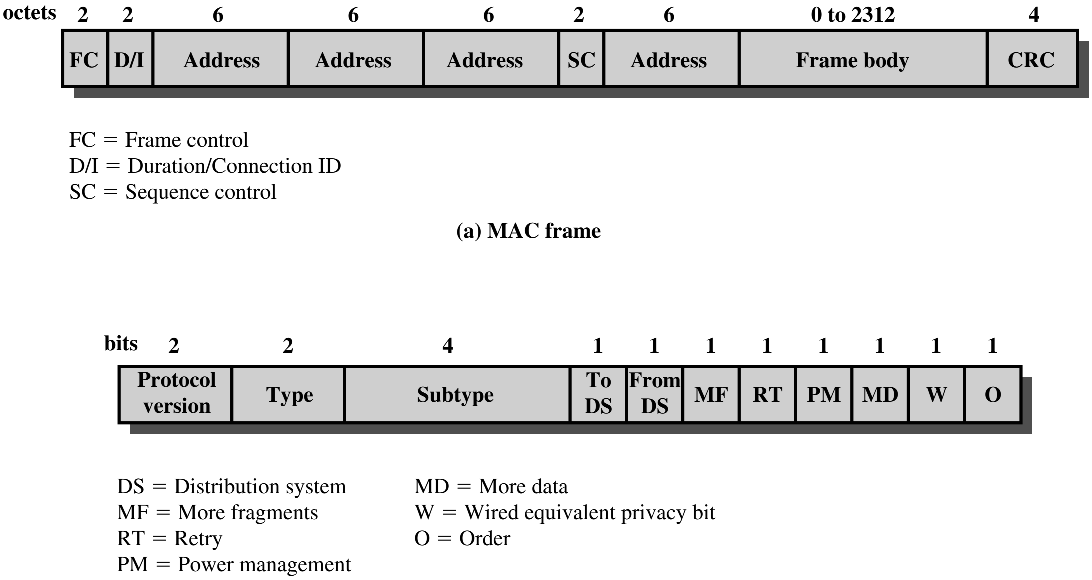

(a) MAC frame

(b) Frame control field

Figure 14.8 IEEE 802.11 MAC Frame Format

- **Addresses:** The number and meaning of the address fields depend on context. Address types include source, destination, transmitting station, and receiving station.
- **Sequence control:** Contains a 4-bit fragment number subfield, used for fragmentation and reassembly, and a 12-bit sequence number used to number frames sent between a given transmitter and receiver.
- **Frame body:** Contains an MSDU or a fragment of an MSDU. The MSDU is a LLC protocol data unit or MAC control information.
- **Frame check sequence:** A 32-bit cyclic redundancy check.

The frame control field, shown in Figure 14.8b, consists of the following fields:

- **Protocol version:** 802.11 version, currently version 0.
- **Type:** Identifies the frame as control, management, or data.
- **Subtype:** Further identifies the function of frame. Table 14.3 Defines the valid combinations of type and subtype.
- **To DS:** The MAC coordination sets this bit to 1 in a frame destined to the distribution system.
- **From DS:** The MAC coordination sets this bit to 1 in a frame leaving the distribution system.
- **More fragments:** Set to 1 if more fragments follow this one.
- **Retry:** Set to 1 if this is a retransmission of a previous frame.
- **Power management:** Set to 1 if the transmitting station is in a sleep mode.
- **More data:** Indicates that a station has additional data to send. Each block of data may be sent as one frame or a group of fragments in multiple frames.
- **WEP:** Set to 1 if the optional wired equivalent protocol is implemented. WEP is used in the exchange of encryption keys for secure data exchange.
- **Order:** Set to 1 in any data frame sent using the Strictly Ordered service, which tells the receiving station that frames must be processed in order.

We now look at the various MAC frame types.

Control Frames

Control frames assist in the reliable delivery of data frames. There are six control frame subtypes:

- **Power save-poll (PS-Poll):** This frame is sent by any station to the station that includes the AP (access point). Its purpose is to request that the AP transmit a frame that has been buffered for this station while the station was in power-saving mode.
- **Request to send (RTS):** This is the first frame in the four-way frame exchange discussed under the subsection on reliable data delivery at the beginning of Section 14.3. The station sending this message is alerting a potential destination, and all other stations within reception range, that it intends to send a data frame to that destination.

Table 14.3 Valid Type and Subtype Combinations

Type Value	Type Description	Subtype Value	Subtype Description
00	Management	0000	Association request
00	Management	0001	Assocation response
00	Management	0010	Reassociation request
00	Management	0011	Reassociation response
00	Management	0100	Probe request
00	Management	0101	Probe response
00	Management	1000	Beacon
00	Management	1001	Announcement traffic indication message
00	Management	1010	Dissociation
00	Management	1011	Authentication
00	Management	1100	Deauthentication
01	Control	1010	Power save - poll
01	Control	1011	Request to send
01	Control	1100	Clear to send
01	Control	1101	Acknowledgment
01	Control	1110	Contention-free (CF)-end
01	Control	1111	CF-end + CF-ack
10	Data	0000	Data
10	Data	0001	Data + CF-Ack
10	Data	0010	Data + CF-Poll
10	Data	0011	Data + CF-Ack + CF-Poll
10	Data	0100	Null function (no data)
10	Data	0101	CF-Ack (no data)
10	Data	0110	CF-poll (no data)
10	Data	0111	CF-Ack + CF-poll (no data)

- **Clear to send (CTS):** This is the second frame in the four-way exchange. It is sent by the destination station to the source station to grant permission to send a data frame.
- **Acknowledgment:** Provides an acknowledgment from the destination to the source that the immediately preceding data, management, or PS-Poll frame was received correctly.
- **Contention-free (CF)-end:** Announces the end of a contention-free period that is part of the point coordination function.

- **CF-end + CF-ack:** Acknowledges the CF-end. This frame ends the contention-free period and releases stations from the restrictions associated with that period.

Data Frames

There are eight data frame subtypes, organized into two groups. The first four subtypes define frames that carry upper-level data from the source station to the destination station. The four data-carrying frames are as follows:

- **Data:** This is the simplest data frame. It may be used in both a contention period and a contention-free period.
- **Data + CF-Ack:** May only be sent during a contention-free period. In addition to carrying data, this frame acknowledges previously received data.
- **Data + CF-Poll:** Used by a point coordinator to deliver data to a mobile station and also to request that the mobile station send a data frame that it may have buffered.
- **Data + CF-Ack + CF-Poll:** Combines the functions of the Data + CF-Ack and Data + CF-Poll into a single frame.

The remaining four subtypes of data frames do not in fact carry any user data. The Null Function data frame carries no data, polls, or acknowledgments. It is used only to carry the power management bit in the frame control field to the AP, to indicate that the station is changing to a low-power operating state. The remaining three frames (CF-Ack, CF-Poll, CF-Ack + CF-Poll) have the same functionality as the corresponding data frame subtypes in the preceding list (Data + CF-Ack, Data + CF-Poll, Data + CF-Ack + CF-Poll) but without the data.

Management Frames

Management frames are used to manage communications between stations and APs. The following subtypes are included:

- **Association request:** Sent by a station to an AP to request an association with this BSS. This frame includes capability information, such as whether encryption is to be used and whether this station is pollable.
- **Association response:** Returned by the AP to the station to indicate whether it is accepting this association request.
- **Reassociation request:** Sent by a station when it moves from one BSS to another and needs to make an association with the AP in the new BSS. The station uses reassociation rather than simply association so that the new AP knows to negotiate with the old AP for the forwarding of data frames.
- **Reassociation response:** Returned by the AP to the station to indicate whether it is accepting this reassociation request.
- **Probe request:** Used by a station to obtain information from another station or AP. This frame is used to locate an IEEE 802.11 BSS.

- **Probe response:** Response to a probe request.
- **Beacon:** Transmitted periodically to allow mobile stations to locate and identify a BSS.
- **Announcement traffic indication message:** Sent by a mobile station to alert other mobile stations that may have been in low power mode that this station has frames buffered and waiting to be delivered to the station addressed in this frame.
- **Dissociation:** Used by a station to terminate an association.
- **Authentication:** Multiple authentication frames are used in an exchange to authenticate one station to another, as described subsequently.
- **Deauthentication:** Sent by a station to another station or AP to indicate that it is terminating secure communications.

Security Considerations

IEEE 802.11 provides both privacy and authentication mechanisms.

The Wired Equivalent Privacy Algorithm

With a wireless LAN, eavesdropping is a major concern because of the ease of capturing a transmission. IEEE 802.11 incorporates WEP to provide a modest level of security. To provide privacy, as well as data integrity, WEP uses an encryption algorithm based on the RC4 encryption algorithm.

Figure 14.9a shows the encryption process. The integrity algorithm is simply the 32-bit CRC that is appended to the end of the MAC frame (Figure 14.8a). For the encryption process, a 40-bit secret key is shared by the two participants in the exchange. An initialization vector (IV) is concatenated to the secret key. The resulting block forms the seed that is input to the pseudorandom number generator (PRNG) defined in RC4. The PRNG generates a bit sequence of the same length as the MAC frame plus its CRC. A bit-by-bit exclusive-OR between the MAC frame and the PRNG sequence produces the ciphertext. The IV is attached to the ciphertext and the resulting block is transmitted. The IV is changed periodically (as often as every transmission). Every time the IV is changed, the PRNG sequence is changed, which complicates the task of an eavesdropper.

At the receiving end (Figure 14.9b), the receiver retrieves the IV from the data block and concatenates this with the shared secret key to generate the same key sequence used by the sender. This key sequence is then XORed with the incoming block to recover the plaintext. This technique makes use of the following property of XOR:

$$\mathrm{A} \oplus \mathrm{B} \oplus \mathrm{B} = \mathrm{A}$$

Thus, if we take the plaintext, XOR it with the key sequence, and then XOR the result with the key sequence, we get back the plaintext. Finally, the receiver compares the incoming CRC with the CRC calculated at the receiver to validate integrity.

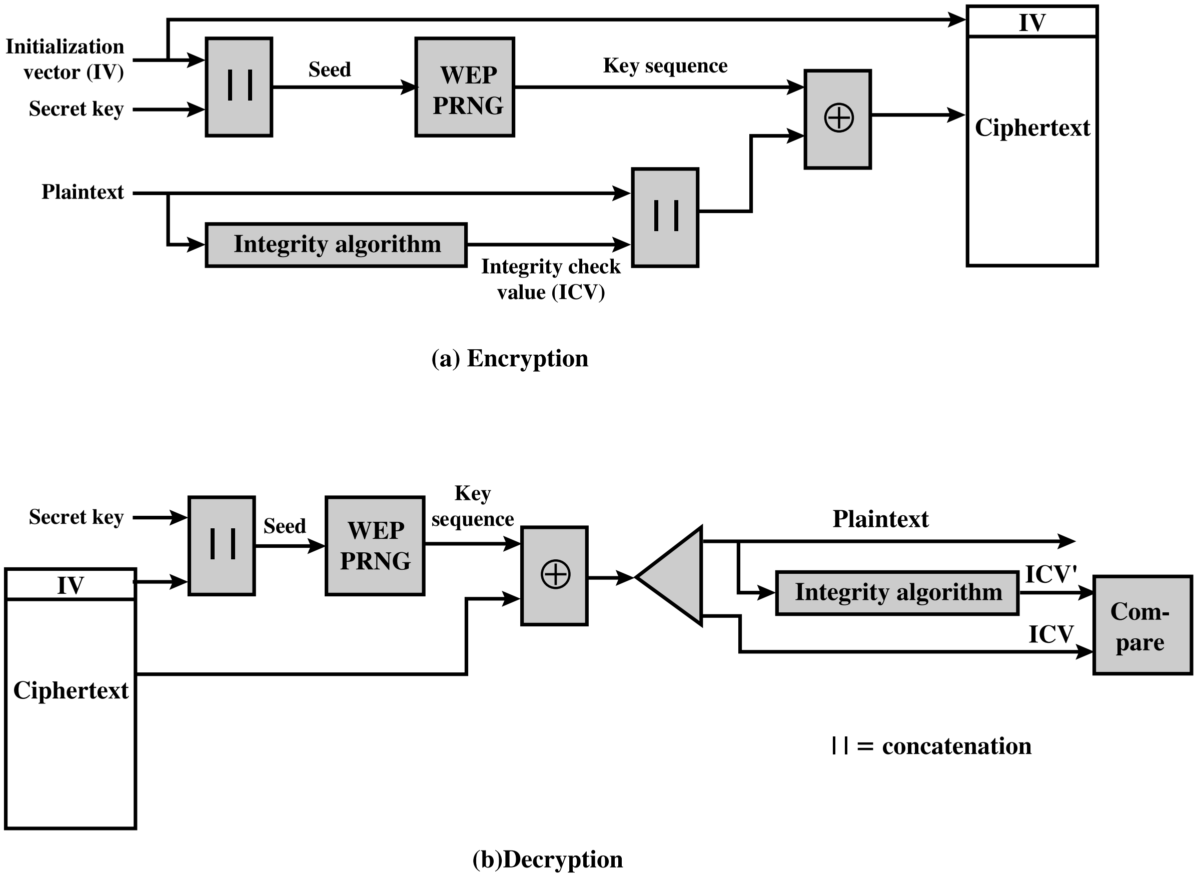

(a) Encryption

(b)Decryption

Figure 14.9 WEP Block Diagram

Authentication

IEEE 802.11 provides two types of authentication: open system and shared key. **Open system authentication** simply provides a way for two parties to agree to exchange data and provides no security benefits. In open system authentication, one party sends a MAC control frame, known as an authentication frame, to the other party. The frame indicates that this is an open system authentication type. The other party responds with its own authentication frame and the process is complete. Thus, open system authentication consists simply of the exchange of the identities between the parties.

Shared key authentication requires that the two parties share a secret key not shared by any other party. This key is used to assure that both sides are authenticated to each other. The procedure for authentication between two parties, A and B, is as follows:

1. A sends a MAC authentication frame with an authentication algorithm identification of "Shared Key" and with a station identifier that identifies the sending station.
2. B responds with an authentication frame that includes a 128-octet *challenge text*. The challenge text is generate using the WEP PRNG. The key and IV used in generating this challenge text are unimportant because they do not play a role in the remainder of the procedure.
3. A transmits an authentication frame that includes the challenge text just received from B. The entire frame is encrypted using WEP.
4. B receives the encrypted frame and decrypts it using WEP and the secret key shared with A. If decryption is successful (matching CRCs), then B compares the incoming challenge text with the challenge text that it sent in the second message. B then sends an authentication message to A with a status code indicating success or failure.

14.4 IEEE 802.11 PHYSICAL LAYER

The physical layer for IEEE 802.11 has been issued in three stages; the first part was issued in 1997 and the remaining two parts in 1999. The first part, simply called IEEE 802.11, includes the MAC layer and three physical layer specifications, two in the 2.4-GHz band and one in the infrared, all operating at 1 and 2 Mbps. IEEE 802.11a operates in the 5-GHz band at data rates up to 54 Mbps. IEEE 802.11b operates in the 2.4-Ghz band at 5.5 and 11 Mbps. We look at each of these in turn.

Original IEEE 802.11 Physical Layer

Three physical media are defined in the original 802.11 standard:

- Direct-sequence spread spectrum operating in the 2.4-GHz ISM band, at data rates of 1 Mbps and 2 Mbps

- Frequency-hopping spread spectrum operating in the 2.4-GHz ISM band, at data rates of 1 Mbps and 2 Mbps
- Infrared at 1 Mbps and 2 Mbps operating at a wavelength between 850 and 950 nm

Table 14.4 summarizes key details.

Table 14.4 IEEE 802.11 Physical Layer Specifications

(a) Direct sequence spread spectrum

Data rate	Chipping code length	Modulation	Symbol rate	Bits/symbol
1 Mbps	11 (Barker sequence)	DBPSK	1 Msps	1
2 Mbps	11 (Barker sequence)	DQPSK	1 Msps	2
5.5 Mbps	8 (CCK)	DBPSK	1.375 Msps	4
11 Mbps	8 (CCK)	DQPSK	1.375 Msps	8

(b) Frequency-hopping spread spectrum

Data rate	Modulation	Symbol rate	Bits/symbol
1 Mbps	Two-level GFSK	1 Msps	1
2 Mbps	Four-level GFSK	1 Msps	2

(c) Infrared

Data rate	Modulation	Symbol rate	Bits/symbol
1 Mbps	16-PPM	4 Msps	0.25
2 Mbps	4-PPM	4 Msps	0.5

(d) Orthogonal FDM

Data rate	Modulation	Coding rate	Coded bits per subcarrier	Code bits per OFDM symbol	Data bits per OFDM symbol
6 Mbps	BPSK	1/2	1	48	24
9 Mbps	BPSK	3/4	1	48	36
12 Mbps	QPSK	1/2	2	6	48
18 Mbps	QPSK	3/4	2	96	72
24 Mbps	16-QAM	1/2	4	192	96
36 Mbps	16-QAM	3/4	4	192	144
49 Mbps	64-QAM	2/3	6	288	192
54 Mbps	16-QAM	3/4	6	288	216

Direct-Sequence Spread Spectrum

Up to seven channels, each with a data rate of 1 Mbps or 2 Mbps, can be used in the DS-SS system. The number of channels available depends on the bandwidth allocated by the various national regulatory agencies. This ranges from 13 in most European countries to just one available channel in Japan. Each channel has a bandwidth of 5 Mhz. The encoding scheme that is used is DBPSK for the 1-Mbps rate and DQPSK for the 2-Mbps rate.

Recall from Chapter 7 that a DS-SS system makes use of a chipping code, or pseudonoise sequence, to spread the data rate and hence the bandwidth of the signal. For IEEE 802.11, a Barker sequence is used.

A Barker sequence is a binary $\{-1, +1\}$ sequence $\{s(t)\}$ of length n with the property that its autocorrelation values $R(\tau)$ satisfy $|R(\tau)| \leq 1$ for all $|\tau| \leq (n - 1)$. Further, the Barker property is preserved under the following transformations.

$$s(t) \rightarrow -s(t) \qquad s(t) \rightarrow (-1)^t s(t) \qquad \text{and} \qquad s(t) \rightarrow -\, s(n-1-t)$$

as well as under compositions of these transformations. Only the following Barker sequences are known:

$n = 2$	$+ +$
$n = 3$	$+ + -$
$n = 4$	$+ + + -$
$n = 5$	$+ + + - +$
$n = 7$	$+ + + - - + -$
$n = 11$	$+ - + + - + + + - - -$
$n = 13$	$+ + + + + - - + + - + - +$

The 11-chip Barker sequence is used. Thus, each data binary 1 is mapped into the sequence $\{+ - + + - + + + - - -\}$, and each binary 0 is mapped into the sequence $\{- + - - + - - - + + +\}$.

Important characteristic of Barker sequences are their robustness against interference and their insensitivity to multipath propagation.

Frequency-Hopping Spread Spectrum

Recall from Chapter 7 that a FH-SS system makes use of a multiple channels, with the signal hopping from one channel to another based on a pseudonoise sequence. In the case of the IEEE 802.11 scheme, 1-MHz channels are used. The number of channels available ranges from 23 in Japan to 70 in the United States.

The details of the hopping scheme are adjustable. For example, the minimum hop rate for the United States is 2.5 hops per second. The minimum hop distance in frequency is 6 MHz in North America and most of Europe and 5 MHz in Japan.

For modulation, the FH-SS scheme uses two-level Gaussian FSK for the 1-Mbps system. The bits zero and one are encoded as deviations from the current carrier frequency. For 2 Mbps, a four-level GFSK scheme is used, in which four different deviations from the center frequency define the four 2-bit combinations.

Infrared

The IEEE 802.11 infrared scheme is omnidirectional (Figure 13.6) rather than point to point. A range of up to 20 m is possible. The modulation scheme for the 1-Mbps data rate is known as 16-PPM (pulse position modulation). In this scheme, each group of 4 data bits is mapped into one of the 16-PPM symbols; each symbol is a string of 16 bits. Each 16-bit string consists of fifteen 0s and one binary 1. For the 2-Mbps data rate, each group of 2 data bits is mapped into one of four 4-bit sequences. Each sequence consists of three 0s and one binary 1. The actual transmission uses an intensity modulation scheme, in which the presence of a signal corresponds to a binary 1 and the absence of a signal corresponds to binary 0.

IEEE 802.11a

The IEEE 802.11a specification makes use of the 5-GHz band. Unlike the 2.4-GHz specifications, IEEE 802.11 does not use a spread spectrum scheme but rather uses orthogonal frequency division multiplexing (OFDM). Recall from Section 11.2 that OFDM, also called multicarrier modulation, uses multiple carrier signals at different frequencies, sending some of the bits on each channel. This is similar to FDM. However, in the case of OFDM, all of the subchannels are dedicated to a single data source.

The possible data rates for IEEE 802.11a are 6, 9, 12, 18, 24, 36, 48, and 54 Mbps. The system uses up to 52 subcarriers that are modulated using BPSK, QPSK, 16-QAM, or 64-QAM, depending on the rate required. Subcarrier frequency spacing is 0.3125 MHz. A convolutional code at a rate of 1/2, 2/3, or 3/4 provides forward error correction.

IEEE 802.11b

IEEE 802.11b is an extension of the IEEE 802.11 DS-SS scheme, providing data rates of 5.5 and 11 Mbps. The chipping rate is 11 MHz, which is the same as the original DS-SS scheme, thus providing the same occupied bandwidth. To achieve a higher data rate in the same bandwidth at the same chipping rate, a modulation scheme known as complementary code keying (CCK) is used.

The CCK modulation scheme is quite complex and is not examined in detail here. Figure 14.10 provides an overview of the scheme for the 11-Mbps rate. Input data are

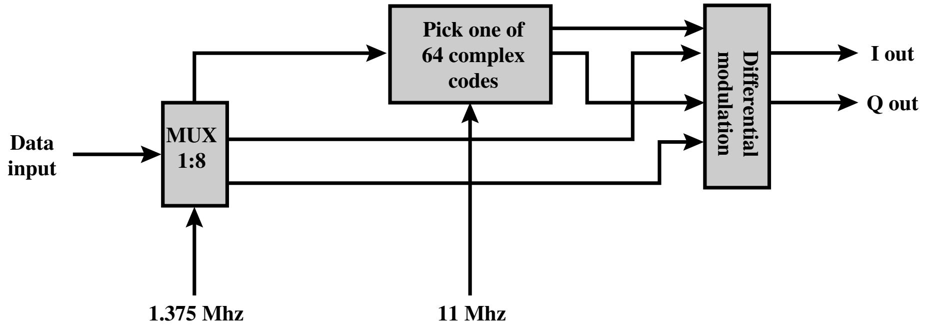

Figure 14.10 11-Mbps CCK Modulation Scheme

treated in blocks of 8 bits at a rate of 1.375 MHz (8 bits/symbol × 1.375 MHz = 11 Mbps). Six of these bits are mapped into one of 64 codes sequences based on the use of the 8 × 8 Walsh matrix (Figure 7.17). The output of the mapping, plus the two additional bits, forms the input to a QPSK modulator.

14.5 RECOMMENDED READING AND WEB SITES

[OHAR99] is an excellent technical treatment of IEEE 802.11. [GEIE99] also provides detailed coverage of the IEEE 802.11 standards, and numerous case studies. [CROW97] is a good survey article on the 802.11 standards. Neither of the last two references covers IEEE 802.11a and IEEE 802.11b. [GEIE01] has a good discussion of IEEE 802.11a.

CROW97 Crow, B., et al. "IEEE 802.11 Wireless Local Area Networks." *IEEE Communications Magazine*, September 1997.

GEIE99 Geier, J. *Wireless LANs.* New York: Macmillan Technical Publishing, 1999.

GEIE01 Geier, J. "Enabling Fast Wireless Networks with OFDM." *Communications System Design*, February 2001. (www.csdmag.com)

OHAR99 Ohara, B., and Petrick, A. *IEEE 802.11 Handbook: A Designer's Companion.* New York: IEEE Press, 1999.

Recommended Web sites:

- **The IEEE 802.11 Wireless LAN Working Group:** Contains working group documents plus discussion archives.
- **Wireless Ethernet Compatibility Alliance:** An industry group promoting the interoperability of 802.11 products with each other and with Ethernet.

14.6 KEY TERMS AND REVIEW QUESTIONS

Key Terms

access point (AP)
Barker sequence
basic service set (BSS)
binary exponential backoff
complementary code keying (CCK)
coordination function
distributed coordination function (DCF)
distribution system (DS)
extended service set (ESS)
logical link control (LLC)
MAC protocol data unit (MPDU)
MAC service data unit (MSDU)
medium access control (MAC)
open system authentication
point coordination function (PCF)
shared key authentication
wired equivalent privacy (WEP)

Review Questions

1. List and briefly define the IEEE 802 protocol layers.
2. What is the difference between a MAC address and an LLC address?
3. List and briefly define LLC services.
4. What is the difference between an access point and a portal?
5. Is a distribution system a wireless network?
6. List and briefly define IEEE 802.11 services.
7. How is the concept of an association related to that of mobility?
8. What characteristics of a wireless LAN present unique security challenges not found in wired LANs?
9. Which form of authentication is more secure and why: open system authentication or shared key authentication?