Chapter 4: IEEE 802.11 In Depth

Overview

Describe and apply the following concepts surrounding WLAN frames

- Terminology Review: Frames, Packets, and Datagrams
- Terminology Review: Bits, Bytes, and Octets
- Terminology: MAC and PHY

Understand IEEE 802.11 CSMA/CA
Understand and compare frame types and formats

- IEEE 802.11 Frame Format Versus IEEE 802.3 Frame Format
- Layer 3 Protocol Support by IEEE 802.11 Frames
- Jumbo Frame Support (Layer 2)
- MTU Discovery and Functionality (Layer 3)

Identify, explain, and apply the frame and frame exchange sequences covered by the IEEE 802.11 standard (as amended)

- Active (Probes) and Passive (Beacons) Scanning
- Dynamic Rate Switching

Summarize the processes involved in authentication and association

- The IEEE 802.11 State Machine
- Open System Authentication, Shared Key Authentication, and Deauthentication
- Association, Reassociation, and Disassociation

Identify and apply regulatory domain requirements Define, describe, and apply IEEE 802.11 coordination functions and channel access methods and features available for optimizing data flow across the RF medium

- DCF and HCF Coordination Functions
- EDCA Channel Access Method
- RTS/CTS and CTS–to–Self Protocols
- Fragmentation

In the preceding chapter, you discovered the basics of modulation and coding at the Physical layer.

This chapter will reveal more information related to the Physical layer of IEEE 802.11 communications, but it will also begin the process of revealing the Medium Access Control (MAC) layer functionality.

You will learn about the terms related to IEEE 802.11 networks and then move on to understanding the collision domain management technologies as well as the framing techniques used.

Finally, you'll learn about authentication and association as well as data flow optimization.

Terminology Review

Understanding any technology requires the use of a shared language. This means that not only the grammar of that language, but the definitions must be agreed upon, if you are to have meaningful dialog.

This section provides an encyclopedic overview of many terms that will be used throughout

this chapter and the remaining chapters of this book. Rather than simply provide a definition, we will provide a detailed description of the functionality of these terms and an explanation of how they are related to each other.

IEEE 802.11 CSMA/CA

Ethernet networks (IEEE 802.3) use a form of collision management known as collision detection (CD).

Wireless networks use a different form of collision management known as collision avoidance (CA). The full name of the physical media access management used in wireless networks is carrier sense multiple access/collision avoidance or CSMA/CA. We'll investigate CSMA/CD the MAC method used in Ethernet networks and then we'll look at CSMA/CA.

The essence of CSMA/CA is that collisions can happen many places on the medium, at any time during a transmission, and likely cannot be detected by the transmitter at its location. Listening for evidence of a collision while transmitting is thus worthless and not a part of the protocol.

This is because transmissions cannot be aborted early. Collisions are only inferred as one possible explanation for failure to receive an immediate acknowledgment (ACK) after transmitting a frame in its entirety.

The frame must be retransmitted completely. Under these circumstances there is much value in collision avoidance, and there is much of it in the IEEE 802.11 protocols.

The *carrier sense* in CSMA means that the devices will attempt to sense whether the physical medium is available before communicating.

The *multiple access* indicates that multiple devices will be accessing the physical medium. In a CD implementation of CSMA, when a collision is detected, both devices go silent for a pseudo–random period of time.

Since the time period is different for each device, they are not likely to try communicating at the same time again. This process helps recover from collisions and avoid another collision. In a CSMA/CD implementation, collisions occur because devices can begin communicating at the same time even though they both listened for "silence" on the physical medium. There was indeed silence, but both devices broke the silence at the same moment.

CSMA/CA is used in wireless networks and it was also used in early Apple LocalTalk networks, which were wired networks that were common to Apple devices. Collision avoidance is achieved by signaling to the other devices that one device is about to communicate.

This would be like saying "Listen, for the next few minutes, because I will be talking" in a telephone conversation. You are avoiding the collision by announcing that you are going to be communicating for some time interval.

CSMA/CA is not perfect due to hidden node problems, which will be covered in detail in but it provides a more efficient usage of a medium like RF medium than would CSMA/CD.

Carrier Sense

Carrier sense is the process of checking to see if the medium is in use or busy. In IEEE 802.11 WLANs, there are two kinds of carrier sense that are performed: virtual carrier sense and physical carrier sense.

Physical carrier sense uses *clear channel assessment (CCA)* to determine if the physical medium is in use. CCA is accomplished by monitoring the medium to determine if the amount of RF energy detected exceeds a particular threshold.

Due to the nature of WLAN architectures, there is no requirement for all stations to be able to hear all other stations existing in the same basic service set (BSS). This is because the wireless access point forms a kind of hub for the BSS.

A station may be able to hear the access point and the access point may be able to hear the other station, but the two stations may not be able to hear each other.

This results in what is commonly known as the *hidden node problem.* For this reason, wireless networks must use other forms of carrier sense to deal with medium access control.

The other form of carrier is virtual carrier sense, which uses a network allocation vector (NAV). The NAV is a timer in each station that is used to determine if the station can utilize the medium.

If the NAV has a value of 0, the station may contend for the medium. If the NAV has a value greater than 0, the station must wait until the timer counts down to 0 to contend for the medium.

Stations configure their NAV timers according to Duration fields in other frames using the medium. For example, if a station detects a frame with a specific duration set in the Duration field, it will set the NAV timer to this duration and will then wait until that time has expired before contending for access.

To be clear, both the physical carrier sense and the virtual carrier sense must show that the medium is available before the station can contend for access. In other words, if the NAV timer reaches 0 and the station uses CCA to detect activity on the medium only to find there is such activity, the station still cannot transmit.

In this case, another frame may be pulled from the medium and used to set a new NAV timer value for countdown. While it may seem that this would prevent a station from ever communicating, the rate of frame transfer is so high that all of these actions usually take place in far less than 1 second.

Note An additional form of carrier sense that is not often written about is what you might call phantom frame sensing. In this scenario, the PHY reads an incoming PLCP header length value and loses the incoming signal completely. However, since the header length was read, the device can still defer to the rest of the phantom frame.

Some systems uses a technique called polling to avoid collisions, this is called N-Stream. Mikrotik and Tsunami is such systems.

Interframe Spacing

After the station has determined that the medium is available, using carrier sensing techniques, it must observe *interframe spacing (IFS)* policies.

IFS is a time interval in which frames cannot be transmitted by stations within a BSS. This space between frames ensures that frames do not overlap each other. The time interval differs, depending on the frame type and the applicable IFS type for that frame.

While the IFS implementation in IEEE 802.11 systems can result in the appearance of QoS, it should not be confused with *IEEE 802.11e* or any Layer 3 or higher QoS solution.

IFS is an 802.11 feature that allows for dependent frames to be processed in a timely manner.

For example, a standard 802.11 data frame is transmitted using the distributed IFS (DIFS) interval and the ACK to this data frame is sent back using the short IFS (SIFS) interval.

Because the ACK uses a SIFS interval, the ACK frame will take priority over any other data frames that are waiting to be transmitted.

This way, the original station that transmitted the data frame will receive the ACK frame and not attempt to resend the data frame. In other words, the frame–to–IFS interval relationships that are specified in the IEEE 802.11 standard ensure that frames will be processed in their proper sequence.

I've mentioned some of the IFS types defined by the IEEE 802.11 standard already. These IFS types include the following types, which will now be covered in more detail:

- SIFS
- PIFS
- DIF
- EIFS

SIFS

Short interframe spacing (SIFS) is the shortest of the available IFS parameters. Frames that are specified to use SIFS will take priority over frames that are specified to use PIFS, DIFS, or extended IFS (EIFS).

This priority function is simply a result of the IFS length. Since the SIFS is the shortest IFS, stations that are waiting to send a frame that is specified to use a SIFS interval will have a shorter wait time and will therefore have access to the wireless medium (WM) before other stations with frames specified for other IFS types.

SIFS is used for many different frames, including:

- ACK frames immediately following the receipt of a data frame
- Clear to Send (CTS) frames sent as a response to Request to Send (RTS) frames
- Data frames that immediately follow CTS frames
- With the exception of first exchange and error conditions, all frame exchanges made in Point Coordination Function (PCF) mode
- With the exception of the first fragment, all fragment frames that are part of a fragment burst

As technically defined by the IEEE 802.11 standard as amended, the SIFS time interval is the time from the end of the last symbol of the previous frame to the beginning of the first symbol of the preamble of the subsequent frame as seen at the air interface.

The accuracy level required is ±10% of the slot time for the PHY in use. For example, the actual SIFS time interval must be within 2 µs of the specified time interval for the DSSS PHY. Slot times for the various PHYs are listed in the later section of this chapter "Random Backoff Times." The SIFS times for the various PHYs are listed here:

- FHSS 28 µs
- DSSS 10 µs
- OFDM 6 µs
- HR/DSSS 10 µs
- ERP 10 µs

PIFS
Point (coordination function) interframe spacing (PIFS) is neither the shortest nor longest interval; it results in a priority greater than DIFS, but less than SIFS.

When an access point needs to switch the network from *Distributed Coordination Function (DCF)* mode to PCF mode, it will use PIFS frames. PCF is an optional part of IEEE 802.11 and has not been implemented in any market devices.

The PIFS duration interval is equal to the SIFS interval for the PHY and one slot time duration for the PHY. For example, DSSS has a 20 µs slot time and a 10 µs SIFS interval.

Therefore the PIFS interval in a DSSS PHY will be 30 µs. For another example, the OFDM PHY has a 9 µs slot time and a 16 µs SIFS interval. Therefore the PIFS interval in an OFDM PHY is 25 µs.

DIFS
Distributed (coordination function) interframe spacing (DIFS) is the longest of the three IFS types covered so far. It is used by standard data frames. The greater delay interval ensures that frames specified for SIFS and PIFS intervals are able to transmit before DIFS data frames.

The DIFS interval is calculated as the PHY's SIFS interval plus two times the PHY's slot time. Based on the same numbers used in the previous paragraphs for the PIFS interval calculations and this new algorithm for calculating the DIFS interval, the DSSS PHY has a DIFS interval of 50 µs and the OFDM PHY has a DIFS interval of 34 µs.

EIFS
Extended interframe spacing (EIFS) is used when a frame reception begins but the received frame is incomplete or is corrupted based on the Frame Check Sequence (FCS) value.

When the last frame of the station received was corrupted, the station uses EIFS for the next frame that it transmits. The EIFS interval is the longest of the IFS intervals and is calculated based on the following more complex algorithm:

$$EIFS = SIFS + (8 * ACKsize) + preamble\ length$$

where the time calculation is the amount of time in microseconds that it takes to transfer the 8 ACKs, preamble, and PLCP header. As you can see, the EIFS is more than the DIFS and SIFS combined.

Contention Window
The IFS delay interval is not the end of the wait for devices that are seeking time on the WM.

After the IFS delay interval has passed, the device must then initiate a random backoff algorithm and then contend for the WM if the Distributed Coordination Function is in effect.

This random backoff algorithm is processed and applied using the *contention window.*

Note The phrase *contention window* has caused much confusion, but it is the phrase in use in the IEEE 802.11 standard. This "window" is actually a range of integers from which one is chosen at random to become the backoff timer for the immediate frame queued for transmission.

Random Backoff Times

All stations having a frame to transmit choose a random time period within the range specified as the contention window. Next, the predefined algorithm multiplies the randomly chosen integer by a *slot time.*

The slot time is a fixed length time interval that is defined for each PHY, such as DSSS, FHSS, or OFDM. For example, FHSS uses a slot time of 50 µs, and DSSS uses a slot time of 20 µs. The slot times for each of the currently ratified PHYs are listed here:

- FHSS 50 µs
- DSSS 20 µs
- OFDM 9 µs
- HR/DSSS 20 µs
- ERP Long Slot Time 20 µs
- ERP Short Slot Time (802.11b compatible) 9 µs

As you can see, there are definite variations among the different PHYs supported in the IEEE 802.11 standard as amended. Though the IEEE 802.11n amendment is still in the draft stage at this writing, it is expected to use the standard 9 µs slot time used in existing PHYs that support OFDM.

Contention Explanation
Now that you have most of the pieces to the media contention puzzle, you can begin to put them together in order to understand how a wireless station decides when it should try to communicate on the WM.

In order to understand this, imagine that a station has a data frame that it needs to transmit on the WM. This data frame will be required to use the DIFS, since it is a standard data frame.

Furthermore, imagine that the station uses carrier sense to determine that a frame is currently being transmitted. For discussion's sake, let's assume that the station detected that the frame being transmitted had a Duration/ID field value of 20 µs.

The station sets its NAV to count down the 20 µs and waits. The NAV reaches 0, and the station uses carrier sense and detects that the WM is silent. At this time, the station must wait for the DIFS interval to expire, and since the station is using the DSSS PHY, it waits for 50 µs.

Next, the station waits for the random backoff time period to expire, and when it does, the station uses carrier sense and detects that the WM is silent. The station begins transmitting the data frame. All of this assumes the network is using the

Distributed Coordination Function, which is the only contention management functionality that has been widely implemented in hardware at this time.

Collision Avoidance

Ultimately, the carrier sense, IFS, and random backoff times are used in order to decrease the likelihood that any two stations will try to transmit at the same time on the WM.

The IFS parameters are also used in order to provide priority to the more time sensitive frames such as ACK frames and CTS frames.

The CCA (PHY and MAC), IFS, variable contention window, and random backoff times, together, form the core of the Distributed Coordination Function.

Even with all of these efforts, a collision can still occur. In order to deal with these scenarios, *acknowledgment,* or *ACK,* frames are used.

An ACK frame is a short frame that uses the SIFS to let the sending device know that the receiving device has indeed received the frame. If the sending device does not receive an ACK frame, it will attempt to retransmit the frame.

Since the retransmitted frame will be transmitted using the rules and guidelines we've talked about so far, chances are that the next frame or one of the next few will make it through.

Frame Types

Management Frames
Management frames are used to manage access to wireless networks and to move associations from one access point to another within an extended service set (ESS).

Control Frames
Control frames are used to assist with the delivery of data frames and must be able to be interpreted by all stations participating in a BSS. This means that they must be transmitted using a modulation technique and at a data rate compatible with all hardware participating in the BSS.

Data Frames
Data frames, the third frame type in our discussion, are the actual carriers of application level data. These frames can be either standard data frames or Quality of Service (QoS) data frames for devices supporting the IEEE 802.11e amendment.

Layer 3 Protocol Support by IEEE 802.11 Frames

Layer 3 protocols include the IP protocol, and wireless networks introduce new problems to the IP protocol that did not exist with wired networks. The primary new problem is also one of the main benefits of wireless networking: mobility.

Users can walk around your facility with their HR/ DSSS PocketPC or their laptop using some other IEEE WLAN technology, and as they walk, they may move beyond the range of the original access point to which they connected.

When this happens, their device may find a new access point with a stronger signal. This new access point may actually exist on a different LAN segment with different IP configurations.

The result is that the user loses his or her existing IP address and gets a new one from the DHCP service on the new segment.

Since many IP based applications rely on the TCP protocol and since the TCP protocol is a Connection oriented protocol, this release and reassignment of IP addresses can result in the loss of connection in the higher layers of the OSI model. For example, if the user was engaged in a Voice over IP (VoIP) or an FTP session, she will need to reconnect and restart that session. This is not only frustrating for the users, but can lead to corrupted or lost data.

Jumbo Frame Support (Layer 2)

As you've learned earlier in this chapter, the frame payload for an IEEE 802.11 frame (the MSDU) has a maximum size limit of 2304 bytes. However, both IP at Layer 3 and IEEE 802.3 at Layer 2 support a default data unit of 1500 bytes. This limit is also known as the MTU (maximum transfer unit).

In certain scenarios, this limit of 1500 bytes can be prohibitive to network performance. For this reason, *jumbo frames* have been introduced into the hardware of many vendors.

Jumbo frames are technically any frames allowing an MTU of greater than 1500. *Baby giant frames,* in Cisco terminology, are frames that support an MTU of up to 1552 bytes with a total frame size of 1600 bytes after headers and trailers are added. Most vendors max out jumbo frames at around 9000 bytes.

For example, the Cisco Catalyst 4000s and 6000s support a jumbo frame MTU of 9198 bytes and 9216 bytes, respectively. After headers and trailers are added, the maximum frame size is 9216 bytes for the Catalyst 4000s and 9234 bytes for the Catalyst 6000s. These are examples of jumbo frames, and you will not be tested on the specifics of what a certain vendor's hardware supports; however, you should know that jumbo frames are frames with an MTU larger than 1500 bytes and many vendors support sizes of more than 9000 bytes.

IEEE 802.11 Frames and Frame Exchange Sequences

While you will not be required to understand every detail of the frames and frame exchanges that occur on a WLAN, you will need to understand the basics of frame exchange sequences and the flow of creating a WLAN, accessing a WLAN, and disconnecting from a WLAN.

In this section you will learn about the beacon management frame, which is used to announce information about an available WLAN. You will also learn about the methods used by client stations to find these WLANs.

Finally, you'll learn how *dynamic rate switching* works so that a balance of data rate and efficient WM utilization is reached. First, you will learn about many of the functions or services provided by the IEEE 802.11 MAC.

MAC Functions

The IEEE 802.11 MAC provides the following functions:

Scanning Before a station can participate in a BSS, it must be able to find the access points that provide access to that service set. Scanning is the process used to discover BSSs or to discover access points within a known BSS.

Synchronization Some IEEE 802.11 features require all stations to have the same time. Stations can update their clocks based on the time stamp value in beacon frames.

Frame Transmission Stations must abide by the frame transmission rules of the BSS to which they are associated. These rules are the Distributed Coordination Function in all known systems at this time.

Authentication Authentication is performed before a station can be associated with a BSS. This will be covered in more detail in the later section of this chapter "Authentication and Association Processes."

Association Once authentication is complete, the station can become associated with the BSS. This includes discovery of capability information in both directions from the station to the access point and from the access point to the station. Association is covered in more detail in the later section of this chapter "Authentication and Association Processes."

Reassociation When a user roams throughout a service area, that user may reach a point where one access point within an ESS will provide a stronger signal than the currently

associated access point. When this occurs, the station will reassociate with the new access point.

Data Protection Data encryption may be employed to assist in preventing crackers from accessing the data that is transmitted on the WM.

Power Management Since the transmitters/receivers (transceivers) in wireless client devices consume a noteworthy amount of power, power management feature are provided that assist in extending battery life by causing the transceiver to sleep for specified intervals.

Fragmentation In certain scenarios, it is beneficial to fragment frames before they are transmitted onto the WM. This type of scenario most often occurs as a result of intermittent interference. Fragmentation is covered in more detail later in this chapter.

RTS/CTS Request to Send/Clear to Send is a feature of IEEE 802.11 that will help prevent hidden node problems and allow for more centralized control of access to the WM. RTS/CTS is covered in more detail later in this chapter.

Beacon Management Frame

The *beacon management frame* is a special type of frame used in IEEE 802.11 networks. This frame is often referred to as the beacon, since this is the frame subtype specified in IEEE 802.11 as amended.

In an ad hoc wireless network (IBSS), all the stations take turns broadcasting the beacon frame.

This is because there is no access point in an independent basic service set (IBSS). Beacon frames can be used by client stations seeking wireless network to join, or these client stations may use other frames known as *probe request* and *probe response* frames. Ad-Hoc is same as wireless workgroup.

Both methods will be covered in the following sections, "Active Scanning" and "Passive Scanning."

Active scanning uses probe request and probe response frames instead of the beacon frame to find a WLAN to join. A client station can use either of two general methods to find the WLAN. The first is to specify the SSID of the network being sought, and the second is to seek for any BSS that may be able to hear and respond to the probe request.

If the SSID is specified in the probe request frame transmitted by the requesting station, all access points that are configured with a matching SSID should respond, assuming they receive the probe request frame.

It is certainly possible that a set of access points using the same SSID could cover an area large enough that all of the access points will not receive the probe request transmitted from a specific location in that area.

The response from the access points that hear the probe request is a probe response frame. The probe response frame contains the same basic information that the beacon frame contains with the exception of the Traffic Indication Map.

If a probe request is transmitted onto the WM having a wildcard SSID (a null value for the SSID), all access points that receive the probe request will respond with a probe response containing their SSIDs. This is the standard behavior.

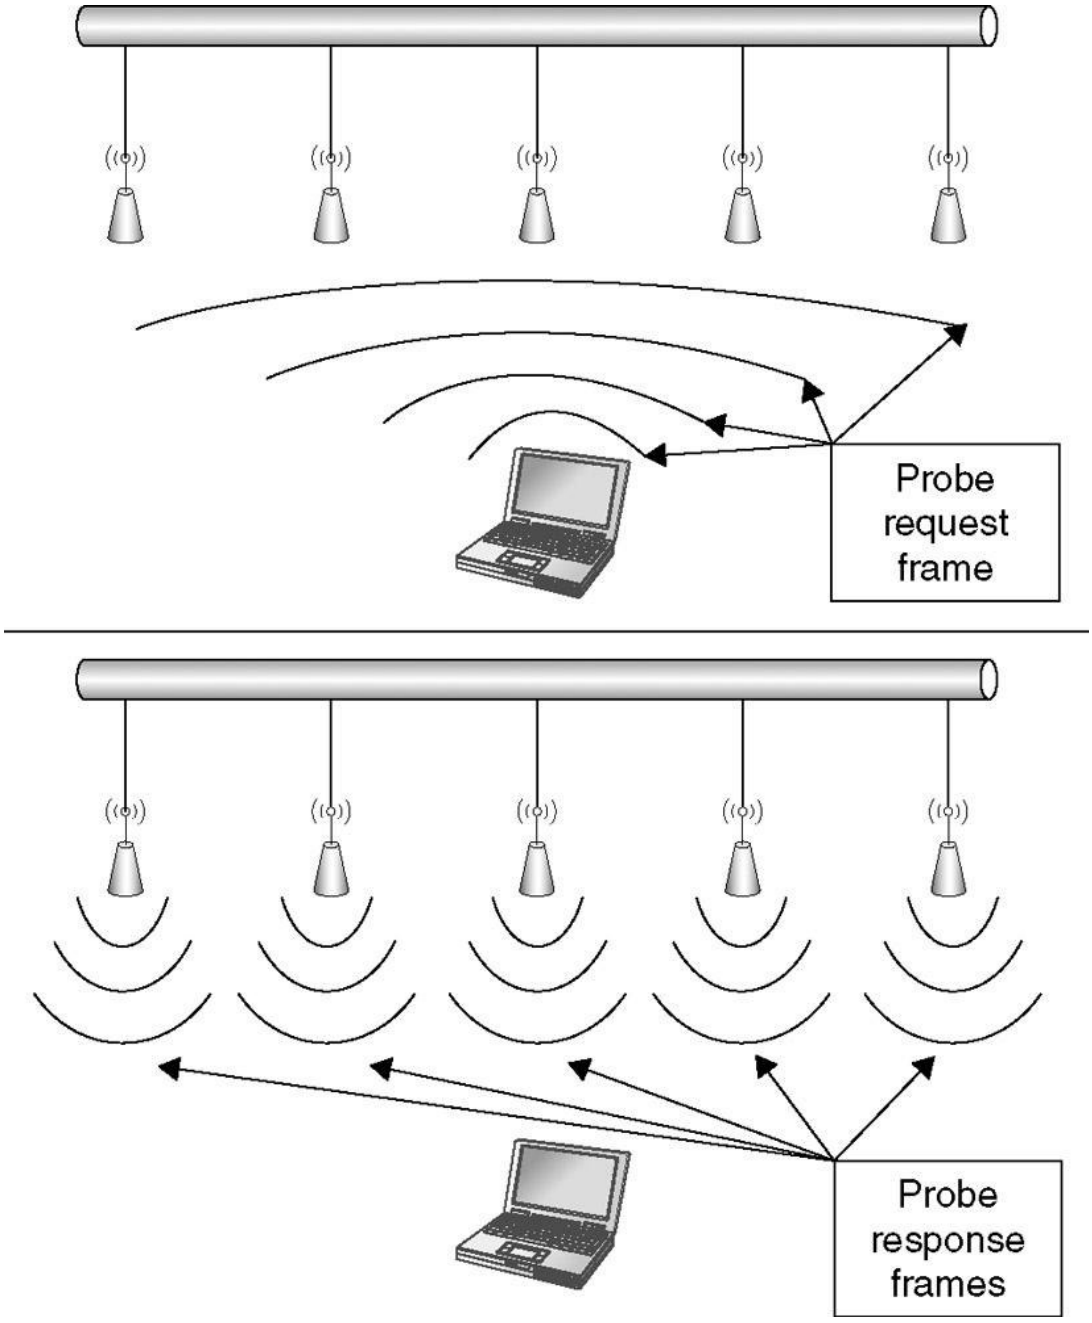

The exact details of the active scanning process are a bit more complex than the simple overview presented up to now. In fact, active scanning involves channel switching and scanning each channel in a station's channel list. The basic process is outlined here:

1. Switch to a channel.
2. Wait for an incoming frame or for the ProbeDelay timer to expire.

3. If the ProbeDelay timer expires, use DCF for access to the WM and send a probe request frame.

4. Wait for the MinChannelTime to pass.

a) If the WM was never busy, there is no WLAN on this channel. Move to the next channel.

b) If the WM was busy, wait until MaxChannelTime has expired and then process any probe response frames.

Authentication and Association Processes

Once a station has located the WLAN to which it seeks to connect, which is the first stage of station connectivity, it must go through the *authentication* and *association* processes, which are the second and third stages of connectivity.

This involves being authenticated by the access point or WLAN controller and then determining the PHY, data rate, and other parameters within which the association must operate.

The first step is authentication and the second is association, and each step is covered in sequence in this section.

The IEEE 802.11 State Machine

The *state machine* of the IEEE 802.11 standard can be in one of three states:

- Unauthenticated/Unassociated
- Authenticated/Unassociated
- Authenticated/Associated

Unauthenticated/Unassociated

In the initial state, a client station is completely disconnected from the WLAN. It cannot pass frames of any type through the access points to other stations on the WLAN or the wired infrastructure.

Authentication frames can be sent to the access points. These frames are not sent through the access points, with the exception of a split MAC implementation where a WLAN controller performs the authentication, but are sent to the access point.

The distinction is important. Frames must be transmitted to the access point in order to eventually reach the authenticated and associated stage; however, until the final stage is reached, only authentication and association request frames will be processed by the access point.

Access points, or WLAN controllers, keep a list known as the association table. Vendors report the stage of the station's state machine differently. Some vendors may report that a client that has not completed the authentication process is unauthenticated and other vendors may simply not show the client in the association table view.

Authenticated/Unassociated

The second state of the state machine is to authenticate an unassociated state. To move from the first state to the second, the client station must perform some kind of valid authentication. This is accomplished with authentication frames.
Once this second state is reached, the client station can issue association request frames to be processed by the access point; however, other IEEE 802.11 frame types are not allowed. In most access points, the association table will now show "authenticated" for the client station.

Since the interval between reaching the authenticated and unassociated stage and moving on to the authenticated and associated stage is very small (usually a matter of milliseconds),

you will not see client stations in the state very often. In most cases, you will either see "unauthenticated" or nothing for the first state and associated for the third state.

The only exception to this is what is sometimes called "preauthentication." A station can authenticate with any number of access points, but it can only be associated with one access point at a time.

The access point to which the station is associated must be a single entity in order for other devices on the network to be able to reach that station. In some systems, the station is capable of authenticating with multiple access points so that it can roam more quickly when the need arises.

Authenticated/Associated
The third and final state is the authenticated and associated state. In order for a station to be in this state, it must have first been authenticated and then associated. The process of moving from state 2 (authenticated and unassociated) to this state is a simple four frame transaction.

The client station first sends an association request frame to an access point to which it has been authenticated. Second, the access point responds with an acknowledgment frame. Next, the access point sends an association response frame either allowing or disallowing the association.

The client sends an acknowledgment frame as the fourth and final step. If the third step resulted in an approval of the association request, the client station has now reached the authenticated and associated state and may communicate on the WLAN or through the WLAN to the wired network if encryption keys match and 802.1X is not enabled.

The association response frame includes a status code element. If the status code is equal to 0, the association request is approved or successful. There are three other status codes that may apply: 12, 17, and 18.

A status code of 12 indicates that the association was rejected for some reason outside of the scope of the IEEE 802.11 standard.

A status code of 17 indicates that the access point is already serving the maximum number of client stations that it can support.

Finally, a status code of 18 indicates that the client station does not support all of the basic data rates required to join the BSS.

This last constraint is imposed to ensure that all stations will be able to receive certain frame types that are communicated at the basic data rates.

If they cannot receive these frame types, they cannot participate in the BSS lest they cause interference by not understanding such frames as CTS or by not having the ability to interpret frames at all.

The result would be that the station not supporting the basic data rates would cause interference due to an internal misconception that the WM was clear.

This is a simplification of a complex technical constraint, but it is sufficient for a WLAN administrator to know that a station cannot associate with a BSS if it does not support the basic data rates required.

Thankfully, these data rates are specified by administrators, and assuming the station is standards based and compatible with the PHY being used, this shouldn't be a problem with proper configuration settings.

The key point is to realize that you cannot transmit data frames for processing until you've been associated and you cannot transmit association frames for processing until you've been authenticated.

Now that you understand the three states in which a station can reside, let's explore the details of how the station can become authenticated and then associated.

Authentication

Based on the three states of a wireless station, you know that the second step to joining a WLAN, after discovery through scanning, is authentication. The IEEE 802.11–1999 (R2003) standard specifies two methods of authentication: *Open System authentication* and *Shared Key authentication.*

The first seems it would be used in less secure environments, while the second seems it would be used in more secure environments; however, you will soon learn why the opposite is true.

Open System Authentication

Open System authentication is essentially a null algorithm. In other words, no true authentication (verification of identity) occurs. Additionally, Open System authentication is specified as the default authentication mechanism in Clause 8 of the IEEE 802.11 standard. Figure below provides a step by step sequence of events that transpire in the Open System authentication process.

Step	Station Acting	Action Taken	Results
1	Client	Transmits an authentication frame with the code indicating that Open System authentication should be used.	Authentication frame transmitted
2	Access point	Receives the authentication frame and responds with an ACK frame.	Authentication frame received
3	Access point	Transmits an authentication frame to the client indicating a positive authentication.	Authentication approved
4	Client	Receives the positive authentication frame and response with an ACK frame to the access point.	Approved authentication acknowledged

Access points configured to use Open System authentication will always respond with a positive authentication to any authentication request.

Be careful not to confuse authentication with confidentiality. Data privacy or confidentiality is about protecting transmitted data from interception. Authentication is about verifying identities of senders and receivers on the network.

The point is that WEP is used in authentication (Shared Key) and it can also be used for confidentiality (data encryption). You can use WEP with both Open System authentication and Shared Key authentication for data confidentiality. Just as the next section points out the weaknesses of WEP as an authentication tool, you should consider it weak as a confidentiality tool.

Before you move away from Open System authentication with an assumption that it provides no use, keep the following realities in mind:

- Open System authentication is preferred at hotspots where you want to provide unauthenticated access to the Internet.

- More secure authentication technologies, such as 802.1X, rely on Open System authentication. In other words, Open System authentication leaves the access point open to other layers of security beyond the pre IEEE 802.11i authentication standards.

Shared Key Authentication
Shared Key authentication utilizes the *wired equivalent privacy (WEP)* key for authentication. WEP can also provide encryption of the MSDU, but Clause 8 defines this algorithm as providing protection from *casual eavesdropping* and should be understood as not providing protection from structured attacks.

Due to the weaknesses discovered in the WEP algorithm, very few networks should implement and use Shared Key authentication or WEP encryption today. Certainly, the networks that do utilize these algorithms are insecure and should be upgraded as soon as possible.

When Shared Key authentication is used, the client station and the access point must both use the same WEP key. Access points can store multiple WEP keys so that some stations can communicate using one WEP key and other stations can communicate using another.

The fact that both stations (the client and the access point) share the same key gives rise to the name Shared Key. The Shared Key authentication process is documented in below as a sequence of steps with descriptions of the activities that occur in each step.

Step	Station Acting	Action Taken	Results
1	Client	Transmits an authentication frame with the code indicating that Shared Key authentication should be used.	Authentication frame transmitted
2	Access point	Receives the authentication frame and transmits an ACK frame to the client.	Authentication frame received
3	Access point	Generates challenge text (randomly generated plain text) and transmits it to the client station in an authentication frame.	Authentication challenge transmitted
4	Client	Receives the authentication challenge frame and transmits an ACK frame to the access point.	Authentication challenge received
5	Client	Places the challenge text in another frame while encrypting it with the WEP key and transmits the challenge response to the access point in another authentication frame.	Authentication challenge response transmitted
6	Access point	Receives the challenge response authentication frame from the client and transmits an ACK frame to the client.	Authentication challenge response received
7	Access point	Decrypts the encrypted text in the authentication challenge response frame and, if it matches the original text, authenticates the client by transmitting a positive authentication frame to the client. If it does not match the original text, a negative authentication frame is sent to the client.	Authentication challenge response processed
8	Client	Client station received the results and transmits an ACK frame to the access point.	Authentication process complete

Deauthentication
Deauthentication frames are known as advisory frames. This is because they advise the network of something and the network cannot prevent that thing from occurring.

In other words, a standard IEEE 802.11based access point cannot deny a deauthentication frame. This frame would be transmitted to the access point (or other members of the IBSS in

an ad hoc network) and the receiving device would simply acknowledge the deauthentication.

This would also result in a lowering of the state machine's state in the access point's association table.

A deauthentication frame will include the address of the station being deauthenticated and the address of the station with which the deauthenticating station is currently authenticated.

The deauthentication frame will have a reason code of 3, which indicates that the deauthenticating station is either leaving or has left the BSS or ESS.

Remember that authentication must happen before association can take place; for this reason, a deauthentication frame effectively disassociates and deauthenticates the transmitting client station from the access point.

Association, Reassociation, and Disassociation

After authentication comes association. As was stated earlier, a station can be authenticated with multiple access points, but it can be associated with only one. There are three frames related to association: association frames, reassociation frames, and disassociation frames.

Association

The process of association is very simple. Four frames are transmitted between the client station and the access point station.

The first frame is an association request frame, which is followed by an acknowledgment frame from the access point.

The third frame is an association response frame, which is followed by an acknowledgment frame from the client station. It is extremely rare for a client station to successfully authenticate and then fail to associate.

This is because the client station can usually determine if it is compatible with the BSS by inspecting the beacon frames or probe response frames sent from the access points.

Reassociation

Reassociation occurs when a client station roams from one access point to another within an ESS.

Because reassociation is part of the roaming process, it will be covered in more detail in the next chapter. An immobile station may also reassociate with its access point in order to change its Robust Security Network Association (RSNA).

Disassociation

Like the deauthentication frame, a *disassociation* frame is an advisory frame in that the access point cannot deny the disassociation. The disassociation service is the component of the MAC layer that is responsible for processing a disassociation. This is one of the 13 architectural services of the IEEE 802.11 MAC layer. The full list of services is provided below with link to the station type that contains the service.

Service	Station Type
Authentication	All stations
Deauthentication	All stations
Association	Distribution System Service
Disassociation	Distribution System Service
Reassociation	Distribution System Service
Distribution	Distribution System Service
Integration	Distribution System Service
MSDU delivery	All stations
Data confidentiality	All stations
DFS	All stations
TPC	All stations
Higher–layer timer synchronization	All stations
QoS traffic scheduling (optional)	All stations and DSS

Regulatory Domain Requirements

Amendments *d, h,* and *j* define a Country element for the IEEE 802.11 standard. This code must be present in any beacon frame where the dot11MultiDomainCapabilityEnabled attribute is equal to TRUE.

In other words, if the access point supports operation in multiple regulatory domains and is IEEE 802.11 compliant, it must include a Country code in order for connecting stations to configure their available channels, minimum output power, and maximum output power settings.

Without investigating the 802.11d, 802.11h, and 802.11j amendments, you would miss this Country code element. It is through the implementation of the Country element that a BSS or ESS can identify the local regulatory domain and automatically configure itself to adhere to the regulations of the local regulatory domain.

Data Flow Optimization Across the RF Medium

Because the RF medium is not a physical medium that can implement limited access by disallowing physical connections, some method must be used to control access to the medium.

Two methods are defined in the IEEE 802.11–1999 (R2003) standard, and more specifications are defined in IEEE 802.11e, which is an amendment that details QoS in wireless networks.

The two methods that have been described since the earliest IEEE 802.11 standard was released are DCF and PCF.

DCF

The Distributed Coordination Function (DCF) is the WM access method that was described earlier in this chapter as CSMA/CA. At least, DCF is the IEEE 802.11 implementation of CSMA/CA.
This means that DCF is inclusive of the carrier sensing mechanisms, interframe spacing, and backoff timers discussed earlier. DCF is said to be a *distributed* coordination function because the coordination of access to the WM is distributed among the wireless stations.

Using the various methods covered in this chapter, all the stations work together to provide cooperative access to the WM without the need for a centralized medium access controller. However, DCF is not the only RF MAC method. PCF and HCF are both valid access methods as well.

PCF

The *Point Coordination Function* has not been covered in this chapter. In this section, we will provide a very brief overview. Though PCF is defined in the IEEE 802.11 standard, it has not been implemented in any widely used devices.

There are also no known plans for its implementation at this time. Many wireless vendors speak of PCF in their literature and then state that it is not implemented due to lack of industry support or because of extra overhead incurred by implementing it.

Unlike DCF, PCF centralizes access to the WM. There is one *point* (station) in the WLAN that is responsible for controlling access to the WM. This point is the access point. In order to implement PCF, you will need both an access point and a client station that support the PCF specifications within IEEE 802.11.

When an access point is configured to use PCF, it will actually use both DCF and PCF. This will be accomplished by alternating between a *contention period (CP)* and a *contention free period (CFP).*

The CP is the window of time when DCF (CSMA/CA) is used to control access to the WM, and the CFP is the window of time when both DCF and PCF are used. Technically, there is more involved than just switching between CFP and CP windows.

The access point goes through a CFP *repetition interval* process that involves the following cycle:

1. The access point waits for the duration of a PIFS.

2. The access point sends a beacon frame announcing the CFP is about to begin. The CFP window begins and the access point polls each station that indicated a desire to participate in the CFP.

3. The CP window begins and the stations contend for access to the WM using normal DCF rules.

This process begins again at step 1 and continues to repeat itself. The interesting thing to note is that PCF provides an apparent benefit that, up to this time, has not really been needed.

The reason is that, though DCF uses random backoff timers that have a probability of giving one station more access to the WM than another during a given window of time, there seems to be a good balance among the stations that access the WM using DCF.

There is a new use of WLANs that has shown some real problems with DCF, and it is unknown whether PCF could help solve these problems, since it is not available. (PCF could even help with the hidden node problem if it were available.)
The new use is Voice over WLAN or VoWLAN. As a solution to this problem and others, many vendors are beginning to turn to IEEE 802.11e as a potential solution.

IEEE 802.11e and WMM

Many networking technologies require very low latency. In fact, latency issues have even been a problem in some wired networks.

The holy grail of networking today is convergence: voice and data on the same medium. One way to provide lower latency is to dedicate a medium to a single pair of devices; however, this is cost prohibitive.

The alternative is to somehow identify the higher priority information and make sure that information gets preferential access to the medium. This is the heart of QoS.

IEEE has released a solution to the QoS problem in the form of the IEEE 802.11e amendment.

IEEE 802.11e specifies the use of EDCA and HCF. Two new station types are introduced by IEEE 802.11e: QoS access points (QoS AP) and QoS stations (QoS STA).

A QoS AP is an access point that can support the QoS facility. A QoS STA is a station that supports the QoS facility and can act as a standard station when associated with a non QoS AP.

The QoS facility is inclusive of the following components that distinguish a QoS STA from a non−QoS STA:

- QoS functions
- Channel access rules
- Frame formats and frame exchanges
- Managed objects

EDCA
EDCA is the IEEE 802.11e enhancement of DCF. Eight traffic categories, or priority levels, are defined by EDCA. The traffic having the higher priority level will gain access to the WM before traffic having a lower priority level.

Ultimately, EDCA does not provide a guarantee of access to the WM; however, it does increase the probability over DCF that a higher priority frame will be transmitted before a lower priority frame.

These eight traffic categories are defined by the User Priority (UP) value. This value can be from 0 to 7. The UP values are identical to those used in 802.1D. Clause 6 of the IEEE 802.11e amendment further explains these UP values and their interpretation.

HCF
HCF provides a preemptive capability to the QAP (QoS AP) that was not available to an access point with PCF. A PCF access point, if it were available, would have the ability to preempt other stations in the BSS during the contention free period; however, it could not preempt other stations during the contention period. HCF adds this capability.

This preemption should not be thought of as interrupting a station's frame transmittal, but rather ensuring that the QAP will be able to transmit on the WM next.

Wireless Multimedia (WMM)
While IEEE 802.11e was being developed, the WiFi Alliance released their Wireless Multimedia (WMM) extensions certification.

This certification is based on the draft IEEE 802.11e standard and was released to provide QoS for Voice over WLAN.

The WMM certification will continue to be updated and redefined to mean the latest interoperable QoS features available from multiple chip vendors.

RTS/CTS and CTS to Self Protocols

DCF provides a CSMA/CA implementation for WLANs using distributed coordination. PCF could have provided CSMA/CA through centralized or point coordination. Sometimes, you need something different from what is offered by either DCF or PCF.

Instead of the access point polling the stations to see which station needs to communicate, the stations can tell the access point they need to communicate and then wait for the access point to give them the go ahead.

This method is called *Request to Send/Clear to Send (RTS/CTS)*.

A scenario can happen on a WLAN when the hidden node problem occurs. In this situation, there are two or more clients that can hear the access point and be heard by the access point, but they can't hear each other.

Therefore, when a frame is sent from one of the client stations (STA1) to the access point, the other client station (STA2) might not be able to sense that it is transmitting using physical sensing.

This results in STA2 transmitting a frame at the same time, causing corruption or cancellation of the other station's frame. It's as if the frames reached the access point and were told, "No vacancy."

RTS/CTS is like calling ahead and making reservations. Like calling ahead to make reservations, it requires extra overhead every time. If you are having problems like hidden nodes, enabling RTS/CTS can help resolve them.

If you are not, then "calling ahead" will only add unnecessary overhead to your WLAN. RTS/CTS works according to the following process:

1. A station wishing to transmit using RTS/CTS sends a *Request to Send frame* to the access point.

2. When the access point receives the RTS request, it sends a *Clear to Send frame* to the WLAN as a broadcast.

3. The stations in the vicinity all hear the duration in either the Request to Send frame or the Clear to Send frame and know to stay silent.

4. The original requesting station transmits its frame and receives an acknowledgment during this quiet window.

RTS/CTS can function in an infrastructure BSS or an independent basic service set (IBSS). In the BSS, the RTS/CTS exchange is between the client stations that wish to send or receive data and the access point, and either may initiate the exchange. In the IBSS, the RTS/CTS exchange is between the two communicating client stations.

The non involved stations hear the exchange and set their N AV timers to cooperate with the RTS/CTS process.

An additional implementation of Clear to Send is found in the IEEE 802.11g amendment for the ERP PHY. This implementation provides for a CTS to self. Essentially, the station using the ERP PHY can communicate using OFDM and faster data rates than older stations such as those using the HR/DSSS PHY.

In order for these stations to coexist, the station with the ERP PHY will transmit a CTS frame that was not preceded by an RTS frame. This frame will be transmitted using modulation that can be understood by the stations with the non−ERP PHYs.

Those stations will go silent as they honor the duration value in the CTS frame. During this silent period, the ERP based station will transmit its OFDM modulated signal without further concern for the non−ERP PHYs.

Fragmentation
As you read earlier in this chapter, IEEE 802.11 frames can support an MSDU of 2304 bytes.

However, because TCP/IP uses an MTU of 1500 bytes as a default and most data is transmitted using TCP/IP, most WLAN frames are considerably smaller than the maximum potential frame size.

There are situations where these frames of roughly 1500 bytes (1508 after the LLC) may still be too large. The IEEE 802.11 MAC layer has the functionality to fragment the data even more.

This is done through a process known as *fragmentation.*

You can configure a setting in your wireless client software called the *fragmentation threshold.* This is a setting, in bytes, that determines when the MAC layer will fragment the frames into smaller frames.

A higher fragmentation threshold usually means less fragmentation protocol overhead but less resilience against interference. A lower fragmentation threshold usually means more fragmentation protocol overheads but more resilience against interference.

If at any moment there is no interference to cause retransmission, the additional fragmentation protocol overhead reduces data throughput.

If at the next moment there is interference causing retransmission, the additional fragmentation protocol overheads increases data throughput above what it would have been without the fragmentation.

Your job, as a WLAN administrator, will require you to find the optimum threshold for a given connection. In the real world, default fragmentation thresholds are normally left as is (no fragmentation) and they are only changed when there is a communications problem.

This is due to the fact that a wireless client may perform better with a lower fragmentation threshold in one area and a higher fragmentation threshold in another area.

Dynamic Rate Switching
Dynamic rate selection, dynamic rate switching, automatic rate shifting, and dynamic rate shifting all refer to IEEE 802.11 Section 9.6 Multirate support, but whatever you call it, it is the process of reducing or increasing the data rate to the next supported data rate as the quality of the RF signal changes.

Remember that signal strength attenuates over distance. This results in a weaker signal at a longer distance than is available at a shorter distance. Other factors, such as absorption into materials in the service area, can also result in a weaker signal at a point equidistant from the access point as another point with a stronger signal.

Whatever the reason for reduced signal quality, the data rate is lowered to provide more effective use of the WM.

Consider that modulation schemes used in the DSSS PHY, for example, change fewer attributes of the RF signal fewer times in order to modulate data onto the signal than do the modulation schemes used in the OFDM or ERP PHYs.

As the quality of the signal degrades, it becomes more and more difficult to demodulate the more complex modulation schemes. By slowing down the data rate, either with a different or the same modulation, it becomes easier to demodulate the data.

A standards based device will only change the data rate to a supported data rate of the standard. For example, a HR/DSSS PHY will shift from 11 to 5.5 Mbps, but will not shift from 11 to 6 Mbps because 6 Mbps is not supported by the HR/DSSS PHY.

In the same way, an ERP PHY will shift from 48 to 54 Mbps, but it will never shift from 48 to 51 Mbps, since 51 Mbps is not a supported data rate.

Summary

This chapter provided an overview of terms related to the IEEE 802.11 MAC and PHY and then presented the MAC layer services in detail. You learned about frames and frame exchanges as well as the process used to go from MSDU to PPDU for transmission on the WM. Finally, you learned about the different methods used to control access to the WM, including DCF, PCF, and RTS/CTS.

Review Questions

1. The IEEE 802.11 state machine supports three states. These three states are different combinations of which two functions?

A. Authentication
B. Accounting
C. Association
D. Interframe spacing

2. The time interval that a station must wait before beginning a random backoff time, when using DCF, is called what?

A. Data rate
B. Interframe spacing
C. Country code
D. Gap window

3. When there is interference in an area, which of the following settings may improve performance if it is modified?

A. Fragmentation threshold
B. DIFS
C. Carrier sense
D. Contention window

4. The unit of data being manipulated is the PSDU. It is being converted into a PPDU. In which layer of the OSI model is this data unit currently?

A. Data Link layer
B. Application layer
C. Physical layer
D. Transport layer

5. What is the name of the sublayer that is actually responsible for modulating data onto the WM?

A. PLCP
B. MAC
C. LLC
D. PMD

6. Which of the following amendments introduce Country elements into the IEEE 802.11 standard? (Choose all that apply.)

A. 802.11e
B. 802.11d
C. 802.11j
D. 802.11i

Answers

1. A, C. The correct answers are authentication and association. The three states are unauthenticated and unassociated, authenticated and unassociated, and authenticated and associated.

2. B. The correct answer is interframe spacing (IFS). There are multiple IFS intervals, including SIFS (shortest), PIFS (neither shortest nor longest), and DIFS (longer than PIFS).

3. A. The correct answer is fragmentation threshold. By decreasing the fragmentation threshold, you reduce the amount of time it takes to transfer a frame and, therefore, increase the likelihood that the frame will be transmitted before interference occurs.

4. C. The correct answer is Physical layer. The Physical layer is where the PSDU becomes a PPDU for transmission on the WM.

5. D. The correct answer is PMD. The PLCP is the communications layer between the MAC and the PMD that allows for the use of a similar MAC layer with all of the PHYs in IEEE 802.11 as amended. The MAC layer is where the upper–layer data units are converted into IEEE 802.11 frames.

6. B, C. The correct answers are 802.11d and 802.11j. In addition to these two, 802.11h introduces a Country entity as well.